

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Extremely Small and Incredibly Fast: Combining Spectroscopy and Microscopy to Reveal Local Excited State Dynamics in Disordered Semiconductors

Permalink

<https://escholarship.org/uc/item/98z0c79v>

Author

Folie, Brendan

Publication Date

2018

Peer reviewed|Thesis/dissertation

**Extremely Small and Incredibly Fast: Combining Spectroscopy and Microscopy
to Reveal Local Excited State Dynamics in Disordered Semiconductors**

by

Brendan Folie

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Naomi Ginsberg, Chair

Professor Feng Wang

Professor Phillip Geissler

Fall 2018

**Extremely Small and Incredibly Fast: Combining Spectroscopy and Microscopy
to Reveal Local Excited State Dynamics in Disordered Semiconductors**

Copyright 2018
by
Brendan Folie

Abstract

Extremely Small and Incredibly Fast: Combining Spectroscopy and Microscopy to Reveal Local Excited State Dynamics in Disordered Semiconductors

by

Brendan Folie

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Naomi Ginsberg, Chair

Despite years of intense study, even today there are many novel semiconductor materials whose unique properties raise fundamental scientific questions about the relationship between structure and function. Two examples are organic semiconductors and semiconducting nanostructures. Organic semiconductors can be solution processed at room temperature, leading to thin films that are flexible, inexpensive, and more sustainable when compared to inorganic alternatives. Yet solution processing leads to complex microstructures, and the resulting polymorphism, interfaces, and defects all affect the material's behavior in complex ways. In semiconducting nanostructures, the novel behavior arises due to their nanometer-scale dimensions. Confinement and enhanced surface interactions affect the electronic structure and dynamics in sometimes unexpected ways and can produce new physics.

Chapters 1 and 2 provide the necessary background for this dissertation. Chapter 1 provides an overview of semiconductors and what makes organic semiconductors distinct. The dynamic processes that are relevant for subsequent chapters are introduced. Chapter 2 describes a home-built transient absorption (TA) microscope that was used in each project discussed here. TA is an important tool for studying a material's ultrafast excited state dynamics, and by combining TA with microscopy we have been able to investigate how those dynamics depend on the local morphology.

Chapters 3 and 4 discuss the use of TA microscopy to probe the ultrafast electronic dynamics of individual crystalline domains of organic semiconductors. In Chapter 3 we find that different domains of the material diF-TES-ADT can display different dynamics. We fit a kinetic model to the observed behaviors, quantify the amount of heterogeneity, and propose that the observations are due to polymorphism. In Chapter 4 we study singlet fission in single crystalline domains of the material TIPS-Pentacene. During singlet fission, a singlet exciton splits into two triplet excitons, initially forming a short-lived and poorly understood correlated triplet pair. By exploiting the inherent anisotropy within an individual crystalline domain we obtain unprecedented insight into nature of the correlated triplet pair. We resolve a disagreement related to its formation timescale, quantify the triplet-triplet binding energy,

and measure how the charge transfer character of the correlated triplet pair perturbs its electronic structure and hence absorption spectrum.

Chapter 5 takes a step back to study the self-assembly process in the organic semiconductor rubrene. We use a variety of spatially- and temporally-resolved techniques to study the growth and morphology of spherulites, a kinetically trapped polycrystalline structure that could allow for both rapid and uniform charge transport. We use time-resolved wide-angle *in situ* X-ray scattering, wide-angle X-ray microdiffraction, and atomic force microscopy to identify shear strain localized along lines of crystalline misorientation. This strain templates upon nucleation and is kinetically trapped only at higher annealing temperatures, possibly promoting the spherulite structure. Steps to extend this project are also described, using scanning transmission electron microscopy to measure the crystal structure with nanoscale resolution and TA microscopy to connect this local structure with electronic dynamics.

Chapter 6 looks at nanowires of the inorganic perovskite CsPbBr_3 , in which quantum confinement along two dimensions changes the excited-state dynamics. We align many nanowires into micron-scale bundles, which allows them to be studied with optical techniques while still retaining anisotropic behavior. We use TA microscopy to study the polarization-resolved ultrafast electronic dynamics of the nanowires, and find among other things that the nanowire geometry splits the degeneracy of the lowest excited states. We also use stroboscopic interferometric scattering to observed exciton diffusion within nanowire bundles, and measure the diffusivity both along and between the nanowires.

Taken together, this dissertation illustrates specific examples of the relationship between semiconductor structure and function. Small changes in processing conditions can drastically affect morphology, and small changes in morphology, whether they be defects, interfaces, or tweaks to the unit cell, can drastically affect electronic properties. Understanding these relationships is vital to incorporating novel semiconducting materials into next-generation optoelectronic devices.

To Frances Mayer, the best teacher I ever had

Contents

Contents	ii
List of Figures	iv
List of Tables	xii
1 Introduction	1
1.1 Description of Semiconductors	1
1.2 Organic Semiconductors Defy Common Approximations	7
1.3 Semiconductor Devices	11
1.4 Ultrafast Dynamic Processes In Semiconductors	14
1.5 Structure of Subsequent Chapters	19
2 Transient Absorption Microscopy (TAM)	20
2.1 Transient Absorption is a Multipurpose Tool	20
2.2 Combining Transient Absorption with Microscopy	27
2.3 Ultraslow Evolution of Transient Absorption Microscope	30
3 TAM Discerns Inter-Domain Heterogeneity in diF-TES-ADT	34
3.1 Characteristics of diF-TES-ADT films	34
3.2 Measuring Heterogeneity with Single-Domain TAM Data	35
3.3 Physical Origins of Heterogeneity	45
3.4 Future Directions	47
4 TAM Probes the Correlated Triplet Pair in Single Domains of TIPS-Pentacene	49
4.1 Singlet Fission and the Correlated Triplet Pair	50
4.2 Kinetic Model of TAM Data	53
4.3 Triplet Pair Separation and Binding Energy	73
4.4 Effects of Charge Transfer Character on the Electronic Structure of Bound Triplets	76
4.5 Implications for Devices	83

5	TAM (and friends) Study Self-Assembly of Rubrene Spherulites	85
5.1	Fabrication and Basic Characterization of Rubrene Spherulites	86
5.2	In-Situ X-Ray Scattering of Self-Assembling Rubrene Spherulites	100
5.3	Microdiffraction—Measuring Spherulite Crystal Structure with Fine Spatial Resolution	113
5.4	Local Structure and Dynamics of Rubrene Spherulites Probed Independently Via Several Imaging Techniques	116
5.5	Rubrene Spherulite Growth is Potentially Driven by Kinetically Trapped Strain	120
5.6	Putting It All Together: Steps Toward Multimodal Imaging	136
6	TAM (and friends) Investigate Quantum Confinement and Dimensionality Effects in Inorganic Perovskite Nanowires	138
6.1	An Introduction to Lead Halide Perovskite Semiconductors	139
6.2	Properties and Deposition of CsPbBr ₃ Nanowire Bundles	141
6.3	Effects of Shape Anisotropy on Diffusion in Nanowires	147
6.4	Effects of Anisotropy on Electronic Structure and Ultrafast Dynamics in Nanowires	155
6.5	Investigating Fluorescence Anisotropy in CsPbBr ₃ Nanowires	164
6.6	Future Directions	171
7	Concluding Remarks	173
	Bibliography	175
A	Alignment Procedure for Transient Absorption Microscope	197
B	Laview Program for Controlling Transient Absorption Microscope	205
C	Code for Simulating Rubrene Crystallization	217
C.1	Main Simulation Code	217
C.2	Helper Functions	228
D	Code for Identifying Rubrene Morphology	234
D.1	Jupyter Notebooks	234
D.2	Helper Functions	239

List of Figures

1.1	Illustration of doping in semiconductors. Unit cell of (a) undoped silicon, (b) silicon doped with phosphorus to add a negative charge carrier, and (c) silicon doped with boron to add a positive charge carrier.	2
1.2	Illustration of band structure formation. Source: [5]	3
1.3	Cartoon illustration of band structure. The band gap is the energy difference between the top of the valence bands and the bottom of the conduction bands.	4
1.4	Using a parabolic approximation (dashed red line) to the dispersion (black lines) in order to calculate effective mass of a charge carrier quasi-particle.	5
1.5	Cartoon of (a) Wannier-Mott and (b) Frenkel excitons, showing the real space representation (top) and band structure (bottom). Wannier-Mott excitons are larger and weakly bound—the energy levels (orange-red lines) are only slightly below the conduction band. Frenkel excitons are smaller and strongly bound.	6
1.6	Representative examples of the two types of organic semiconductors: (a) small molecule and (b) polymer. The materials are (a) TIPS-Pentacene and (b) MEH-PPV [9].	8
1.7	Bonding orbitals in polyacene molecules. (a) Pentacene consists of five acene rings. The in-plane sp^2 orbitals are shown in red. (b) A side-on view showing the out-of-plane p orbitals in green. (c) A crystallized polyacene derivative. The π orbitals on neighboring molecules overlap, leading to strong interactions.	10
1.8	A prototypical photovoltaic cell schematic. An incident photon (red) creates an electron and hole that are separated at an interface. Electrodes collect the charge carriers and create current.	12
1.9	Diagram of a top-gate field effect transistor.	13
1.10	Jablonski diagram of the various states (thick lines) and processes (arrows) important for this thesis. The vertical position of each state denotes its energy. See text for definitions.	15
1.11	Nonlinear relaxation processes in inorganic semiconductors . Filled/open circles are electrons/holes, arrows are transitions, and zigzag lines are phonons. (a) Biexciton recombination—one exciton relaxes due to interaction with a second exciton. In this example a phonon is required to conserve momentum. (b) Auger recombination involving two electrons and one hole.	18

2.1	Diagram of a transient absorption setup. Ultrafast pump and probe pulses are focused and overlapped with time delay τ in the sample. The pump is blocked and the probe is sent into a detector. Chopping the pump allows one to measure the differential change in probe transmission due to excitation.	21
2.2	Origin of ground state bleach	22
2.3	Diagram of stimulated emission	23
2.4	Diagram of excited state absorption	23
2.5	Hypothetical (a) TA decay data and (b) fit.	24
2.6	Example TA spectra, illustrating how time-zero (black line) can depend on probe wavelength.	26
2.7	Illustration of how polarization-resolved TA takes advantage of structural anisotropy in a material. The green structure approximates the shape of a single acene molecule. In its crystalline form all nearby molecules are oriented in the same way. Rotating the probe polarization (orange arrows) allows for selective coupling to different transition dipole moments (red/blue arrows).	28
2.8	Iterations of our group's TA Microscope: (a) single-color, (b) two-color, and (c) broadband.	32
3.1	(a) Structure and (b) absorption spectrum (black) of diF-TES-ADT thin film. Red curve in (b) is the excitation laser spectrum.	35
3.2	Optical images of diF-TES-ADT dropcast films. Fairly uniform single crystalline domains are visible in (b).	35
3.3	Polarization-resolved TAM data of three different diF-TES-ADT domains (labeled A-C), visualized in two ways. (a)-(c) ΔT vs. time where each series corresponds to a fixed probe polarization. Lines are global fits to a sum of three exponentials. Gray dots are polarization-averaged data, which themselves are fit to a sum of three exponentials. (d)-(f) ΔT vs. probe polarization for a fixed delay time of 1 ps. Lines are fits to a cosine squared plus offset.	36
3.4	Polarization dependence of global fit parameters A_1 , A_2 , A_3 , and C for Domain A.	39
3.5	Histograms of time constants τ_1 (a,c) and τ_2 (b,d) from global fits of TA signal collected from individual domains of diF-TES-ADT (a,b) and TIPS-Pn (c,d). Time constants are reported as fractions of their average values.	40
3.6	Histograms of normalized offset in single domains within thin films of (a) diF-TES-ADT and (b) TIPS-Pn. Bin size is 0.05.	40
3.7	Bulk TA on diF-TES-ADT for two selected probe wavelengths, 550 nm (blue) and 671 nm (red). Pump fluences are comparable to those used in TAM. TA was measured at 20 kHz using a Light Conversion PHAROS regenerative amplifier pumping a non-collinear parametric amplifier (NOPA, also Light Conversion). The sample was translated horizontally after each scan, cycling between 20 discrete points spaced by 0.1 mm.	42
3.8	TAM from three domains of diF-TES-ADT (same data as in Figure 3.3), along with the kinetic model fit (solid lines).	45

3.9	Schematic illustration of possible variations in structure and non-azimuthal orientation of small molecule crystal domains. (a) Structural variations, as inferred for diF-TES-ADT, involve changing the intermolecular spacing and angle with respect to the substrate. (b) TIPS-Pn appears to maintain a uniform intermolecular spacing and angle with respect to the substrate. (c) The non-azimuthal orientation of diF-TES-ADT crystal grains appears to be variable even if independent of lattice spacing.	47
4.1	Structure of TIPS-Pentacene. (a) Molecular structure [9], (b) optical image of a crystalline thin film with an 8 μm laser spot marked in red and a uniform area inside the blue dashed lines, and crystalline packing structure looking (c) top-down and (d) side-on. Crystal axes are indicated, as are TDMs along the short (red) and long (blue) axis the pentacene core.	52
4.2	Absorption spectrum of TIPS-Pn (dark blue), pump laser spectrum (red), and non-degenerate probe wavelengths used in this study (vertical bars).	53
4.3	Degenerate polarization-resolved TA of a single crystalline domain of TIPS-Pn, represented by (a) fixing probe polarization and plotting the signal over time, and (b) fixing the delay time and plotting the signal vs. polarization. Curves are fits to the kinetic model. The polarization-resolved transmission is also shown, in red.	54
4.4	Pump power dependence of TA signal at several delay times for (a) Domain 1 and (c) Domain 3. Pump powers are chosen so that we remain in the linear regime, even at long delay times.	55
4.5	Explanation and example of normalized offset, ζ . (a) Cartoon depicting how ζ is calculated. (b) Measured ζ in TIPS-Pn for four data sets, two with a degenerate 700 nm probe and two with a 694 nm probe. The dashed line is a guide to indicate the plateau of ζ at long times.	56
4.6	TA signal vs. probe polarization at $\tau = 1.4$ ns and $\lambda = 694$ nm, fit to a sum of two TDMs (Equation 4.2).	57
4.7	Attempted fit of TA data to an ultrafast singlet fission model, fixing the relative strength and orientation of the triplet ESA TDM. The model is incapable of fitting the data.	59
4.8	A Jablonski diagram of the kinetic model for TIPS-Pn	61
4.9	Polarization-resolved TAM and fits with selected non-degenerate probes	63
4.10	Polarization-resolved TAM and fits with all non-degenerate probes	64
4.11	Long time TA signal and fits for two selected probe wavelengths	65
4.12	Degenerate TA data and kinetic model fits for Domains 2 (a) and 3 (b)	67
4.13	Transition dipole moments of TIPS-Pn: (a) absorbance (for unpolarized light) and (b) orientation. Dashed line in (b) is a guide to the eye. Orientations are relative to the absorption dipole at 700 nm.	69
4.14	Idealized probe transmission vs. polarization	70

4.15	Comparison of $\Delta T/T$ (blue) with ΔT (orange) showing the simplicity of the latter. $\Delta T/T$ does not follow a \cos^2 shape (dotted line). ΔT does fit well to \cos^2 (dashed line), and including the shape correction factor (solid line, Equation 4.14) leads to a slight improvement	73
4.16	Diagram of the hopping model used to estimate triplet binding energy.	74
4.17	A visualization of the lowest energy triplet exciton wave function in TIPS-Pn, calculated by solving the BSE.	75
4.18	Contour plot of τ_{dis} for various binding energy landscapes. Vertical bars indicate the range of binding energies that can produce dissociation timescales within our measured range of 200-400 ps.	76
4.19	Energy levels for low-lying singlet and triplet excitons, computed with BerkeleyGW.	80
4.20	Comparison of the polarization-averaged linear absorption spectrum calculated with BerkeleyGW and measured on a disordered polycrystalline film of TIPS-Pn.	81
4.21	Absorption spectra of excited singlet and triplet states S_1 and T_1 computed with BerkeleyGW.	81
4.22	Coefficient of the uncertain-sign term in Equation 4.27, for unpolarized light incident along the z -axis, and broadening $\sigma = 100$ meV.	82
5.1	Structure of rubrene: (a) single molecule [9] and (b) orthorhombic crystal looking from the side and (c) top-down.	87
5.2	Structure of (a) generic spherulite and (b) rubrene spherulite.	88
5.3	Cartoon schematic of crossed polarizers.	88
5.4	Rubrene spherulites viewed through crossed polarizers in (a) our microscope, (b) our collaborator's microscope in reflection mode, (c) and in transmission mode (courtesy of Jordan Dull, Princeton University).	89
5.5	Attempted rubrene spherulites made with solution processing techniques.	90
5.6	Rubrene spherulite made with thermal evaporation.	92
5.7	Picture of thermal annealing setup.	93
5.8	Rubrene spherulites evaporated and annealed at Berkeley (170°C), using purified rubrene.	93
5.9	Rubrene spherulites evaporated at Princeton and annealed with a range of temperatures. Nuclei, spherulites, and platelets are noted. Discontinuities are places where the crystalline orientation abruptly changes.	94
5.10	Absorption and photoluminescence of (a) rubrene and (b) rubrene endoperoxide.	95
5.11	Photoluminescence (a) image and (b) spectra of a proper rubrene spherulite. 405 nm excitation light preferentially couples to oxidized rubrene.	95
5.12	Photoluminescence (a) image and (b) spectra of a thermally evaporated rubrene spherulite excited at 473 nm.	96
5.13	Photoluminescence (a) image and (b) spectra of a solvent vapor annealed rubrene spherulite.	97
5.14	AFM image of TPTPA on substrates: (a) ITO and (b) bare glass.	97

5.15	AFM image of (a) rubrene spherulite showing growth direction, (b) rubrene platelet showing interfaces (red arrows), (c) spherulite edge showing lip, and (d) linecut across spherulite edge.	98
5.16	AFM images of a (a) rubrene spherulite and (b) a discontinuity emanating from the nucleus. (c) Polarized optical microscopy image of the spherulite.	99
5.17	AFM images of (a) small and (b) large rubrene platelet nuclei. (c) Polarized optical microscopy image of the structure shown in panel b.	100
5.18	Illustration of grazing-incidence wide-angle X-ray scattering (GIWAXS)	101
5.19	GIWAXS pattern of oriented orthorhombic rubrene	103
5.20	GIWAXS patterns of disordered (a) orthorhombic and (b) triclinic rubrene	104
5.21	Temperature rise of rapid heater stage, measured with the stage controller.	105
5.22	Rubrene spherulites annealed <i>in situ</i> at the GIWAXS beamline. Optical cross-polarized transmission Images taken above ITO for temperatures (a) 160 °C, (b) 165 °C, (c) 170 °C, (d) 175 °C, (e) 180 °C, and (f) above glass at 170 °C.	106
5.23	GIWAXS pattern of rubrene spherulites annealed <i>in situ</i> at 170 °C (5 second exposure).	107
5.24	Time-resolved GIWAXS of rubrene spherulites during <i>in situ</i> annealing at 170.	107
5.25	(a) GIWAXS image of (002) peak of rubrene spherulite film after being annealed <i>in situ</i> at 170 °C. Slices of the peak along the (b) q_z and (c) q_y direction.	108
5.26	Time-resolved slices of the (002) peak of rubrene spherulites annealed <i>in situ</i> at 170 °C.	108
5.27	(a) Linecuts of the (002) diffraction peak for different values of the sample orientation angle, α . (b) Schematic showing incoming and outgoing X-rays, sample tilt, and crystalline orientation.	110
5.28	Linecuts of (002) peak of rubrene film for several values of incident angle, α . To the left of the jagged line the horizontal axis is stretched for visibility. The dashed line corresponds to studying vertically-oriented crystallites.	111
5.29	One possible type of shear strain on rubrene's crystal structure and how it would affect both the real and reciprocal space lattices.	112
5.30	(a) Microdiffraction image of (002) peak of rubrene spherulite annealed at 170 °C. Linecuts of the peak along the (b) q_z and (c) q_y direction.	113
5.31	Microdiffraction on a rubrene spherulite. The beam was rastered from left to right, and the (002) peak at each point used to produce maps of the q_y and q_z slices. The q_z slice is consistent, but the q_y slice varies around the discontinuity, indicating shear strain.	114
5.32	Microdiffraction on a rubrene platelet. The beam was rastered from left to right, and the (002) peak at each point used to produce maps of the q_y and q_z (not shown) slices. The peak does not vary as a function of spatial position.	115
5.33	TAM spectra of a rubrene spherulite at three time delays. Figure courtesy Jenna Tan.	117

5.34	Infrared scanning-scattering near-field optical microscopy with synchrotron source. Infrared spectra (amplitude and phase) are shown after averaging over two crystalline regions of a rubrene film on gold. Maps of the two regions are shown on the right, after summing the signal amplitude over all frequencies. Figure courtesy of Omar Khatib, Raschke lab, CU Boulder.	118
5.35	Diffraction pattern at one point on a rubrene spherulite as measured with TEM. Figure courtesy of Steven Zeltmann, Minor lab, UC Berkeley.	119
5.36	Strain map of rubrene spherulite, calculated by measuring the shift in several diffraction peaks. An example diffraction image is shown in the upper-left. Figure courtesy Steven Zeltmann.	119
5.37	The pixel interfacial energy as a function of misorientation, $\Delta\theta$, for $\epsilon_2 = 0$ (solid line) and $\epsilon_2 > 0$ (dashed line).	125
5.38	Example energy profiles created by (a) two adjacent interfaces, as a function of the orientation of the middle pixel. The orange and blue dashed lines are the energies due to the two interfaces; the solid green line is the total energy. In all plots, $E_0 = -50$ and $\theta_{crit} = 0.5$. Values of ϵ_2 and θ_{offset} are: (b) 0, 0.2, (c) 0, 0.3, (d) 100, 0.2, (e) 100, 0.29, (f) 100, 0.35.	126
5.39	Snapshots of simulated rubrene crystallization.	128
5.40	Results of simulated rubrene crystallization for different values of critical angle.	129
5.41	Results of simulated rubrene crystallization for different values of temperature.	129
5.42	Example diagram of a single tree in a random forest classifier.	130
5.43	Training data used for random forest classifier.	131
5.44	Finding regions in an image of crystalline rubrene. (a) Regions one might identify by eye, (b) regions found with agglomerative clustering, (c) those regions overlaid on the image.	132
5.45	Testing the accuracy of a random forest classifier. (a) Confusion matrix and (b) accuracy bar chart for $d = 4$ and $n_{trees} = 80$	133
5.46	Random forest classifier applied to training images.	134
5.47	Results of 200 random forest classifier models built on subsets of the training data. (a) Histogram of the results and (b) scatter plot of the predicted amount of spherulite vs. platelet.	135
5.48	Ratio of spherulite to platelet coverage for the ITO region of rubrene films annealed <i>in situ</i> at temperatures from 160 to 180 °C.	136
6.1	CsPbBr ₃ orthorhombic unit cell.	141
6.2	Cartoon of CsPbBr ₃ nanowire bundle. Yellow lines represent ligands.	142
6.3	TEM images of (a) 10 nm and (b) 3 nm CsPbBr ₃ nanowire bundles. Courtesy of Jianmei Huang.	143
6.4	Optical transmission images of (a) 10 nm and (b) 3 nm CsPbBr ₃ nanowire bundles deposited on glass. Several distinct bundles are circled.	144
6.5	Absorption (solid lines) and photoluminescence (dashed lines) of 10 nm (orange) and 3 nm (blue) nanowire bundles.	144

6.6	Normalized time-resolved photoluminescence of 10 nm and 3 nm nanowire bundles in solution.	145
6.7	Stability of CsPbBr ₃ 10 nm nanowire bundles. The PL is stronger for thicker bundles, but does not diminish over the course of several weeks.	146
6.8	Optical image of a thick CsPbBr ₃ nanowire.	146
6.9	stroboSCAT on a CsPbBr ₃ 10 nm nanowire bundle. (a) iSCAT image of bundle and (b)-(d) stroboSCAT images for three time delays.	148
6.10	Longitudinal diffusion in CsPbBr ₃ 10 nm nanowire bundles. (a) Selected line cuts along the long axis of a bundle. (b) Width of the excitation spot, fit to an exponentially decaying diffusivity.	148
6.11	Sanity checks related to exponentially decaying diffusivity. (a) Attempted (failed) fit of longitudinal diffusion to a power law. (b) Diffusion decay time for a range of pump fluences, showing no trend. (c) Diffusion in a thick nanowire.	149
6.12	Transverse line cuts showing inter-wire diffusion.	150
6.13	Trap-limited decay model in nanowires (Equations 6.7 and 6.9). (a) Distribution at a fixed time fit to a Gaussian and (b) Evolution of σ^2 fit to a decaying exponential.	152
6.14	Mapping simulation parameters to those returned by the fit. (a) Fit vs. actual diffusivity, and (b) a contour plot of diffusion decay rate, $1/\tau_d$, vs. trapping rate and diffusivity.	153
6.15	Box plots of (a) diffusivity and (b) inverse trapping rate for eleven CsPbBr ₃ 10 nm nanowire bundles.	154
6.16	Calculating trap density from trapping rate. (a) Simulation results for trapping time vs. trap density, showing the linear fit with slope -1 on a log-log plot. (b) Box plot of inverse linear trap density.	154
6.17	(a) Polarized absorption spectra of a CsPbBr ₃ nanowire bundle. 0° corresponds to the longitudinal axis. (b) Wavelength at half-max plotted vs. polarization, showing the oscillatory trend for nanowires (red). The same measurement on a thick nanowire (black) shows no polarization-dependence. Figure courtesy of Jenna Tan.	156
6.18	Transient absorption spectroscopy of 10 nm CsPbBr ₃ isolated nanowires in cyclohexane.	157
6.19	(a) Evolution Associated Spectra (EAS) of CsPbBr ₃ 10 nm nanowire TA data fit to a four-component sequential model. (b) Normalized EAS.	158
6.20	Evolution associated spectra (EAS) for CsPbBr ₃ 10 nm nanowires in solution, excited with several pump powers.	159
6.21	$1/\Delta OD$ vs. time for the GSB in CsPbBr ₃ 10 nm nanowires. A linear relationship (dashed lines guide the eye) indicate second order dominates initially.	160
6.22	Transient absorption spectra at $\tau = 1$ ps and $\tau = 10$ ps (offset for clarity), comparing solutions of bundled and isolated CsPbBr ₃ 10 nm nanowires.	161
6.23	(a) TA plot of CsPbBr ₃ 3 nm nanowires in solution. (b) Decay of GSB signal, comparing 3 nm and 10 nm nanowires.	162

6.24	Selected normalized TA slices for an individual 10 nm CsPbBr ₃ nanowire bundled probed with polarized light. Rotating the probe polarization from perpendicular to parallel (relative to the longitudinal axis of the nanowires) shifts the TA spectrum about 5 meV.	162
6.25	Normalized TA slices for an individual 3 nm CsPbBr ₃ nanowire bundle probed with polarized light. Rotating the probe polarization from perpendicular to parallel (relative to the longitudinal axis of the nanowires) shifts the TA spectrum about 5 meV. Gaussian fits to the GSB peaks are shown in an inset.	163
6.26	Transmission efficiency through the fluorimeter monochromator for horizontally polarized light relative to vertically polarized light. The black dashed line indicates $\eta_H = \eta_V$	165
6.27	Time-resolved polarized PL of 10 nm CsPbBr ₃ nanowires: (a) bundled and (b) isolated.	166
6.28	Fluorescence anisotropy of CsPbBr ₃ nanowires: (a) 10 nm isolated nanowires in cyclohexane excited at 465 nm, (b) 10 nm isolated nanowires in cyclohexane excited at 408 nm, (c) 3 nm bundles in toluene excited at 408 nm. The first letter of the legend indicates excitation polarization (Vertical or Horizontal), and the second letter indicates emission polarization.	167
6.29	Defect formation and aggregation in 3 nm CsPbBr ₃ nanowires. (a) PL spectrum of a fresh sample, that same sample 5 days later (red-shifted), and then after shaking (a bump appears at 520 nm). (b) TRPL on the 5 day-old sample at both 465 nm (blue) and 520 nm (orange), showing long-lived species that emit at 520 nm.	170
A.1	Diagram of TAM experiment	197
A.2	Images of focused and overlapped (a) pump and (b) probe.	201
A.3	Transmission scans to find the sample. (a) Rough scan in X and Y , (b) fine scan in Y	203
A.4	Time-0 scan at $\lambda = 517$ nm.	203
B.1	Left side of “Setup” screen	206
B.2	Right side of “Setup” screen	207
B.3	Left side of “Microscope Control” screen	209
B.4	Right side of “Microscope Control” screen	210
B.5	“Wavelength Control” screen	215

List of Tables

2.1	Specs of TA Microscope stages	33
3.1	Fits of TAM data to three exponential for the three domains shown in Figure 3.3. Results are shown for both a global fit (all polarizations fit simultaneously with one set of time constants) and a fit to polarization-averaged data. Disagreement between domains is indicative of heterogeneity. Error bars are 1σ	38
3.2	Fitting coefficients for the TAM data collected on three diF-TES-ADT domains and shown in Figure 3.3	38
3.3	Time constants from fits of bulk TA data at two probe wavelengths. The most general form of the fit is $A_2e^{-t/\tau_2} + A'_2e^{-t/\tau'_2} + A_3e^{-t/2300\text{ps}} + C$. 1σ confidence intervals are shown. Though the fit parameters are phenomenological, their physical origin can be understood via the kinetic model.	43
3.4	Time constants and singlet-triplet pair energy splitting calculated from fitting domains A, B, and C to the kinetic model.	44
3.5	Transition dipole moments, calculated using the kinetic model, for GSB/SE (A_0), ESA from the hot excited state (A_1^n), thermalized excited state (A_1^0), and from triplets (A_T). Values in parentheses are the phases for TDMs in domain B after being rotated by 60° . This is to make the A_0 TDM line up with those in domains A and C, and facilitate comparison with the other TDMs.	46
4.1	Measured first timescale in TIPS-Pn, arranged in decreasing order of pump photon energy.	60
4.2	Results of kinetic model fit to degenerate probe data (time scales).	67
4.3	Results of kinetic model fit to degenerate probe data (transition dipole moments).	68
4.4	Exciton creation efficiency, ξ , measured for three domains and three probe wavelengths. We do not have reliable data for Domain 2 at 700 nm.	71
5.1	Relative importance of features in random forest classifier. The “full model” uses all training data, whereas the “subset models” each use a subset of the training data, as described later in this section. The local area around each pixel was defined with $d = 4$	133
A.1	Part numbers of important components of TA microscope	204

A.2 Abbreviations used in Figure A.1 204

Acknowledgments

I would like to thank the many talented and generous scientists I have had the pleasure of working with over the last five years, chief among them Naomi Ginsberg. Thank you for your guidance and thorough feedback, as well as for fostering a positive, supportive working environment.

Thank you to the many current and former Ginsberg group members who have made working in the basement so enjoyable: Clarice Aiello, Connor Bischak, Ben Cotts, Lucas Ginsberg, Craig Hetherington, Dave Kaz, Brandon Luo, Dannielle McCarthy, Sam Penwell, Jake Precht, Namu Ramesh, Trevor Roberts, Claire Stachelrodt, Hannah Stern, Rebecca Wai, Hannah Weaver, Hao Wu, and Eniko Zsoldos. A special thanks to Rodrigo Noriega and Cathy Wong for their patient mentoring when I was a novice researcher, to Milan Delor for the many illuminating discussions, and to Jenna Tan for being an eager and excellent student—I am excited to see how you extend this work.

I have been fortunate to work with many wonderful collaborators at Berkeley and beyond. Thank you to Sivan Refaely-Abramson, Jonah Haber, Jeffrey Neaton, Jianmei Huang, Minliang Lai, Mengyu Gao, Peidong Yang, Dan Oron, Eran Rabani, Kranthi Mandadapu, Sasha Efros, Pete Sercel, John Lyons, Noam Bernstein, Steven Zeltmann, Andrew Minor, Omar Khatib, Eric Muller, Markus Raschke, Anna Hailey, Jordan Dull, Barry Rand, Camelia Stan, and Teresa Chen. A special thanks for Chenhui Zhu and Alex Liebman-Pelaez at the Advanced Light Source, for believing in the rubrene self-assembly project and volunteering so much of their time and resources to make it a success.

Thank you to the people who keep Berkeley running smoothly, chief among them Anne Takizawa, Donna Sakima, and Joelle Miles in the physics department, and Carl Lamey, Roy Washington, and Jessie Woodcock in the chemistry department.

Thank you to Bin Liu, Annie Weathers, and Jun Zhang, at New York University, who gave me my first research position.

Thank you to the many incredible professors at Harvey Mudd College, especially Tom Donnelly. Thank you for giving me the chance to work on such interesting problems, for teaching me how to ask the right scientific questions, and for being a kind and generous mentor.

Thank you to my family for all of their love and support. Thank you to my parents for teaching me to love reading, and also teaching me that being a good person is even more important than intellectual accomplishment. Thank you for letting me move across the country when I was just 17. Thank you to my sister and cousins—having you all close by is part of what makes this state feel like home.

Cassie, thank you for sharing your life with me. I love you.

Chapter 1

Introduction

1.1 Description of Semiconductors

Semiconductors are the backbone of modern electronics, and though Wolfgang Pauli once dismissed their behavior as a mere "dirt effect" [1], that behavior has spawned a rich field of physics in which many open questions remain, even after nearly a century of research. In this chapter I will introduce the basic properties and behaviors of semiconductors, paying particular attention to the ways in which their traits can be influenced by local structure and morphology.

1.1.1 Basic Theory

Semiconductors are characterized by a room temperature conductivity much lower than that of a conductor, but much larger than that of an insulator. This intermediate conductivity comes about because the gap in energy between the valence band (the highest energy electrons in the material's ground state) and the conduction band (the lowest energy states in which electrons can move freely) is larger than the thermal energy $k_B T$, but not excessive. The value of this energy gap is generally on the order of hundreds of meV to several eV [2]—large enough so that thermal excitation into the conduction band is rare, but small enough so that electrons can be excited using light in the near-infrared (IR) to near-ultraviolet (UV) regime, or through a modest applied potential.

Under ambient conditions, the fact that the valence-conduction energy gap ΔE is much greater than thermal energy means that relatively few electrons are in a conductive state (the fraction is $e^{-\Delta E/k_B T}$). However the conductivity can be tuned with *doping*: adding small amounts of impurities to the material that change the electronic structure and make it easier for conducting states to be populated [3]. Silicon provides a prototypical example: in a pure silicon crystal each atom has four electrons that participate in bonding (see Figure 1.1a), but some of the silicon atoms can be replaced with phosphorus, which has five electrons in its outer shell. This fifth electron is not strongly bound to a particular atom, and can be more

easily excited into a conduction state (Figure 1.1b). This is called an n-type semiconductor, due to the easy conduction of negative charge.




Figure 1.1: Illustration of doping in semiconductors. Unit cell of (a) undoped silicon, (b) silicon doped with phosphorus to add a negative charge carrier, and (c) silicon doped with boron to add a positive charge carrier.

A second type of doping in silicon can be achieved by substituting boron. Boron has three electrons in its outer shell, so it is not able to bond properly with one of its four silicon neighbors. A nearby electron can fill the gap, but then it will leave its own missing bond behind, which another electron can fill, and so on (Figure 1.1c). In this way the lack of a bond can move through the material as if it were a particle. In fact this lack of a bond is known as a hole, and it behaves as a positive charge carrier. Semiconductors in which positive holes can easily be excited are known as p-type.

I will now discuss why the valence band, conduction band, and gap between them exist in the first place. Consider a single atom with Z electrons. By the Pauli exclusion principle no orbital can be more than singly occupied [4], so the Z lowest-energy orbitals are filled, and the higher energy orbitals are unfilled. The lowest energy transition that can occur is for the electron in the highest occupied orbital to be promoted to the lowest unoccupied orbital. Denote the excited state as $|\psi^*\rangle$ and let it have energy E_0 . Now consider bringing two atoms together so that their atomic orbitals overlap. This overlap leads to an interaction between electrons on the two atoms, which we model simply as a coupling between the excited state of atom 1, $|\psi_1^*\rangle$, and the excited state of atom 2, $|\psi_2^*\rangle$. The Hamiltonian, H , couples them through the off-diagonal matrix element $\langle\psi_1^*|H|\psi_2^*\rangle = \lambda \neq 0$, and the system Hamiltonian is:

$$H = \begin{pmatrix} E_0 & \lambda \\ \lambda & E_0 \end{pmatrix} \quad (1.1)$$

Diagonalizing this Hamiltonian leads to eigenvalues $E = E_0 \pm \lambda$. That is, the one excited state of energy E_0 has split into two excited states with energy splitting 2λ . A similar splitting happens for all of each atom's atomic states. As many atoms are added to the solid, they split into many energy levels that become so close in energy as to appear like one continuous energy band. This process is diagrammed in Figure 1.2 [5]. The filled atomic

orbitals form the valence bands and the unfilled orbitals form the conduction bands. The energy difference between the highest point on the valence band and the lowest point on the conduction band is known as the *band gap*.




Figure 1.2: Illustration of band structure formation. Source: [5]

In an ideal crystal there is translation symmetry, meaning that momentum is a well defined quantum number. The material's band structure can therefore be represented using a diagram like the one in Figure 1.3. Each line corresponds to a distinct energy band, and each point on the line corresponds to a specific electron state. The position on the vertical axis denotes the state's energy and its position on the horizontal axis denotes the state's momentum. The bands on the bottom are valence bands and the bands above them are conduction bands. The band edge transition is denoted with an arrow.

1.1.2 Quasiparticles: The Quirky Cast of Pertinent Characters

Solids are complicated systems in which each electron and proton is strongly coupled via the electrostatic force to many nearby particles. In principle no behavior can be understood without taking many-body interactions into account, and yet solid-state phenomena are often described in terms of independent excitations, known as *quasi-particles*, that have well defined properties. This picture works because at room temperature the excitations that occur are sparsely populated minor perturbations to the ground state, and behave almost independently of each other. This intuition can be made exact using Green's functions, which show that although quasi-particles are not exact eigenstates of the fully interacting Hamiltonian, they have almost-precisely defined energies and hence long lifetimes, making them useful approximations [6, 7]. Several important quasi-particles are described below.

Charge carriers

Holes, first described in Section 1.1.1, are quasi-particles. The behavior arises due to a "missing" electron moving between sites—it is inherently many-body and there is no actual

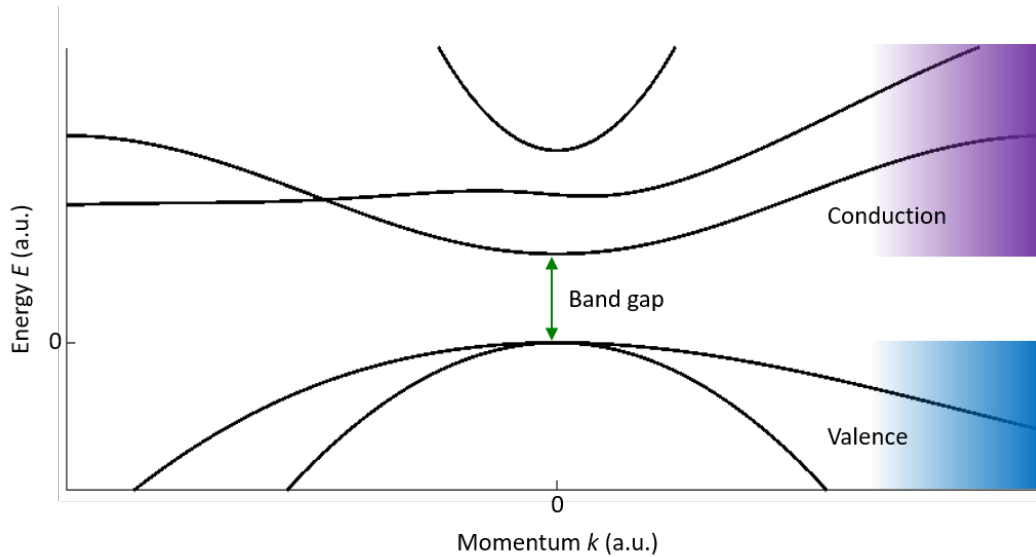


Figure 1.3: Cartoon illustration of band structure. The band gap is the energy difference between the top of the valence bands and the bottom of the conduction bands.

hole particle. But it is simpler to imagine a positive particle moving around the system, and it is also correct in that it results in accurate predictions of material behavior. The hole has well defined energy, momentum, charge, mass, mobility, and other properties.

Excited electrons, despite sharing a name with a very real particle, are also quasi-particles in solids. The surrounding potential changes the way the real electrons behave, but rather than thinking about this many-body problem it is again easier to consider each excited electron as an independent particle that has properties which differ from those of a free electron. Consider mass. Mass affects the relationship between an electron's energy E and momentum k , or its *dispersion*, as

$$E = \frac{\hbar^2 k^2}{2m}. \quad (1.2)$$

In a solid, interactions with other electrons and protons modify this dispersion relation as can be seen in the band structure (Figure 1.3). Although the bands are clearly non-parabolic, they are often well-approximated by parabolas near the top of the valence/bottom of the conduction bands, as shown in Figure 1.4. This parabolic approximation (red dashed line) can be written using a modified form of Equation 1.2, $E = \frac{\hbar^2 k^2}{2m^*}$. The value of m^* is known as the *effective mass* of the charge carrier [3].




Figure 1.4: Using a parabolic approximation (dashed red line) to the dispersion (black lines) in order to calculate effective mass of a charge carrier quasi-particle.

Excitons

Charge carriers in solids often exist due to optical excitation from the valence band to the conduction band, hence an electron and hole are created simultaneously. Electrostatic attraction between electron and hole produces a new, neutral quasi-particle known as an exciton. In materials with strong dielectric screening the attraction is weak and the exciton delocalizes over a large area. These are called “Wannier-Mott” excitons, and they can be modeled with a hydrogenic Hamiltonian, leading to a progression of bound states with energies scaling as $-1/n^2$. The binding energy tends to be on the order of 10 meV [3], meaning that Wannier-Mott excitons easily dissociate into free charge carriers. In materials with weak dielectric screening the electron-hole attraction is much stronger, often around 0.1 - 1 eV. The result is a smaller, tightly-bound Frenkel exciton that generally does not dissociate unless some driving force is applied. Excitons in organic semiconductors (see Section 1.2) are usually of the Frenkel variety. These two limiting cases are diagrammed in Figure 1.5.

It is often important to consider the spin state of an exciton. Both electron and hole are spin-1/2 particles, so their combination must have total spin $S = 0$ or $S = 1$ [4]. There are three possibilities for the spin-1 state depending on whether the spin along the z -axis, S_z , is -1, 0, or +1, hence $S = 1$ is known as the *triplet* state. There is only one spin-0 state, known as the *singlet*. Using the notation $|\alpha\beta\rangle$, where α is the spin- z component of the electron and β is the spin- z component of the hole, these four states are laid out below [4].

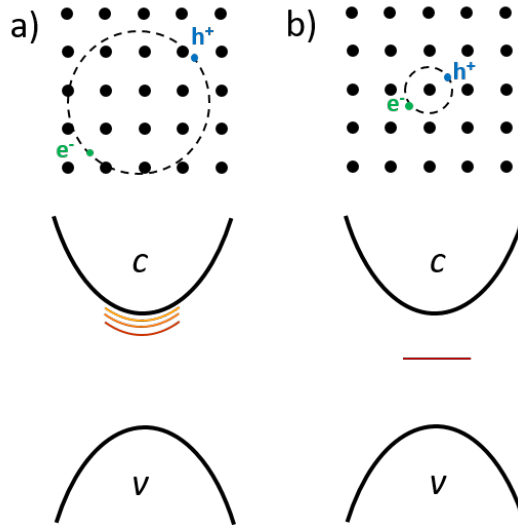


Figure 1.5: Cartoon of (a) Wannier-Mott and (b) Frenkel excitons, showing the real space representation (top) and band structure (bottom). Wannier-Mott excitons are larger and weakly bound—the energy levels (orange-red lines) are only slightly below the conduction band. Frenkel excitons are smaller and strongly bound.

	S = 0 (Singlet)	S = 1 (Triplet)
$S_z = -1$		$ \downarrow\downarrow\rangle$
$S_z = 0$	$\frac{1}{\sqrt{2}}(\uparrow\downarrow\rangle - \downarrow\uparrow\rangle)$	$\frac{1}{\sqrt{2}}(\uparrow\downarrow\rangle + \downarrow\uparrow\rangle)$
$S_z = +1$		$ \uparrow\uparrow\rangle$

Singlet and triplet states generally differ in energy, even if the Hamiltonian is not explicitly spin-dependent. This is because the total wave function for fermions must be anti-symmetric under exchange, so an anti-symmetric singlet spin state must be paired with a symmetric spatial wave function, while a symmetric triplet spin state must be paired with an anti-symmetric spatial wave function. In an anti-symmetric wave function two particles can not occupy the same spatial position simultaneously, hence electrons in the triplet state do not “see” each other or repel one another as much. The exciton is therefore both lower in energy and smaller in extent. This can be formalized for weak excitons: enforcing anti-symmetry results in an additional, repulsive Hamiltonian term that is only non-zero for singlets. If $|\phi_{ck}\phi_{vq}\rangle$ is the state with an electron of momentum k in the conduction band and a hole of momentum $-q$ in the conduction band, then the exchange Hamiltonian term has the form [8]:

$$\langle \phi_{ck'} \phi_{vq'} | H_{ex} | \phi_{ck} \phi_{vq} \rangle = 2 \langle \phi_{ck'} \phi_{vq} | \frac{e^2}{r} | \phi_{cq'} \phi_{vk} \rangle.$$

Ultimately this leads to an exchange energy that depends on the overlap integral between the electron and hole wave functions. Hence, smaller excitons generally have larger singlet-triplet energy differences.

Phonons

Phonons are quantized lattice vibrations. As each atom in the lattice jostles about it pushes and pulls on the nearby atoms, and these disturbances propagate as if each atom were connected to its neighbors via springs. As long as the shaking is not too large in magnitude, vibrations can be described with a set of normal modes that have well-defined momentum (corresponding to the spatial frequency of vibration) and energy (temporal frequency). Therefore a band structure can be computed for phonons just as for electrons.

For a crystal with a single atom per unit cell, the three spatial dimensions give rise to three *acoustic* phonon bands: two for transverse vibrations and one for longitudinal vibrations. In the low-momentum limit (phonon wavelength \gg unit cell size), the relationship between energy ω and momentum k is $\omega = ck$ for some constant c . The speed of sound, $\partial\omega/\partial k$, is therefore constant. Adding a second atom to the unit cell allows for *optical* phonon branches in which the two types of atoms oscillate out of phase. These modes are important in ionic crystals; they can be high in energy even with zero momentum, and are so named because they are generally excited with light [3].

Polarons

The coupling of a phonon and a charge carrier results in a polaron quasi-particle. Intuitively a charge carrier exerts a force on the surrounding lattice, causing it to deform, and this deformation can be described as a superposition of phonon modes. Similar to excitons, stronger interactions (such as in ionic crystals) lead to more localized polarons. This deformation moves through the lattice with the charge carrier, and it has the effect of hindering the carrier's motion. If the deformation is strong enough then the charge carrier can become self-trapped [3].

1.2 Organic Semiconductors Defy Common Approximations

Organic semiconductors are composed of molecules that are made up primarily of carbon and hydrogen. They come in two types: those that are composed of small molecules and those that are composed of polymer chains. Prototypical examples of each are shown in Figure 1.6. As the name suggests, small molecule semiconductors (Figure 1.6a) consist of small organic




Figure 1.6: Representative examples of the two types of organic semiconductors: (a) small molecule and (b) polymer. The materials are (a) TIPS-Pentacene and (b) MEH-PPV [9].

molecules, each of which has a well-defined structure. They can be amorphous or pack into a crystalline form [10]. Polymers, on the other hand, are long chains of indefinite length with carbon backbones (Figure 1.6b). They can be coerced into a more ordered form, but generally form a tangled web [11, 12]. Organic semiconductors are studied for many reasons: the ability to control side groups makes their properties tunable, they can be processed using inexpensive techniques, they form flexible layers [13], and their physics is distinctly unusual. These properties are elaborated upon below, with a focus on small molecule semiconductors as they are the subject of Chapters 3, 4, and 5.

1.2.1 Properties of Organic Semiconductors

Being built of carbon-based molecules, organic semiconductors display several distinct behaviors. For example, carbon is not very polarizable. Its dielectric screening is weak, meaning that excitons are of the tightly-bound Frenkel variety. Furthermore, the bonding between molecules is governed by van der Waals forces [14], which arise due to correlated fluctuations in the instantaneous electric dipole moments of uncharged molecules [15], and are generally on the order of 1 eV for small organic molecules [16]. This is much weaker than the covalent or ionic bonds found in most inorganic semiconductors, and this difference leads to a cascade of new behaviors.

Weak intermolecular forces produce soft materials. Organic semiconductors are essentially plastic, and like plastic they can be bent and stretched as the intermolecular bonds are broken and re-formed. In addition to being important for making flexible devices, this softness also has implications for the phonon structure. Without strong forces keeping them in their equilibrium positions, organic molecules vibrate slowly and with large amplitudes. Compared to inorganic materials the speed of sound is low, the phonon energies are low, and the displacement amplitudes are high. Large displacements of the molecules can strongly affect the electronic states: electron- and exciton-phonon coupling are high [17], which affects the electronic dynamics. There's also a second class of phonons one must consider: intramolec-

ular vibrations. These are vibrations of a single molecule, as opposed to the entire lattice [18]. Because intramolecular bonds are stronger than intermolecular bonds, intramolecular phonons are higher in energy, but they still have large displacement amplitudes and a strong effect on electron dynamics.

Because nuclei in organic semiconductors have significant freedom of motion, they generally rearrange in response to changes in electron density. The nuclear coordinates in the excited state can look quite different from those in the ground state. And yet because electrons rearrange themselves much faster than nuclei can (this is the Born-Oppenheimer approximation, and it is valid because electrons are much lighter than protons [8]), when an organic semiconductor is excited the nuclear coordinates are *not in their lowest energy state*, even for a band edge excitation. This excess energy is converted to intramolecular phonons as the molecules rearranges themselves. The same process also happens when an exciton relaxes to the ground state, and the net effect is that the emission spectrum is lower in energy than the absorption spectrum, a phenomenon known as the *Stokes shift* [19].

Large nuclear motions in organic semiconductors have another important effect: they disrupt translational symmetry, which makes momentum a less-well-defined quantum number and localizes the electrons and excitons. This results in increased electron-hole overlap, which leads to stronger exchange forces. In some cases the exchange force lowers the triplet exciton energy all the way to half that of the singlet exciton, enabling singlet fission (see Chapter 4). Another effect is that excitons in organic semiconductors cannot be entirely described by delocalized band states. The other extreme is to model the states as being localized on a single molecule. In this case the valence and conduction bands become the highest occupied molecular orbital (HOMO) and lowest unoccupied molecular orbital. However the localized picture does not perfectly capture their behavior either, especially for some of the more ordered small organic molecules [20, 21]. This intermediary status makes organic semiconductors difficult to model, but also leads to novel physics.

Finally, organic semiconductors cannot be understood without considering their morphological richness. Because they are pliable they can often adopt both amorphous and crystalline forms within the same sample, and there are often multiple crystalline polymorphs. This variety leads to a rough energetic landscape for excitons and charge carriers—defects and interfaces act as trap sites, decreasing carrier mobility [11]. One particularly important class of morphologies are those that display π -stacking, the out-of-plane overlap of p -orbitals. This can be seen in molecules with an acene core—a string of hexagonal carbon rings attached on their edges, as shown in Figure 1.7a. Of carbon’s four bonding electrons, three form hybridized sp^2 orbitals that comprise strong $C - R$ bonds, and a fourth exists in a π -orbital that is delocalized perpendicular to the plane of the acene core. In crystal structures with parallel acene cores (see Figure 1.7b), the π orbitals on neighboring molecules overlap. The result is a more stable crystal structure and more efficient motion of energy and charge along the π -stacking direction [22].




Figure 1.7: Bonding orbitals in polyacene molecules. (a) Pentacene consists of five acene rings. The in-plane sp^2 orbitals are shown in red. (b) A side-on view showing the out-of-plane p orbitals in green. (c) A crystallized polyacene derivative. The π orbitals on neighboring molecules overlap, leading to strong interactions.

1.2.2 Value of Solution Processing and the Scientific Questions it Raises

From a device point of view, one of the most interesting things about organic semiconductors is that many of them can be solution processed [11]. This means they can be dissolved in solvents and deposited or printed to form a film. Solution processing is rapid and inexpensive, and it can be done without sophisticated equipment. However, its relatively informal nature also makes it more difficult to control, which leads to its own complications and scientific questions. For example, a disordered macrostructure can result, with multiple crystalline polymorphs, amorphous regions, interfaces, and defects all present. It therefore becomes necessary to ask how this heterogeneity affects the film's properties. Do different crystalline domains display different electronic structure? How do interfaces affect charge transport? Spatial resolution in characterization becomes critical, because bulk measurements do not provide much in the way of insight if one does not know how much variability there is between different regions of the film, or if a few individual features have an outsized effect on the overall behavior. This question of heterogeneity will be explored explicitly in Chapter 3. In Chapter 6 I will explore the slightly different but related question of how physical confinement within a nanoscale structure affects electronic dynamics.

In addition to studying the heterogeneity that arises from solution processing, it is also

valuable to go back one step and study the process of film formation. Self-assembly often occurs over the scale of seconds as solvent molecules evaporate. Further film annealing can be performed by applying heat or solvent vapor, giving the molecules the ability to reorient and rearrange. How should one model complicated intermolecular and statistical forces driving self-assembly? How can one tune the deposition and annealing parameters to enforce a desired morphology, and how do the ultimate properties of the film emerge from this morphology? There are many open questions along this causal chain from deposition conditions to morphology to behavior, and a better understanding is crucial to making higher efficiency organic semiconductor devices. The challenge is made especially difficult by the wide range of both time and length scales: one must observe from nanometers (individual molecules) to millimeters (the entire film) and femtoseconds (ultrafast dynamics upon photoexcitation) to seconds (film formation). Self-assembly will be explored in Chapter 5.

1.3 Semiconductor Devices

Semiconductors form the active layer in a wide variety of electronic devices, some of which are described below. When using organic semiconductors, solution processing (Section 1.2) can make it particularly difficult to control the semiconductor morphology and resulting device properties.

1.3.1 Photovoltaics

A photovoltaic (PV) cell converts light energy into electrical energy. With band gaps on the order 1 eV, semiconductors absorb light across a wide range of the visible spectrum and convert that energy into an electron and hole. The electron and hole can produce useful work in the form of current, but only if they are extracted before recombining. As shown in Figure 1.8, this requires an interface between two materials. In a homojunction cell the same material is used on both sides of the interface, but with different doping. A typical inorganic cell uses n-type and p-type silicon, for example. In a heterojunction two materials are used, one that has a propensity to accept electrons and one that donates electrons (accepts holes) [23]. This language is more common for describing organic photovoltaics (OPV). Because of the strong binding energies found in organic semiconductors, light absorbed in an OPV typically produces an exciton instead of free charges. Excitons lifetimes are on the order of nanoseconds, which can make it more difficult to extract current in an OPV and reduce their efficiency. For this reason, in polymer OPVs the donor and acceptor are often blended together to minimize the distance required to diffuse to an interface [24].

A vast amount of solar energy is available to provide electricity for human activities, but there are practical limits to the conversion efficiency of PVs. Significant loss comes from the fact that photons with energy below the band gap are not absorbed, and those with energies above the band gap tend to thermalize before being extracted (Section 1.4). This fact is primarily responsible for the Shockley-Queisser limit of 34% for single junction




Figure 1.8: A prototypical photovoltaic cell schematic. An incident photon (red) creates an electron and hole that are separated at an interface. Electrodes collect the charge carriers and create current.

solar cells [25], though there are ways to circumvent it (see Chapter 4). The efficiency is also influenced by how reliably excitons can find their way to an interface for separation, a process that depends strongly on the morphology of the material and the energetic landscape that morphology creates.

There are several different semiconductor technologies used for photovoltaics, ranging widely in efficiency and cost. Commercial silicon for ordinary use is around 20% efficient [26], while multi-junction solar cells have been made with up to 46% efficiency [27]. Organic and perovskite based photovoltaics, both of which will be discussed in subsequent chapters, have achieved efficiencies in the lab of 15% [28] and 22% [29], respectively. Although OPVs are hindered by their short exciton diffusion lengths, they have the advantage of large absorption cross sections—absorption lengths can be on the order of 100 nm [30], meaning very little material is required and less energy is required for processing.

1.3.2 Light Emitting Diodes

Light emitting diodes (LEDs) are essentially PVs in reverse—free charge carriers are injected through opposite electrodes, re-combine into an exciton, and relax by emitting light. The energy of the emitted light is set by the band gap and hence can be narrow, making LEDs useful for lighting and display applications. Using organic materials in an LED (OLED) has several advantages. OLEDs are lightweight and flexible, and the output color can be tuned by modifying the band gap, either by changing the local structure or adding polar dopant molecules to the active layer [31].

1.3.3 Field Effect Transistors

Field effect transistors (FETs) are crucial for a wide variety of electronics, particularly displays. Like all transistors they are essentially switches, controlling the flow of current through a channel [32]. Exact configurations vary, but all configurations involve three electrodes: source, drain, and gate. The source and drain are connected via a semiconductor through which current flows. The gate is separated from the semiconductor with a dielectric layer [33]. One possible configuration is shown in Figure 1.9.

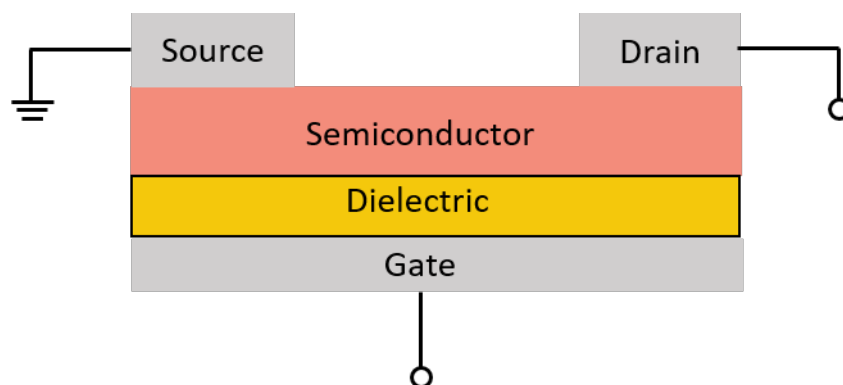


Figure 1.9: Diagram of a top-gate field effect transistor.

Applying voltage between source and drain sends current through the semiconductor. However, a transistor must also be able to amplify or quench this current, and that's where the dielectric and gate come in. An accumulation layer forms at the semiconductor-dielectric interface, which gets populated with mobile charge carriers. Applying voltage to the gate creates an electric field at the interface, which tunes the number of charge carriers in the semiconductor. At one extreme there are few charge carriers and little current flows. Increasing the voltage past a certain threshold voltage causes the current to grow linearly, and as the voltage is further increased the current eventually saturates. There are three primary figures of merit for an FET: threshold voltage, on/off ratio (current in the saturated state divided by current in the off state), and carrier mobility of the semiconductor. The threshold voltage is a measure of the trap density, the on/off ratio is a measure of how cleanly on and off states can be distinguished for use in digital electronics, and the mobility is a measure of both how much power the transistor will draw and how quickly it can be switched between the on and off states. [13, 33, 32].

Active layers of commercial FETs have most commonly been made with silicon, but extensive research on organic FETs (OFETs) has brought their mobilities up to roughly $10 \text{ cm}^2/\text{V} \cdot \text{s}$ [34, 35]. Though far below the roughly $1000 \text{ cm}^2/\text{V} \cdot \text{s}$ [36] displayed by crystalline silicon, this is high enough to be practical for applications outside of the lab. As with all organics they offer the benefits of inexpensive deposition, tunable characteristics, and thin active layers that can be used to make flexible electronic devices [33, 13]. Most carrier transport occurs within a very thin (about 1 nm) layer at the semiconductor-dielectric

interface, so it becomes especially important to create strong in-plane $\pi - \pi$ bonds that will facilitate high mobility [12, 11]. Understanding the impact of morphology on transport is essential to creating high-performing, reliable OFETs.

1.3.4 Sensors

Semiconductors find use in a tremendous variety of sensors, including those that measure light (photodiodes), pressure, temperature, voltage, chemical concentration and acceleration [37]. In general, sensors work by using some external stimulus to create current in the semiconductor, hence an understanding of carrier motion within the active layer can lead to higher performing devices [14].

1.3.5 Lasers

Lasers produce coherent light, which has innumerable uses both as a scientific tool and in devices ranging from medical equipment to bar code scanners. A well-defined optical gap is necessary to produce laser light, hence the gain medium must be a semiconductor. It must also be capable of sustaining a population inversion, and in solid state lasers it is beneficial to have a large oscillator strength, high quantum yield and long excitation lifetime. Many organic semiconductors fit this description. First, they tend to absorb strongly. Second, the vibrational structure and Stokes shift create a natural four-level system with the transition of interest from the vibrationally relaxed excited state to the vibrationally excited ground state. Because thermalization is so fast (see Section 1.4), a population inversion can be created even if absolute excitation levels are low [38]. One class of materials that shows potential as a gain medium are hybrid organic-inorganic halide perovskites, thanks in part to their low trap densities even when solution processed [39]. Perovskites will be discussed further in Chapter 6.

Besides a gain medium, the other key component of a laser is the resonator cavity. The resonator cavity physically defines the allowed standing modes that can exist, and hence sets the output frequency of the laser light. A cavity can be as simple as two mirrors on either side of the gain medium, but there is also interest in creating lasers out of nanostructures in which the gain medium itself acts as the resonant cavity, using smooth facets on opposite sides of the structure to contain light. Both nanowires [40, 41, 42] and nanocrystals [43] can be used, enabling fine control of the lasing properties by tuning the material's size and shape.

1.4 Ultrafast Dynamic Processes In Semiconductors

When a semiconductor is perturbed out of equilibrium, whether by optical absorption or charge injection, a variety of dynamics can occur over femtoseconds to microseconds. Figure 1.10 summarizes several such processes with a Jablonski diagram. Each horizontal line is a

state; the height of the line denotes the state’s relative energy. Arrows represent processes that transfer population from one state to another.

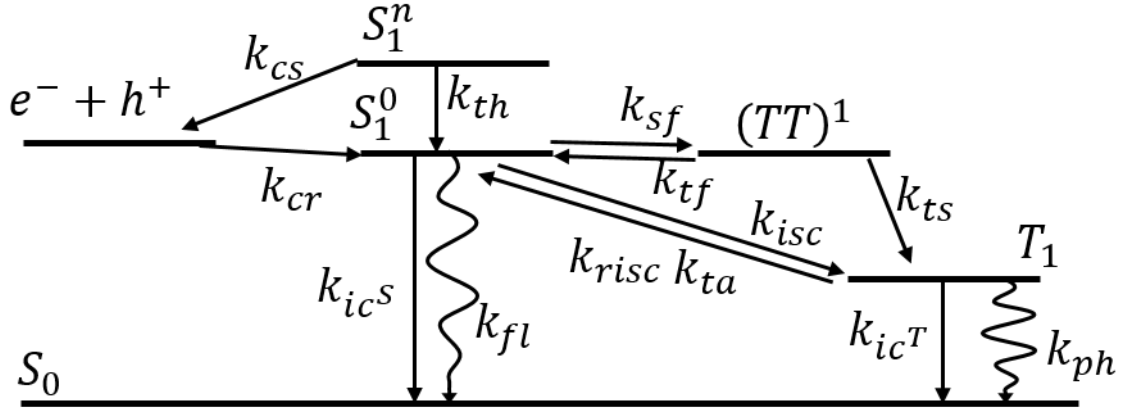


Figure 1.10: Jablonski diagram of the various states (thick lines) and processes (arrows) important for this thesis. The vertical position of each state denotes its energy. See text for definitions.

1.4.1 Thermalization



Consider a an organic semiconductor in its ground state S_0 , optically excited to an excitonic state in the lowest excited state manifold, S_1^n . This state has excess energy above the lowest excited state, S_1^0 , hence it is known as a “hot” or “vibrationally excited” exciton. Yet this heat cannot be contained for long—the exciton rapidly thermalizes (k_{th}) to S_1^0 because of coupling to phonons, transferring its heat to the lattice. Because of the Stokes shift, thermalization can be expected to occur in organic semiconductors even when they are excited resonantly at the band edge. The timescale of thermalization is dictated by the frequency of the relevant phonon modes, and it tends to occur over tens to hundreds of femtoseconds [44].

1.4.2 Linear Recombination



What goes up must come down—excited states eventually lose their energy, sometimes by emitting a photon as photoluminescence (PL). Emission from the singlet is fluorescence (k_{fl}) and from the triplet is phosphorescence (k_{ph}). Fluorescence timescales are generally on the order of nanoseconds [45], while phosphorescence is generally on the order of microseconds or longer [46] because it involves a spin flip [47]. Those excitons that do not emit radiatively generally relax through phonon modes. This is known as internal conversion, and it can occur from the singlet (k_{ics}) or the triplet (k_{ict}). The fraction of excitons that relax radiatively is the *PL quantum yield*, and it can depend strongly on local environment [48].


1.4.3 Charge Separation and Exciton Formation



An exciton may dissociate into free charge carriers (k_{cs}) or two charge carriers may recombine (k_{cr}), under a variety of circumstances. For example if the exciton binding energy is low enough, then an exciton can undergo charge separation. Even a tightly bound exciton, such as those found in organic semiconductors, can still separate at an interface. Charge separation is necessary to generate current in photovoltaics.

Exciton formation generally occurs between two uncorrelated charge carriers, and hence either a singlet or triplet (process not shown in Figure 1.10) state can result, depending on the spin states of the carriers. Spin statistics dictate that triplets and singlets form in a 3:1 ratio [49]. Exciton formation is necessary for LEDs.

1.4.4 Singlet Fission and Triplet Separation



When the triplet energy is roughly one half that of the singlet, a singlet exciton can split into two independent triplet excitons. The process can be divided into two parts. In singlet fission (k_{sf}) the singlet forms a *correlated triplet pair*—two triplets entangled into an overall singlet state, denoted $(TT)^1$. These triplets eventually undergo triplet separation (k_{ts}), losing both spatial and spin coherence [50]. Singlet fission will be discussed in more detail in Chapter 4.

1.4.5 Triplet Fusion and Triplet Annihilation



In the reverse of the above process, two triplets can recombine to form a singlet. When the triplets are a correlated pair formed by fission, this is known as triplet fusion (k_{tf}). If they are uncorrelated triplets that encountered each other randomly, it is called triplet annihilation (k_{ta}) [50]. The terms *geminate* and *non-geminate* recombination might also be used, to refer to triplets that were formed simultaneously and separately, respectively [51].

1.4.6 Intersystem Crossing and Reverse Intersystem Crossing



Intersystem crossing (k_{isc}) is the conversion from singlet to triplet exciton via an electron spin flip. Reverse intersystem crossing (k_{risc}) is used to refer to the reverse process, although it stems from the same physics. Both occur due to spin-orbit coupling—spin and orbital angular momentum are coupled in the electronic Hamiltonian, meaning that while total angular momentum is conserved, eigenstates of the Hamiltonian have mixed spin character. Spin-orbit coupling is stronger for higher energy electrons, so while it is significant in many materials with heavy metal atoms, in most organic semiconductors the timescale is too slow to compete with recombination pathways and is rarely observed [52].

1.4.7 Nonlinear Recombination

Recombination pathways typically involve one exciton and have constant rates. However, in situations with high densities of excitons or charge carriers, certain non-linear recombination pathways become important as well. They tend to be most relevant immediately after photoexcitation, when excitations densities are highest.

Two nonlinear recombination pathways are diagrammed for inorganic semiconductors in Figure 1.11—biexciton recombination and Auger recombination. A biexciton is a bound state between two excitons. As shown in Figure 1.11a, their interaction can cause one exciton to prematurely relax, giving its energy to the second exciton. The second exciton rapidly loses the excess energy through phonons. The relaxation of the first exciton can be either radiative or non-radiative. When radiative it can be distinguished from ordinary fluorescence because it emits photons of slightly lower energy [53, 54], due to the attractive biexciton interaction.

Auger recombination is similar in concept to biexciton recombination, except that it involves three charge carriers. The example with two electrons and one hole is shown in Figure 1.11b. An electron and hole recombine, giving their energy to a second electron, which then relaxes non-radiatively. Both biexciton and Auger recombination may or may not require phonon emission/absorption in order to conserve both energy and momentum [55].




Figure 1.11: Nonlinear relaxation processes in inorganic semiconductors . Filled/open circles are electrons/holes, arrows are transitions, and zigzag lines are phonons. (a) Biexciton recombination—one exciton relaxes due to interaction with a second exciton. In this example a phonon is required to conserve momentum. (b) Auger recombination involving two electrons and one hole.

1.4.8 Diffusion

Diffusion is the motion of quasiparticles due to uncorrelated hops. Exciton diffusion is most important for this thesis, as it is crucial to both charge separation in PVs and light emission in LEDs. Singlet exciton diffusion in organic semiconductors can often be modeled as Förster resonance energy transfer (FRET). FRET occurs due to dipole-dipole coupling between two excitons: one exciton transfers its energy to another as if by photon emission and reabsorption, although the 10 nm length scale is too small for real photons—only virtual photons are involved. An analytic expression for FRET can be worked out using Fermi’s Golden Rule—the matrix element involves the interaction of two dipoles, the energy of which goes as r^{-3} . This matrix element is squared, so the FRET rate goes as r^{-6} . A full expression is given in Equation 1.9 [19]:

$$k_{FRET} = \frac{1}{\tau} \left(\frac{r}{R_0} \right)^{-6} \quad (1.9)$$

$$R_0 = \left[\frac{9c^4 \hbar^4 \eta \kappa^2}{8\pi n^4} \int \frac{A_{abs}(\epsilon) D_{ems}(\epsilon)}{\epsilon^4} d\epsilon \right]^{1/6}$$

τ is the fluorescence lifetime, η is the quantum yield, κ is a transition dipole moment orientation factor, A_{abs} is the acceptor absorption spectrum and D_{ems} is the donor emission spectrum. FRET is most efficient for closely spaced well-aligned molecules with a small Stokes shift.

Because it involves optical pathways, triplet excitons cannot diffuse via FRET; they instead move via Dexter transfer, which is mediated by orbital overlap. The rate of Dexter transfer decays exponentially with donor-acceptor site separation, and is most significant over roughly 1 nm [56]. Triplet diffusivity is generally orders of magnitude slower than singlet diffusivity for the same molecule.

1.5 Structure of Subsequent Chapters

Chapter 2 describes transient absorption microscopy (TAM), a home-built instrument for optically probing ultrafast dynamics with micron-scale spatial resolution. TAM is ideal for studying the heterogeneous, microstructured materials that arise when solution processing organic semiconductors, and is crucial to the subsequent chapters.

Chapters 3 and 4 describe measurements on individual crystalline domains of small molecule organic semiconductors. In the former we observe a variety of behavior across different domains, tie this behavior to variability in the triplet energy levels, and discuss implications for devices. In the latter we exploit the anisotropy within a single domain to optically disentangle the behavior of singlets and triplets. This allows us to make novel measurements of the binding energy and electronic structure of the elusive singlet fission intermediary, the correlated triplet pair.

Chapter 5 investigates the self-assembly of an organic semiconductor during thermal annealing into an intricate spherulite structure. Time-resolved X-ray scattering is used to observe crystallization in real time, and a variety of spatially-resolved techniques are used to investigate the structure and dynamics of the resulting film. These techniques include TAM, microbeam X-ray scattering, scanning transmission electron microscopy, and infrared near-field microscopy. These results are used to inform a simulation that attempts to explain how the film morphology varies with annealing temperature.

Chapter 6 discusses measurements of ultrathin perovskite nanowires. The width of the nanowires is on the scale of the exciton Böhr radius, so quantum confinement effects are relevant in two dimensions. This leads to polarization-dependent optical response, which can be probed with TAM. Highly anisotropic diffusion is also observed.

Taken together, these chapters compose an investigation into the relationship between local structure and electronic dynamics in semiconductor materials.

Chapter 2

Transient Absorption Microscopy (TAM)

2.1 Transient Absorption is a Multipurpose Tool

Transient absorption (TA), a relatively simple nonlinear optical spectroscopy [57], uses one ultrafast laser pulse to excite a sample and a second, time-delayed pulse to probe the ensuing change in absorption. TA is commonly used across many scientific fields to study a wide variety of samples and processes. A minuscule selection of applications includes: tracking the initial energy transfer steps upon light absorption in photosynthetic systems [58], observing a hot-phonon bottleneck in lead-iodide perovskites [59], establishing that singlet fission occurs in an organic semiconductor [60], measuring electron-phonon coupling in metal nanostructures [61], and observing charge separation at the donor/acceptor interface of an OPV [62]. What these and most other applications of TA have in common is that they involve changes in the excited state of a system that occur on ultrafast timescales—femtoseconds to nanoseconds. Indeed, high time resolution is TA’s greatest strength. In this chapter I explain the mechanism of TA and motivate why coupling TA into a microscope is actually a good idea, in spite of the attendant pain and suffering.

2.1.1 Fundamentals

A simple diagram of a transient absorption setup, also known as pump-probe spectroscopy, is shown in Figure 2.1. The first component is the pump laser pulse, which is focused on and excites the sample. In this thesis I focus on optical frequencies, meaning that electronic transitions are being excited and studied, however TA can also be used with infrared light to study vibrational modes [63]. After excitation, the excited state evolves for some delay time τ , whereupon the probe arrives at the excited region of the sample. After passing through the sample the pump is somehow blocked or filtered, and the probe intensity is measured on a detector. If the probe is broad in spectrum, it must be dispersed and measured with a multi-channel detector, but the principle is the same [64]. The presence of excited

species will modify the way the probe interacts with the sample, causing its absorption (or the transmission, which is what is actually measured by the detector) to either increase or decrease. Comparing the probe absorption with and without the pump pulse excitation allows us to determine the dynamics of the excitation. That's where the chopper wheel comes in—it rotates at some frequency f , alternatively blocking and permitting pump pulses to impinge on the sample. The detector measures the probe intensity both with and without the pump, and the difference between the two reports on the properties of the excited state after evolving for time τ .




Figure 2.1: Diagram of a transient absorption setup. Ultrafast pump and probe pulses are focused and overlapped with time delay τ in the sample. The pump is blocked and the probe is sent into a detector. Chopping the pump allows one to measure the differential change in probe transmission due to excitation.

To make the above explanation quantitative, let T_{on} be the probe transmission intensity when the pump is present, and T_{off} be the probe transmission intensity when the pump is absent. The difference in transmission is therefore $\Delta T = T_{on} - T_{off}$. Because the probe transmission can vary from experiment to experiment or from wavelength to wavelength within an experiment, we often normalize by the probe transmission. That is, the reported quantity is $\frac{\Delta T}{T} = \frac{T_{on} - T_{off}}{T_{off}}$.

In addition to being nice and normalized, $\Delta T/T$ also has a simple physical interpretation: it is minus the change in optical density due to pump excitation. To see this, let T be the incident probe transmission strength and OD be the optical density of the un-excited sample. Therefore, $T_{off} = Te^{-OD}$. Let the optical density of the excited sample be $OD + \Delta OD$, and assume that $|\Delta OD| \ll 1$ (as is generally true for transient absorption experiments). Therefore, $T_{on} = Te^{-OD-\Delta OD} \approx Te^{-OD}(1 - \Delta OD)$. We then calculate $\Delta T/T$ as follows:

$$\frac{\Delta T}{T} = \frac{Te^{-OD}(1 - \Delta OD) - Te^{-OD}}{Te^{-OD}} = \frac{-\Delta OD Te^{-OD}}{Te^{-OD}} = -\Delta OD. \quad (2.1)$$

Transient absorption data may be plotted as either $\Delta T/T$ or as ΔOD , but they are identical up to an overall negative sign.

I now consider several physical mechanisms by which pump excitation can alter the probe transmission, and hence create a TA signal. Although a proper treatment requires time-dependent perturbation theory [57], in the incoherent limit (pump and probe are well separated in time), we can take the intuitive approach of thinking about the populations of various states. This is conceptually simple and leads to accurate conclusions.

Ground State Bleach

The ground state bleach (GSB) occurs because excitation bleaches the band edge transition. That is, once the pump has promoted a given electron from the ground to the excited state, the probe cannot induce the same transition on the same electron. This is illustrated simply in Figure 2.2 for a two level system. With the pump off, the probe can be absorbed by the electron (blue) to promote it from the ground state S_0 to the excited state S_1 . However, if the pump has already excited this electron the probe photon cannot be absorbed and transmits through the sample. The change in transmission, ΔT , is positive because more photons are transmitted when the pump is present than when it is absent. If ΔOD is being plotted, the GSB is negative.

In practice the GSB spectrum often has the shape of the band edge transition. It appears immediately after photoexcitation, and its decay can be a useful proxy for the overall decay of the excited state, i.e. return to the ground state.




Figure 2.2: Origin of ground state bleach

Stimulated Emission

Stimulated emission (SE) is often considered in the context of generating laser pulses, but it is also important for interpreting TA data. As long as the excited states are optically coupled to the ground state, probe photons of the appropriate wavelengths can stimulate emission, bosonically stimulating additional photons into the exiting probe field. This is illustrated in Figure 2.3: while ordinarily the probe pulse might transmit through the sample, with the pump on it stimulates emission from the excited state and now there is one additional photon. Hence the transmission measured on the detector is higher than it is without the pump, and ΔT is positive, just as for GSB.

Because GSB and SE have the same sign and both probe the band-edge transition, they tend to overlap initially. However if there is a Stokes shift then as the excitations thermalize the stimulated emission signal will shift to lower energies. This ultrafast red shift is often used as a marker for vibrational relaxation.




Figure 2.3: Diagram of stimulated emission

Excited State Absorption / Photoinduced Absorption

Excited state absorption (ESA) occurs when the probe is absorbed, promoting an electron that has already been excited to the pump to some higher-lying excited state, S_2 . As shown in Figure 2.4, with the pump off the probe might be transmitted, but when the pump is on the probe resonantly couples to the $S_1 \rightarrow S_2$ transition and is absorbed. The transmission is therefore lower when the pump is on, and ΔT is negative. A more general term is photoinduced absorption (PIA), which describes any absorption of the probe that occurs due to the prior excitation by the pump.




Figure 2.4: Diagram of excited state absorption

Electro-absorption

Electro-absorption occurs when excitation leads to strong electric fields within the sample, either because of delocalized excitons or separated charge carriers. The electric field produces

a Stark effect, causing the absorption spectrum to shift. The resulting TA spectrum therefore resembles the derivative of the absorption spectrum. Electro-absorption is most relevant in inorganic materials with high dielectric screening, but it has also been used to observe charge separation in OPVs [62].

Other Considerations

The above picture of populated states goes a long way towards understanding TA spectra, but is by no means complete. The existence of excitations can perturb the electronic structure or prevent other electrons from occupying a specific state [65]. Spectral features tend to overlap each other, and a decay can be misconstrued as a shift or vice versa. Further complications arise due to the non-instantaneous nature of the pulses [66], and when considering the electron states within a proper band structure picture [59]. Careful analysis is often required to understand the physical mechanisms behind TA data.

2.1.2 Interpreting TA decays

A single spectral-channel transient absorption signal might look something like the data sketched in Figure 2.5a. Well before time-zero ($\tau = 0$), when the probe is arriving at the detector before the pump, the only signal is due to scattering of the pump beam into the detector. This non-physical background should be subtracted before doing analysis. Around time $\tau = 0$, when the pump and probe are overlapped temporally, the signal grows in. The width of this growth profile is called the Instrument Response Function (IRF), and is determined by a convolution of the pump and probe pulse profiles – shorter pulses make for a shorter IRF and hence better time resolution.

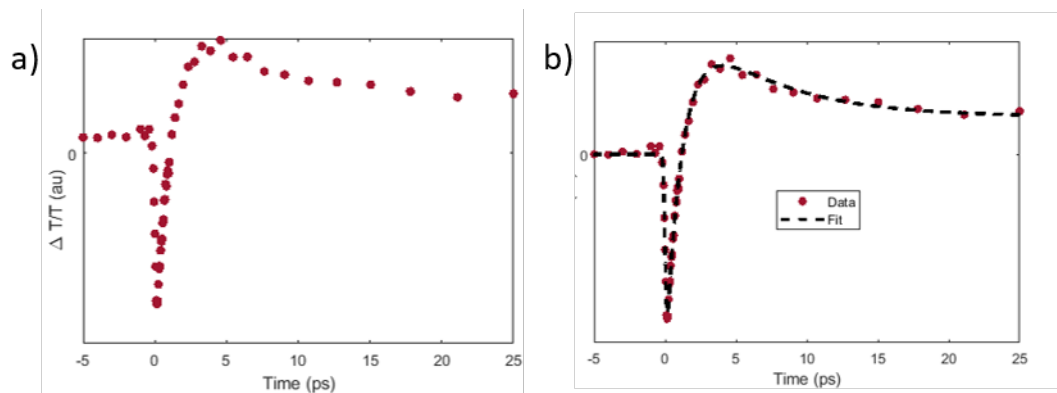


Figure 2.5: Hypothetical (a) TA decay data and (b) fit.

After growing in the signal begins changing as the excited population evolves. In this example the signal changes from negative to positive, and then decays towards its base level before plateauing at a slightly more positive value. What should we make of this decay?

The simplest dynamic picture is that the excitations are evolving from one state to another, and that these different states have different signals. These evolutions are often assumed to have independent, constant rates. Constant-rate kinetics lead to exponential decays, hence TA traces are often modeled using a sum of exponentials. A residual signal at long times indicates the presence of some excited species with a lifetime much longer than what we can measure experimentally. This can be modeled with a constant signal. The curve in Figure 2.5a appears to have two decay components, so the TA dynamics may be modeled as:

$$\frac{\Delta T}{T}(\tau) = A_1 e^{-\tau/t_1} + A_2 e^{-\tau/t_2} + C. \quad (2.2)$$

However, Equation 2.2 does not capture the rise of the signal. There are two ways to approach this problem. If we don't care about getting the best time resolution we can take a simple approach and only model the data after some time $\tau \gg \tau_{rise}$, where τ_{rise} is the temporal width of the signal rise. The more precise method is to multiply the model by a step function and convolve it with the IRF, which is often approximated as a Gaussian of width σ_{rise} . The data would then be fit with:

$$\frac{\Delta T}{T}(\tau) = \left(\frac{1}{\sigma_{rise} \sqrt{2\pi}} e^{-\tau^2/2\sigma_{rise}^2} \right) \otimes \left((A_1 e^{-\tau/t_1} + A_2 e^{-\tau/t_2} + C) \theta(\tau) \right). \quad (2.3)$$

Taking the example data from used above, subtracting the pre-time-zero background and fitting to Equation 2.3 yields the curve shown in Figure 2.5b. It does a good job of fitting the data, and can be used to extract parameters t_1 , t_2 , σ_{rise} , and C . The next task would be to identify the physical origins of the two decays and the long-lived species, which could be done by appealing to the nature of the TA decays or supplementing with additional experiments.

In many cases, especially involving solid samples, an exponential or two will not adequately fit the TA decay. To take an extreme case, if five exponentials are required to achieve a good fit, it's unlikely to be a sign that there are 5 distinct processes being observed. Instead it is a testament to the ability of a function to fit almost any data as long as enough free parameters are included. Hence a more careful approach is needed.

One possible approach is to assume that constant decay rates are still valid, but that there is a network of states with more complicated decay pathways between them. This is known as a kinetic model, and examples can be found in both Chapters 3 and 4. One advantage of kinetic models is that they can be solved exactly—the resulting system of equations can be written as a linear differential equation using matrices and solved exactly (see Chapter 3). One disadvantage is that the possible model permutations grows rapidly as more states are introduced, so great care must be taken and physical motivation used to avoid adding too many free parameters and over-fitting the data.

There are also situations in which a simple system of linear rates is not an appropriate description of the system. For example there could be heterogeneity within the sample, so that instead of probing one decay rate we are probing a distribution of decay rates. In that case a stretched exponential could be used to fit the data. It could also be that linear rates are not appropriate, especially if excitation densities are high and nonlinear decay pathways

are activated. Higher order rate equations can be used to model the data, but adding nonlinearities makes the resulting system of equations generally not solvable analytically, and hence fitting becomes much slower.

2.1.3 Interpreting TA spectra

When using a broadband probe pulse and a multi-channel detector, the TA signal becomes a function of both time *and* wavelength: $\frac{\Delta T}{T}(\tau, \lambda)$. The underlying kinetics do not change however, so the signal can still be fit using a sum of exponentials or a kinetic model. The difference is that the signal amplitudes, A_i , becomes spectra, $A_i(\lambda)$. So for example, instead of using Equation 2.2, one might use the following:

$$\frac{\Delta T}{T}(\tau, \lambda) = A_1(\lambda)e^{-\tau/t_1} + A_2(\lambda)e^{-\tau/t_2} + C(\lambda). \tag{2.4}$$

Two spectral representations of TA data are commonly used: Evolution associated spectra (EAS) and Decay associated spectra (DAS). In a kinetic model, each EAS is the spectrum for a specific state. DAS, on the other hand, correspond to a decay [67]. That is, if there is a decay pathway from state j to state k , which have EAS A_j and A_k , then the DAS for that decay is $A_{jk}^{DAS}(\lambda) = A_j(\lambda) - A_k(\lambda)$. This is useful because in some cases the spectra for the individual states might not be informative, but seeing how the spectrum changes can convey information. For example, if the decay involves a shift of the stimulated emission, then the DAS will display a telltale Gaussian derivative line shape.




Figure 2.6: Example TA spectra, illustrating how time-zero (black line) can depend on probe wavelength.

Another complication arising with TA spectra is that the IRF can be wavelength dependent. Because the refractive index of any transmissive optic varies with wavelength, different components of the probe pulse will travel at different speeds, leading to a *chirped* beam in

which the central frequency varies from the beginning to the end of the pulse. The net effect is that time-zero is a function of wavelength—different wavelengths arrive at the sample at different times. To model the data we must know time-zero at each wavelength, $t_0(\lambda)$. An example is shown in Figure 2.6 taking some data from Chapter 6: color is used to plot ΔOD vs. wavelength and time. The time at which the signal appears is sketched with a black line – it clearly depends on λ . Fortunately an IRF spectrum can be measured and the decay at each λ_i fit using $\text{IRF}(t; \lambda_i)$ [64] without sacrificing temporal resolution. If the IRF is approximated as a δ function, this is equivalent to measuring $t_0(\lambda)$ and shifting all of the measured data points $\Delta OD(t, \lambda) \rightarrow \Delta OD(t - t_0(\lambda), \lambda)$.

2.2 Combining Transient Absorption with Microscopy

Transient Absorption Microscopy (TAM) is similar to TA, but the pump and probe pulses are focused with a microscope objective to a specific spot on the sample. Depending on the objective used, the spot sizes can range in diameter from several micron down to the diffraction limit of hundreds of nanometers. The benefits, challenges, and uses of TAM are discussed below.

2.2.1 Benefits

The most obvious advantage of TAM is that smaller laser beams can study smaller objects. For example if one has a sample of nanostructures dispersed on a substrate, and their separation is larger than the beam size, then one can measure the TA signal of an individual nanostructure [68]. This allows for the observation of *heterogeneity*—variation between different structures that would be averaged out by a bulk measurement. Because a TAM setup generally includes a camera for imaging the sample, it then becomes possible to correlate the TA spectra with the optical image of the sample [69], and possibly with other microscopy techniques as well. Even if the sample is a continuous film, one can still collect TA spectra at different positions and compare them [70].

A possibly less-appreciated aspect of TAM is that it allows for polarization-resolved measurements. One can collect TA spectra for several different probe polarizations, θ , meaning that the TA signal is now three dimensional: $\frac{\Delta T}{T}(\tau, \lambda, \theta)$. This is informative because at the microscopic scale many samples are spectroscopically anisotropic. Consider a polycrystalline material. A bulk measurement interrogates many randomly-oriented crystalline domains, meaning that all directional information is averaged out, and the probe polarization is irrelevant. However a local measurement can probe a single crystalline domain with well-defined crystal axes. Light polarized along different crystal axes will therefore have different interactions with the sample [71, 72, 73, 50].

This concept is illustrated in Figure 2.7. The green skeleton is a sketch of one molecule in a molecular crystal, standing on a gray substrate (based off of the pentacene derivatives studied in Chapters 3 and 4). Some electronic transitions have transition dipole moments (TDMs)




Figure 2.7: Illustration of how polarization-resolved TA takes advantage of structural anisotropy in a material. The green structure approximates the shape of a single acene molecule. In its crystalline form all nearby molecules are oriented in the same way. Rotating the probe polarization (orange arrows) allows for selective coupling to different transition dipole moments (red/blue arrows).

polarized along the long axis of the acene core (blue arrow) and some have TDMs polarized along the short axis of the acene core. Rotating the probe polarization (two possibilities are shown as orange arrows) therefore selectively couples to different TDMs, disentangling their signals and allowing for a better understanding of the underlying dynamics. This is particularly useful in situations where contributions to the TA signal obscure or cancel each other because they have opposite signs.

Another advantage of TAM is that it allows for the observation of spatial dynamics, such as diffusion. By rastering the probe pulse and measuring the TA signal at each position, one can observe exciton migration: the signal strength will drop at the center of the pump volume whereas there will be a corresponding rise at the edge of the pump volume [74]. This change is un-measurably small for traditional macroscopic TA, but using smaller laser spots makes for larger gradients in excitation concentration, and hence more significant diffusion.

2.2.2 Challenges

TA microscopy is not without its drawbacks. The first challenge is weaker signals – shrinking the beam means that significantly smaller regions of the samples are probed, so fewer photons are involved in the measurement and the signal to noise ratio decreases. As an illustration, shrinking the pump beam diameter from $150\ \mu\text{m}$ (typical for traditional TA) to $1.5\ \mu\text{m}$ shrinks the beam area (and hence the excitation volume) by a factor of 10^4 . To get around this issue people often use higher power densities, either by increasing the energy density per pulse (fluence) or by increasing the laser repetition rate. However, increasing the fluence runs the risk of damaging the sample. Even if no damage occurs, creating more excitations per unit volume could push the sample into the nonlinear regime where excitons are interacting

with each other and new decay pathways are opening up. The TA signal is then fluence dependent. As discussed in Section 2.1 more complicated modeling is necessary to capture the dynamics, and that model may not capture the dynamics that would be observed at lower fluences. Increasing the laser repetition rate, on the other hand, can cause long-lived excited species or heat to build up. If the repetition rate is 100 MHz for example, then any species with a lifetime of about 10 ns or longer will still be prevalent when a new pump pulse arrives. As with increasing the fluence, this complicates the interpretation of the data.

Other challenges are due to the mechanics of using a microscope. For example, everything expands/contracts/drifts with slight changes in temperature, and these changes are more pronounced when working in a microscope. If the beam sizes are on the order of microns, then a drift of just 100 nm can lead to a significant change in beam overlap and hence signal strength. The microscope objective also causes problems because of all the glass inside of it. Recall from Section 2.1 that transmissive optics chirp the beam, and sending a broadband probe through several centimeters of objective glass can lead to a beam that is dispersed over 10 ps, as opposed to less than 1 ps in typical TA experiments. This can limit the time resolution. To get around this issue, one group recently built a TA microscope that only uses an objective to collect light post-sample. They use exclusively reflective optics before the sample, and achieved sub-10 fs resolution [75].

2.2.3 Uses

Many early TAM experiments investigated individual metal nanostructures, in part because metals can be excited at high intensities without damage [68]. Studies of silver nanocubes show that excitation energy is rapidly converted to heat through a phonon “breathing” mode, which manifests as an oscillatory TA signal. The oscillations in ensemble measurements are weak and dephase quickly, but single nanoparticles oscillate with a clear frequency. This frequency can be correlated with the oscillation lifetime and particle size, and by building up statistics researchers were able to understand the origin of the breathing mode. Metal nanowires have also been studied, which allows for the exploitation of light polarization — light polarized along the longitudinal axis of the wire interacts differently with it than light polarized along the transverse axis of the wire [68].

Since then, as the technique has become more widespread and laser systems have achieved higher stability, the spatial resolution of TAM has been turned on a wide variety of problems in physics, chemistry, materials science, and biology [76]. That may involve studying a small crystal, such as a single-layer flake of a transition metal dichalcogenide [70, 77] or a single crystalline domain of an organic semiconductor [71]. Other work has probed different spots on a nanowire to investigate the role of strain [69]. In some cases people have gone smaller and examined interfaces, which are crucial for understanding the behavior of electronic devices, but notoriously difficult to investigate due to their small size. In some cases TAM laser pulses may be small enough to study the interface directly, but more common is that TA spectra are measured at multiple positions around the interface, and modeling is used to extract the signal due to the interfacial region. These ideas have been used to observed

a long-lived charge transfer state forming at donor/acceptor interfaces in an OPV [78], to detect and measure the morphology of buried interfaces in an organic semiconductor [72], and to observe quenching at grain boundaries in a perovskite thin film [75].

There is another class of TAM experiments in which the pump and probe are spatially separated in order to track how energy or charge move through a material after excitation. Due to the large number of spatial points studied these measurements generally require extremely high laser repetition rates in order to achieve sufficient signal/noise, but when successful they have resulted in quite interesting results. In one study researchers measured the local thermal diffusivity of nanowires, and were able to quantify how it is affected by substrates [79]. Other notable examples include the observation of long-lived quasi-ballistic hot carrier diffusion in halide perovskite films [74] and the way in which singlets and triplets inter-convert to boost diffusion lengths in an organic semiconductor [80, 81]. The latter is notable because multiple probe wavelengths were employed in a two-color geometry (see Section 2.3). Some probe wavelengths selectively interrogated the singlet and some the triplet.

2.3 Ultraslow Evolution of Transient Absorption Microscope

Over the years, our TA microscope has undergone several upgrades in order to enable new science. Three distinct phases are diagrammed in Figure 2.3. Abbreviations used are: Pol (polarizer), WP (half-wave plate), RR (retroreflector), obj (objective), PMT (photomultiplier tube), BP (bandpass filter), f/S (subharmonic frequency output), BS (beam splitter), NF (notch filter), and DG (diffraction grating).

The generation of the laser pulses is omitted from Figure 2.3 because it has remained constant. We use an 80 MHz mode-locked Ti:sapphire Coherent Vitara oscillator to create a seed pulse, which feeds a Coherent Legend-Elite regenerative amplifier. The regenerative amplifier first stretches the seed pulse in time and then combines it in a cavity with a high-energy laser pulse generated by a Coherent Evolution laser. The cavity contains a Ti:sapphire crystal that acts to amplify the seed pulse. After exiting the cavity the pulse is compressed, resulting in a 5 kHz output that is both ultrafast and high in energy.

Stage 1 (Figure 2.3a), which is used in Chapter 3, was built by Cathy Wong, Sam Penwell, and Ben Cotts. It is characterized by a *degenerate probe* and *non-collinear geometry*. The output of the regenerative amplifier is sent into a Coherent OPerA Solo optical parametric amplifier (OPA) and its output produces both the pump and the probe. This beam is sent through a prism compressor (not pictured) followed by a telescope and optional 50 μm pinhole to act as a spatial filter (not pictured) before being split with a 50:50 beam splitter. The pump line passes through a chopper wheel which is driven by a controller synchronized to some subharmonic of the 5 kHz laser repetition rate. Both beams then go through retroreflectors, but the pump retroreflector is on a 150 mm motorized delay stage,

which can be moved to control the delay between pump and probe with 1 μm (6.67 fs) resolution.

Before entering the microscope, both beams pass through a broadband half-wave plate and a broadband polarizer. The polarizer is set to a desired polarization, and then the half-wave plate can be rotated to vary the beam power. The probe passes through an additional half-wave plate in a motorized mount. Rotating this optic rotates the probe's polarization *without* changing its power or pointing. This feature is crucial for measuring the polarization-resolved TA signal.

The beams are re-combined using a mirror in a half-open mount. The pump reflects off the edge of the mirror while the probe slips by its edge. The two beams are now parallel but non-collinear as they make their way into the objective below the sample. It is crucial that the two beams are small enough to both fit inside the clear aperture of the objective. They focus onto the sample and are re-collimated by a top objective. Because the beams are non-collinear we can use a mask to block the pump while the probe enters a photomultiplier tube (PMT).

The PMT signal is fed into a lock-in amplifier referenced to the chopper repetition rate. The output is ΔT , the change in probe transmission between “pump on” and “pump off.” The PMT signal is also fed into a second lock-in amplifier (not pictured) referenced to the laser repetition rate. This output is the average probe intensity, T .

There are several advantages to this setup. It is simpler to align and more stable than its successors. It also allows for better time resolution—both pump and probe can be sub-50 fs. But there are drawbacks as well. The spatial resolution is limited because the beams must fit side-by-side into the objectives, so they cannot be too large and hence cannot focus as tightly as beams that over-fill the objective. It is also more susceptible to thermal drift. The region of overlap between the focused beams and the sample is limited to just a few microns in all three dimensions. Micron-scale drifts in either objective or the sample can therefore diminish the signal. This problem was addressed by building a thermally insulating box around the microscope, and having the experimental control program regularly adjust the Z position of the sample to re-optimize the signal. Every N scans of all time delays (where N is user-specified and typically chosen so that the check is done every five minutes or so), the delay stage moves to some pre-set delay time and scans the vertical sample position Z to create a plot of $\Delta T/T$ vs. Z . It then moves Z to the position that gave the strongest signal. These data are saved for later investigation, so the user can get a sense of how much drift occurred.

The biggest drawback of Stage 1 is that it provides limited spectral information. Because the probe has limited bandwidth and is detected by a single-channel detector, we only measure electronic transitions within a limited energy range and the signals due to all such transitions are averaged together. The beam energy is generally chosen to correspond with the band gap of the sample so we primarily monitor transitions around the band edge, which are important. But we have limited information about the contributions of ground state bleach vs. stimulated emission, about shifts in transient absorption spectral shapes, and about any excited state transitions at energies above or below the probe energy. This limits

the information that can be gleaned about our samples.

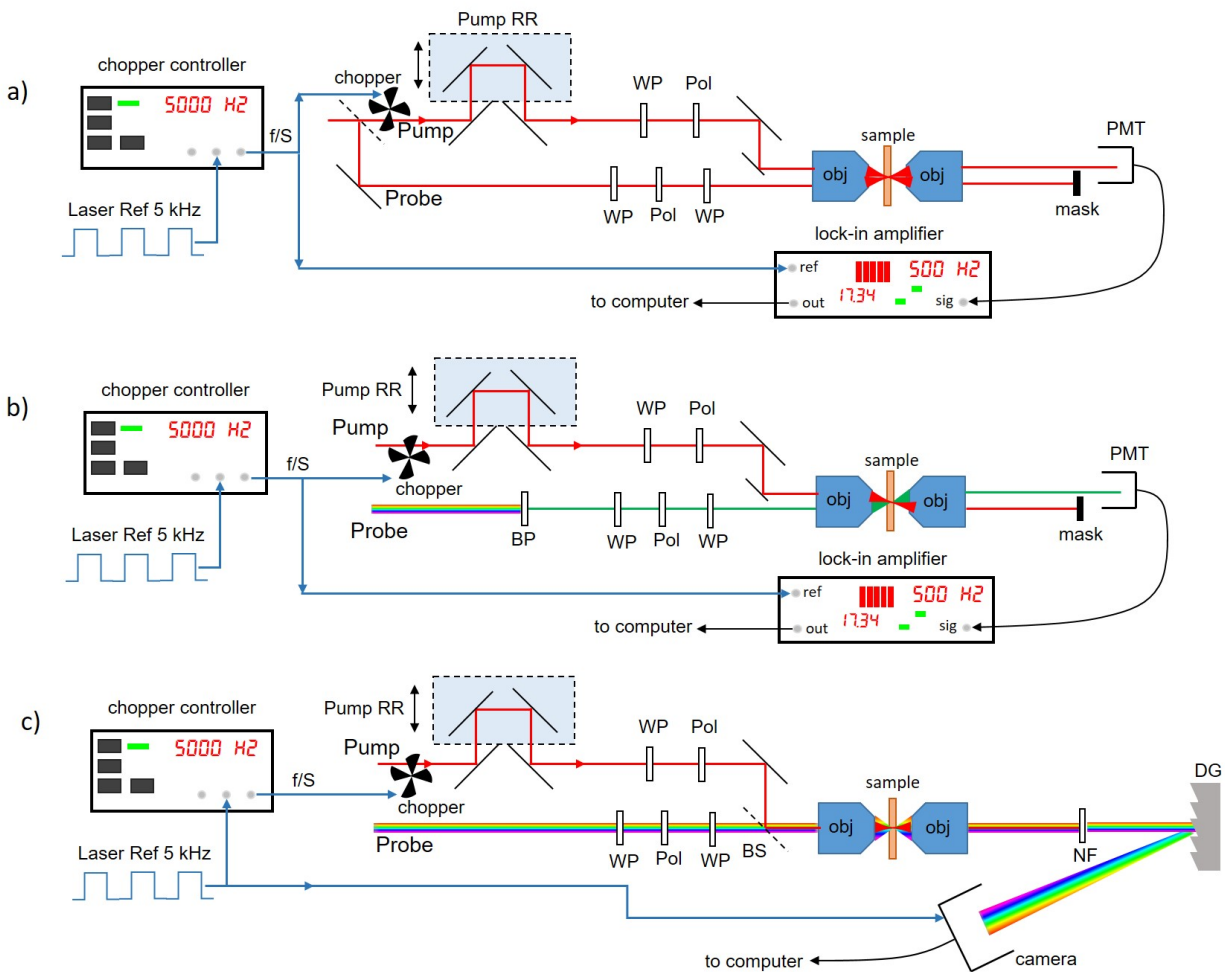


Figure 2.8: Iterations of our group’s TA Microscope: (a) single-color, (b) two-color, and (c) broadband.

Stage 2 (Figure 2.3b), which is used in Chapter 4, is characterized by a *non-degenerate probe* and a *non-collinear geometry*. In addition to generating the pump pulse, part of the output of the regenerative amplifier is diverted and used to generate a broadband pulse in a nonlinear crystal: either sapphire, Nd:YAG, or CaF_2 . This pulse is then filtered with a 10 nm bandpass filter to create the probe. The probe is focused and collected in the same manner as in Stage 1.

The primary advantage to Stage 2 is that we can swap out the bandpass filters to investigate different regions of the spectrum, and hence obtain a fuller picture of the sample’s electronic transitions. Because the pump and probe are spectrally distinct, we can also spectrally filter the pump, which is more effective than the spatial filter used in Stage 1

and produces signals with much lower background levels. However, it is time-consuming to collect data separately for each probe wavelength, and because the scans for different wavelengths are done under non-identical conditions it can be difficult to make quantitative comparisons between them. In addition, the use of bandpass filters limits the resolution in both wavelength and time. The spectral resolution is obviously limited by the 10 nm bandwidth of the filter. Cutting the bandwidth also limits the temporal resolution, because frequency and time form a Fourier transform pair. In practice, time resolutions around 100-200 fs are achieved. Finally, this setup is also prone to thermal drift because it also uses a non-collinear geometry.

Stage 3 (Figure 2.3c), which is used in Chapters 5 and 6, is characterized by a *broadband probe* and *collinear geometry*. A broadband pulse is generated as in Stage 2, but in this case it is combined on a beam splitter with the pump and they both enter the microscope through the center of the objective. A notch filter removes the spectral region around the pump, and a diffraction grating disperses the remaining light onto a line camera.

This setup enables full transient absorption spectroscopy with micron-scale resolution. Using the full clear aperture of the objective we achieve probe spot sizes around 3 μm in diameter ($1/e^2$), which could be improved further with higher-numerical aperture objectives. We achieve sub-nanometer spectral resolution, and temporal resolution is theoretically limited by the pump (in practice it is around 100 fs). The specs of the three stages are tabulated in Table 2.1. The collinear geometry makes the setup easier to align and less sensitive to drift. We still cannot measure the transient absorption spectrum around the excitation wavelength, but this is a problem common to most TA setups, and generally dealt with by shifting the excitation wavelength away from the region of interest. Potential future improvements could involve removing transmissive optics to improve temporal resolution or inserting a microscope cryostat to control the temperature of the samples [82].

Stage	Time resolution (fs)	Probe size (μm)	λ resolution (nm)
1	35	8	N/A
2	150	6	10
3	100	4	0.25

Table 2.1: Specs of TA Microscope stages

Chapter 3

TAM Discerns Inter-Domain Heterogeneity in diF-TES-ADT

One strength of TAM is its ability to measure electronic behavior within a single crystalline domain of a small-molecule organic semiconductor. This technique was initially applied to study TIPS-Pentacene (TIPS-Pn), finding among other things that all domains have similar electronic dynamics [71] (an extension of this research is discussed in chapter 4). In this chapter I present the results of a similar study on diF-TES-ADT, in which we observed evidence of inter-domain heterogeneity. I first discuss the interpretation of TAM data, and discuss ways of quantifying heterogeneity. I then present a kinetic model that allows us to relate the observed heterogeneity to variation in more fundamental material parameters. The physical origins of this variation are discussed. The work here is published in [73].

3.1 Characteristics of diF-TES-ADT films

2,8-difluoro5,11-bis(triethylsilylethynyl)anthradithiophene (diF-TES-ADT) is a pentacene derivative shown in Figure 3.1a [83, 84, 85, 86, 87]. Its solid-state absorption spectrum is shown in Figure 3.1b. The first vibronic peak of the $S_0 \rightarrow S_1$ transition is addressed with a laser pulse centered at 550 nm, shown in red.

Optical images of dropcast diF-TES-ADT samples are shown in Figure 3.2. Samples were formed by purchasing diF-TES-ADT powder from Sigma-Aldrich and dissolving to 5 mg/ml in toluene (the triethylsilyl side groups make the molecule soluble). 150 μL solution was dropcast onto a clean glass coverslip and then covered with a large petri dish to slow the evaporation rate. We observe large crystalline domains, which are promoted by the fluorination of the anthradithiophene core [83]. As these domains are many microns in extent, the micron-scale laser pulses used in TAM can easily interrogate a single domain. The films were found to be stable and optically active in atmosphere for weeks, though they were susceptible to bleaching under laser fluences that reached into the hundreds of $\mu\text{J}/\text{cm}^2$.




Figure 3.1: (a) Structure and (b) absorption spectrum (black) of diF-TES-ADT thin film. Red curve in (b) is the excitation laser spectrum.




Figure 3.2: Optical images of diF-TES-ADT dropcast films. Fairly uniform single crystalline domains are visible in (b).

3.2 Measuring Heterogeneity with Single-Domain TAM Data

3.2.1 Interpreting Polarization-resolved TAM Data

TAM measurements were performed within individual domains in a dropcast thin film of diF-TES-ADT, and example data are shown for three domains in Figure 3.3. The top panels show ΔT as a function of pump-probe delay time, τ . ΔT is shown instead of $\Delta T/T$ because the transmission is a function of probe polarization, therefore using it to normalize the signal presents additional complications (for more discussion see Section 4.2.8). Each curve corresponds to a fixed polarization of the probe relative to the pump. Some probe polarizations yield positive TA signals and some negative, but all decay towards 0 as a function of time. The bottom panels show a slice of ΔT vs. probe polarization for a fixed delay time $\tau = 1$ ps. The TA signal oscillates between positive and negative values in a \cos^2

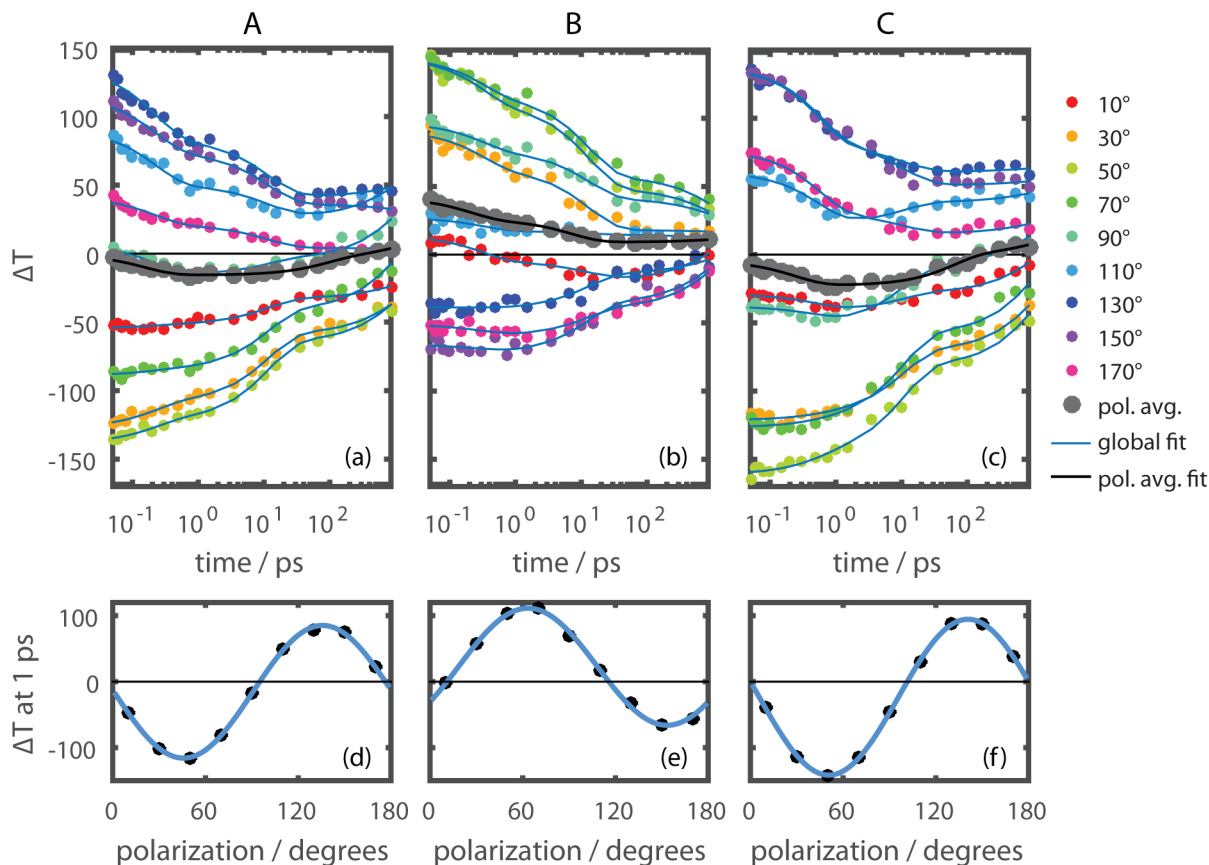


Figure 3.3: Polarization-resolved TAM data of three different diF-TES-ADT domains (labeled A-C), visualized in two ways. (a)-(c) ΔT vs. time where each series corresponds to a fixed probe polarization. Lines are global fits to a sum of three exponentials. Gray dots are polarization-averaged data, which themselves are fit to a sum of three exponentials. (d)-(f) ΔT vs. probe polarization for a fixed delay time of 1 ps. Lines are fits to a cosine squared plus offset.

fashion.

The oscillatory nature of ΔT vs. θ is evidence of the anisotropic nature of single crystalline domains, which allows the probe to selectively couple to distinct optical transitions. As discussed in Chapter 2, this indicates that after photoexcitation the probe is resonant with both excited state absorption (ESA) and ground state bleach (GSB) transitions, which result in negative and positive transient absorption signals, respectively (Figure 2.7). These transitions must have transition dipole moments (TDMs) that have different orientations in the lab-frame, so the coupling of the probe to each of these transitions will change with the probe polarization for a given domain. In comparing different domains for a given probe polarization, domains oriented so that their GSB transition (typically along the polyacene short axis) is aligned with the probe field will have dominantly positive TA signals as compared

with those domains oriented so that their ESA transitions (typically along the polyacene long axis) align with the field. diF-*TES*-ADT has a packing structure with a single molecule per unit cell, so the delocalized crystal TDMs point in the same directions as those of an individual molecule. In the diF-*TES*-ADT domain measured in Figure 3.3a,d, about half of the probe polarizations produce positive signals, which means that GSB dominates ESA, and half of the probe polarizations produce negative signals, where ESA instead dominates GSB. The projection of these two different TDMs onto the sample plane determines their relative abilities to couple to the probe pulse’s polarized electric field. The relative amounts of GSB and ESA can thus be used to characterize and compare the non-azimuthal orientation of different crystalline domains in a film.

3.2.2 Quantifying Heterogeneity

We quantify heterogeneity between domains of diF-*TES*-ADT by considering both the decay timescales and the relative balance of GSB and ESA. First, the timescales: Figure 3.3 shows transients collected at a number of different probe polarizations for each domain. Because the pump polarization remains the same for all measurements in a particular domain, the initial exciton population is the same for each measurement in the polarization series; the differences among the polarized transients in a given domain are due to the polarization-dependent coupling strength between the probe and the available transitions. The physical processes which occur to the photogenerated exciton population depend on the electronic structure and dynamics of the domain, and will be the same irrespective of the probe polarization. Thus, in fitting the data to a weighted sum of exponential decays, the time constants in the transients collected at each probe polarization should be the same, but the amplitudes should change owing to the coupling between the probe and the available transitions. Because the same time constants are used for all transients, this is known as a “global fit.” We performed a global fit for each individual domain. Within a domain the transient for each probe polarization was fit to the function

$$y = A_1 e^{-t/\tau_1} + A_2 e^{-t/\tau_2} + A_3 e^{-t/2300\text{ps}} + C. \quad (3.1)$$

In the fit to y , the time constants τ_1 and τ_2 are global, while the pre-exponential amplitudes A_1 , A_2 , and A_3 , and offset C are different for each polarization. The time constant of the third exponential was set to the fluorescence lifetime of diF-*TES*-ADT, which represents the longest time scale of loss of the singlet excited state population, and which has been measured in films to be 2300 ps [45]. The constant C is due to long-lived non-emissive species such as triplets. The resulting fits are shown as blue curves in Figure 3.3a-c, with extracted time constants shown in Table 3.1. The fit A_1 , A_2 , A_3 , and C values are in Table 3.2, and plots of the dependence of these fit values on probe polarization are shown in Figure 3.4 for a representative sample domain. The amplitudes have a cosine-squared dependence on polarization, as would be expected from the varying coupling of the probe pulse electric field vector to the TDMs.

domain	global		polarization averaged	
	t_1 (ps)	t_2 (ps)	t_1 (ps)	t_2 (ps)
A	0.23 ± 0.02	12.1 ± 0.8	0.16 ± 0.04	90 ± 34
B	0.33 ± 0.05	13.5 ± 1.1	0.21 ± 0.05	7.1 ± 1.4
C	0.60 ± 0.06	11.4 ± 0.9	0.24 ± 0.08	67 ± 19

Table 3.1: Fits of TAM data to three exponential for the three domains shown in Figure 3.3. Results are shown for both a global fit (all polarizations fit simultaneously with one set of time constants) and a fit to polarization-averaged data. Disagreement between domains is indicative of heterogeneity. Error bars are 1σ .

pol		A_1	A_2	A_3	C		A_1	A_2	A_3	C		A_1	A_2	A_3	C
10°	A	408	-169	-244	-246	B	53	-39	23	26	C	372	-145	-243	-251
30°		389	-147	-234	-250		164	-74	-42	-27		378	-140	-259	-266
50°		356	-134	-222	-230		339	-128	-131	-96		288	-105	-186	-206
70°		245	-97	-135	-138		442	-158	-201	-147		178	-70	-91	-104
90°		112	-57	-36	-32		453	-160	-231	-165		69	-36	-5	5
110°		35	-34	34	43		372	-131	-197	-140		23	-30	39	55
130°		29	-37	38	42		199	-74	-102	-68		56	-45	15	31
150°		106	-67	-21	-12		67	-33	-24	-6		153	-82	-58	-49
170°		189	-88	-93	-98		8	-14	20	27		250	-113	-151	-149
avg.		210	-91	-101	-104		236	-90	-99	-67		198	-85	-102	-107

Table 3.2: Fitting coefficients for the TAM data collected on three diF-*TES*-ADT domains and shown in Figure 3.3

The transients for each polarization can be averaged together within a single domain and fit to Equation 3.1. These data and the fits are shown as gray dots in Figure 3.3a-c, and the ensuing time constants are in Table 3.1. This polarization average is analogous to a bulk measurement, provided that each of the domains in the sample were identical in structure and had the same crystalline orientation relative to the substrate-normal. Yet the differences among the time constants of the three polarization-averaged fits demonstrate that even averaging over rotations in the substrate plane is insufficient to reach the homogenized transient obtained for a bulk measurement. These differences therefore report on the heterogeneity of the domains in the film.

The time constants resulting from fits to the averaged data and from the global fits to the polarization dependent series do not agree, and as shown in Table 3.1 the fits to the averaged data generally have more uncertainty. This shows that the dynamics can be more accurately determined by using TAM to measure individual domains with multiple probe polarizations, likely because of the mutual cancellation of ESA and GSB upon averaging. In particular, the intermediate time constant τ_2 is substantially lengthened from about 10 ps to several tens of picoseconds when averaging over polarizations, and the error in the parameter value

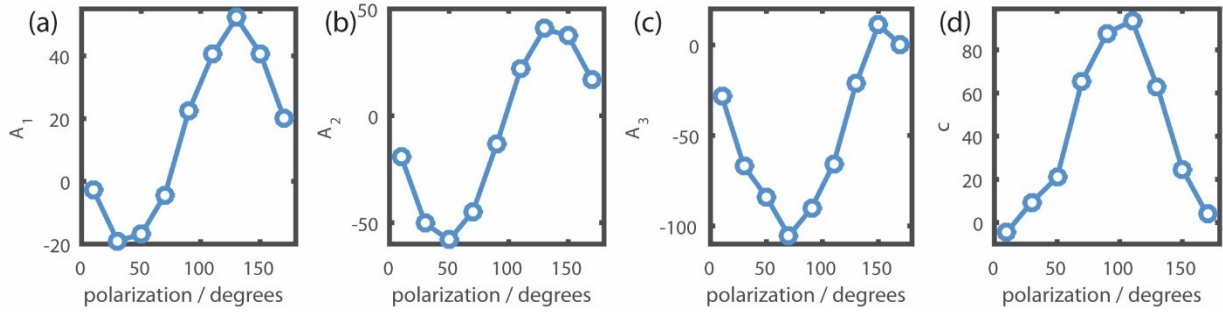


Figure 3.4: Polarization dependence of global fit parameters A_1 , A_2 , A_3 , and C for Domain A.

increases commensurately. This is not surprising, as when compared to bulk measurements or to a single polarization-averaged transient the set of single domain polarized transient data introduces far more fitting constraints than additional fit parameters. Each additional probe polarization introduces constraints equal to the number of time points measured (23), while only adding a free parameter for each decay term (of which there are 4).

As shown in Table 3.1, the three diF-TES-ADT domains studied in Figure 3.3 have similar dynamics, in that there are fast (sub-picosecond), intermediate (tens of picoseconds), and long (2300 ps) time scale processes observed in each one (the physical origins of these timescales are discussed in the next subsection). The time constants extracted from global fits for the three different domains do not, however, agree with one another to within 1σ confidence intervals. Full transients, similar to those in Figure 3.3a-c, were collected for eight different diF-TES-ADT domains, and their time constants were determined by global fits, as described above. Figure 3.5a,b shows histograms of these τ_1 and τ_2 values, respectively. The τ_1 and τ_2 values are reported as fractions of the average values from these eight domains, 0.595 and 15.9 ps, respectively. This can be compared to Figure 3.5c,d, which shows histograms of the fast and intermediate time constants obtained from the global fits of transients collected for twenty seven different TIPS-Pn domains in a previous work [71]. The average τ_1 and τ_2 values for these domains are 0.046 and 3.17 ps. The percentage variability in time constants observed in diF-TES-ADT is several times that observed in TIPS-Pn films, especially apparent for τ_1 . The spread of τ_1 and τ_2 values from eight domains of diF-TES-ADT is larger than the spread of values found after twenty seven measurements of TIPS-Pn. A positive value indicates that GSB dominates while a negative value indicates that ESA dominates, and larger magnitudes indicate stronger imbalance between the two components. For the domains shown in Figure 3.3, these offset parameters are -0.08, 0.18, and -0.03, respectively. This indicates that ESA dominates over GSB in domains A and C, while GSB dominates over ESA in Domain B.

We also consider heterogeneity in the relative balance between GSB and ESA, as can be seen in Figure 3.3. For example GSB dominates over most probe polarizations in Domain B, but ESA slightly dominates in Domains A and C. We quantify this variation by first fitting

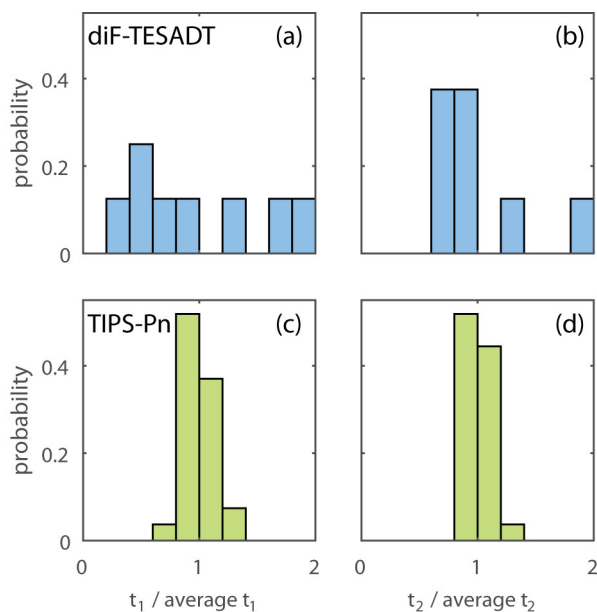


Figure 3.5: Histograms of time constants τ_1 (a,c) and τ_2 (b,d) from global fits of TA signal collected from individual domains of diF-TEs-ADT (a,b) and TIPS-Pn (c,d). Time constants are reported as fractions of their average values.

$\Delta T(\theta)$ at some fixed time to a cosine squared function $A \cos^2(\theta - \theta_0) + C$. In Figure 3.3d-f these data and the fit are shown for time $t = 1$ ps. Next we define the *normalized offset* as the average value of the curve divided by its peak-to-peak amplitude.

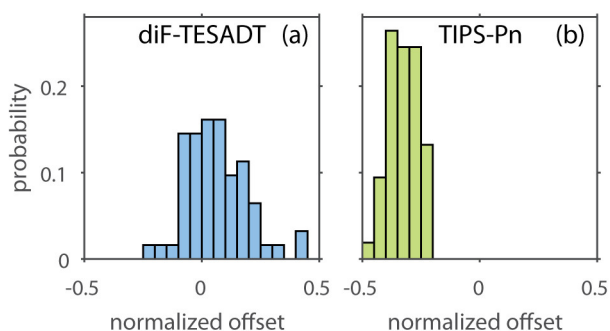


Figure 3.6: Histograms of normalized offset in single domains within thin films of (a) diF-TEs-ADT and (b) TIPS-Pn. Bin size is 0.05.

To visualize this heterogeneity, we construct histograms of the normalized offset measured in different domains. Using TAM, 62 diF-TEs-ADT domains were inspected, and their offset parameters calculated at $\tau = 1$ ps. The histogram of the resulting normalized offsets, shown

in Figure 3.6a, indicates that the normalized offset in the inspected domains can range from -0.22 to +0.42. These domains show more heterogeneity in their normalized offset than TIPS-Pn does: Figure 3.6b shows a histogram of the offset parameters measured for 51 domains of TIPS-Pn, centered more narrowly about roughly -0.35. The distribution of offset parameters in the measured TIPS-Pn domains is more than 2 times narrower than the measured diF-TEs-ADT domain distribution.

3.2.3 Kinetic Model

To gain deeper insight into the electronic processes occurring in diF-TEs-ADT after photoexcitation, we fit the polarization-resolved TAM data to a kinetic model. The model is based on one previously used to fit polarization-resolved TAM data of single TIPS-Pn domains [71], although the details differ. In our model the pump promotes some population from the ground state S_0 to the hot excited state S_1^n . This population relaxes to the thermalized excited state S_1^0 with rate constant k_{th} , which is analogous to the first time constant of our global fits.

The S_1^0 population can relax to the ground state with rate constant k_r , which is fixed to the fluorescence lifetime [45], $k_r = \frac{1}{2300\text{ps}}$. Singlet excitons can also undergo singlet fission, forming a correlated pair of triplets with rate constant k_{sf} . The correlated triplet pair, $(TT)^1$, may annihilate to re-form a single singlet exciton and a ground state occupancy with rate constant k_{ta} . If E_{TT} and E_S are the energies of the triplet pair and singlet exciton, respectively, then these rates are related through detailed balance by

$$k_{ta} = k_{sf}e^{(E_{TT}-E_S)/k_B T}. \quad (3.2)$$

Equilibrium between the singlet and triplet states will be established with rate $k_T = k_{sf} + k_{ta}$. The inverse of k_T is most analogous to the second time constant of our empirical triexponential fits.

Correlated triplet pairs can dissociate into a pair of separated (but still spin-correlated) triplet, $T + T$, with rate k_{ts} . Such triplets will eventually return to the ground state via phosphorescence or by encountering one another to re-form a correlated triplet pair, neither of which are included in our model because they happen on timescales longer than those of the experiment.

Our invocation of singlet fission as the primary electronic process on ultrafast timescales is justified by the evidence for long-lived species in our data. The global fit of the polarization-resolved TAM data required a constant, C , indicating that some optically active species is present beyond the singlet fluorescence lifetime. Triplets are the most likely candidate. Because there are no heavy atoms in diF-TEs-ADT, inter-system crossing is unlikely to occur on these timescales, so singlet fission is the best explanation.

We also observe evidence for a long-lived species in bulk TA measurements on a thin film of diF-TEs-ADT. These measurements were performed with a 75 fs, 554 nm pump focused to 200 μm full width at half maximum, and a broadband probe focused to half that size.

The probe was spatially filtered and passed through a 10 nm bandpass filter to select several probe wavelengths: 550 nm, 600 nm, 640 nm, 671 nm, and 694 nm. TA traces are shown in Figure 3.7 for the 550 nm probe (blue) and the 671 nm probe (red). All other studied probe wavelengths yielded results that were nearly identical to those at 671 nm.

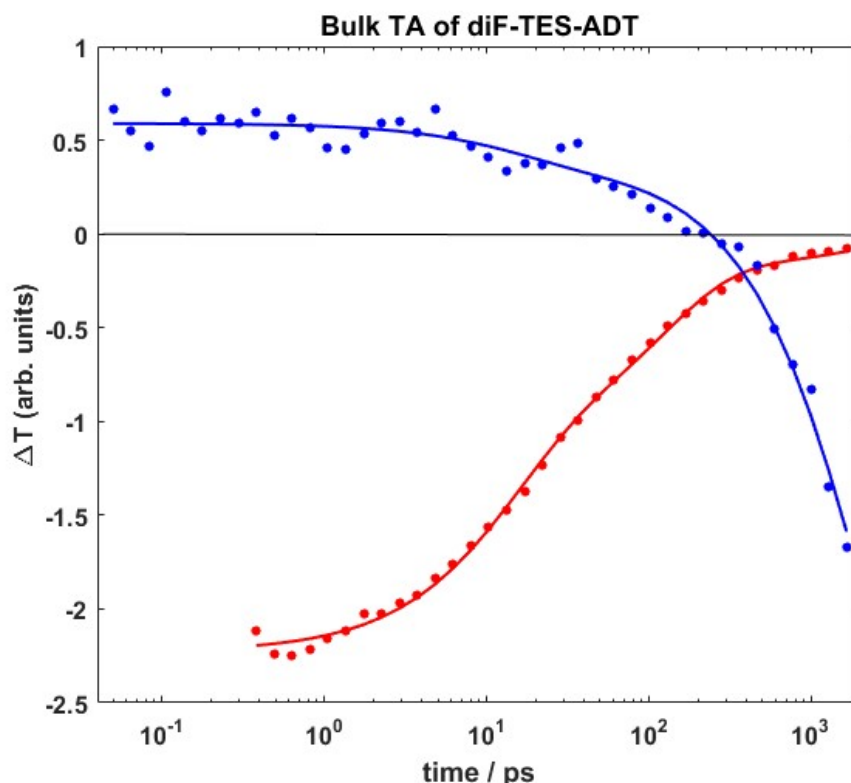


Figure 3.7: Bulk TA on diF-TES-ADT for two selected probe wavelengths, 550 nm (blue) and 671 nm (red). Pump fluences are comparable to those used in TAM. TA was measured at 20 kHz using a Light Conversion PHAROS regenerative amplifier pumping a non-collinear parametric amplifier (NOPA, also Light Conversion). The sample was translated horizontally after each scan, cycling between 20 discrete points spaced by 0.1 mm.

We observe an early GSB on the $S_0 \rightarrow S_1$ transition (550 nm), which competes with an ESA to net a small signal. Over the fluorescence timescale the GSB disappears and what remains is a significant ESA that persists for at least 1 ns. At 671 nm (and other red wavelengths) we observe an initial ESA that appears to diminish over two timescales and then disappear entirely on the fluorescence timescale. Therefore the long-lived triplet state does not have an ESA in the region 600 - 694 nm.

The curves are fit to a sum of exponential decays and the results are shown in Table 3.3. The first timescale of the global fit, τ_1 , is too fast to be seen with this instrument response function. We find that timescales corresponding to τ_2 and fluorescence lifetime $\tau_3 = 2300$ ps

are needed for both curves. The 671 nm probe curve requires an addition timescale, which we denote τ'_2 , whereas the 550 nm curve requires a long-lived constant term, C . The long-lived constant provides more evidence for triplets, while the necessity of an additional decay term τ'_2 could be evidence of triplet separation or that the laser is probing many heterogeneous domains.

probe λ	τ_2 / ps	τ'_2 / ps	A_2	A'_2	A_3	C
550 nm	15 ± 8	N/A	0.21 ± 0.04	N/A	3.8 ± 0.1	-3.5 ± 0.1
671 nm	13 ± 1	120 ± 10	-1.07 ± 0.07	-0.97 ± 0.06	-0.19 ± 0.02	N/A

Table 3.3: Time constants from fits of bulk TA data at two probe wavelengths. The most general form of the fit is $A_2e^{-t/\tau_2} + A'_2e^{-t/\tau'_2} + A_3e^{-t/2300\text{ps}} + C$. 1σ confidence intervals are shown. Though the fit parameters are phenomenological, their physical origin can be understood via the kinetic model.

With the model justified, we can implement it. Let P_α be the population of state α . The populations of the various states obey the following coupled first-order differential equations:

$$\begin{aligned}
 \frac{dP_{S_1^n}}{dt} &= -k_{th}P_{S_1^n} \\
 \frac{dP_{S_1^0}}{dt} &= k_{th}P_{S_1^n} - k_{sf}P_{S_1^0} + k_{tf}P_{(TT)} \\
 \frac{dP_{(TT)}}{dt} &= k_{sf}P_{S_1^0} - k_{tf}P_{(TT)} - k_{ts}P_{(TT)} \\
 \frac{dP_{T+T}}{dt} &= k_{ts}P_{(TT)}
 \end{aligned} \tag{3.3}$$

These equations can be solved analytically by writing them in matrix form [71]. We define:

$$\vec{P} = \begin{pmatrix} P_{S_1^n} \\ P_{S_1^0} \\ P_{(TT)} \\ P_{T+T} \end{pmatrix}, \quad \mathbf{K} = \begin{pmatrix} -k_{th} & 0 & 0 & 0 \\ k_{th} & -k_{sf} & k_{tf} & 0 \\ 0 & k_{sf} & -k_{tf} - k_{ts} & 0 \\ 0 & 0 & k_{ts} & 0 \end{pmatrix} \tag{3.4}$$

and re-write Equation 3.3 as $\frac{d\vec{P}}{dt} = \mathbf{K}\vec{P}$. This equation can be solved if \mathbf{K} is diagonal, which it is not, but we can make it diagonalize it by finding the eigenvalues and eigenvectors. Let \mathbf{V} be a matrix of eigenvectors of \mathbf{K} , and $\mathbf{\Lambda}$ be a diagonal matrix of eigenvalues. Then $\mathbf{V}^{-1}\mathbf{K}\mathbf{V} = \mathbf{\Lambda}$, and we define $\mathbf{U} = \mathbf{V}^{-1}\vec{P}$. Employing this transformation,

$$\begin{aligned}
 \frac{d\vec{U}}{dt} &= \mathbf{V}^{-1}\frac{d\vec{P}}{dt} = \mathbf{V}^{-1}\mathbf{K}\vec{P} = \mathbf{\Lambda}\mathbf{V}^{-1}\vec{P} = \mathbf{\Lambda}\vec{U} \\
 \Rightarrow \frac{d\vec{U}}{dt} &= \mathbf{\Lambda}\vec{U}
 \end{aligned} \tag{3.5}$$

Because $\mathbf{\Lambda}$ is diagonal, Equation 3.5 can be solved as $\vec{U}(t) = e^{\mathbf{\Lambda}t}\vec{U}_i$, where $\vec{U}_i = \mathbf{V}^{-1}\vec{P}_i$ and

$\vec{P}_i = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$ is the initial population distribution (everything in the S_1^n state). Transforming

back to the initial variables, we have an analytic expression for $\vec{P}(t)$:

$$\vec{P}(t) = \mathbf{V}^{-1}e^{\mathbf{\Lambda}t}\mathbf{V}\vec{P}_i. \quad (3.6)$$

With the population levels solved for, we can calculate the GSB, SE, and ESA signals that make up the TA signal by writing expressions for each relevant optical transition. Each of these signals are modeled as \cos^2 functions of probe polarization, θ :

$$A(\theta) = a(\cos(\theta - \theta_0))^2 + c. \quad (3.7)$$

GSB and SE are assumed to have the same TDM, $A_0(\theta)$. The TDMs for ESA from the various states are $A_1^n(\theta)$, $A_1^0(\theta)$, and $A_T(\theta)$. The paired and separated triplets are assumed to have the same ESA. Although it is not enforced by the model, we would physically expect $A_1^n(\theta)$ and $A_1^0(\theta)$ to have the same phase, and this is found to occur.

We apply this model to fit the polarization-resolved TAM data from the three domains shown in Figure 3.3. Results are shown in Figure 3.8. The fit is quite good—it captures the wide variety of observed dynamics. Each domain yields 16 parameters: k_{th} , k_{sf} , ΔE , k_{ts} , as well as a , c , and θ_0 for each of the four TDMs. Although it does not correspond to a single physical process, we report k_T as it more directly related to what we observe in experiment. These parameters are in Tables 3.4 and 3.5.

domain	k_{th}^{-1} / ps	k_{sf}^{-1}	ΔE ($k_B T$)	k_T^{-1} / ps	k_{ts}^{-1} / ps
A	0.42 ± 0.05	18 ± 2	-0.7 ± 0.2	12 ± 1	130 ± 50
B	0.44 ± 0.08	1300 ± 200	4.40 ± 0.04	15 ± 2	∞
C	0.77 ± 0.06	15 ± 2	-0.3 ± 0.1	9 ± 1	130 ± 30

Table 3.4: Time constants and singlet-triplet pair energy splitting calculated from fitting domains A, B, and C to the kinetic model.

The results of the kinetic model fit reproduce the results of the global fit discussed above. They also offer more insight into how Domains A and C are similar, but Domain B is distinct. Considering the timescales, k_{th}^{-1} can be compared to τ_1 of the global fit, and is seen to be comparable, though slightly higher. k_T^{-1} can be compared to τ_2 , and yields excellent agreement. In both domains A and C the triplet pair is about $0.5 k_B T$ below the singlet in energy, and triplets dissociate over roughly 130 ps, indicating fairly efficient singlet fission. Subsequent work has also found that singlet fission occurs in diF-TEs-ADT [88]. The authors found slower timescales for both triplet formation and separation, which may be due to the fact that they studied an amorphous film in which intermolecular coupling is




Figure 3.8: TAM from three domains of diF-TES-ADT (same data as in Figure 3.3), along with the kinetic model fit (solid lines).

expected to be weaker. Returning to these results, domain B seems to have a triplet pair that is much higher in energy than the singlet, and does not show evidence of triplet dissociation. Singlet fission may not occur in this domain, and a different model may be necessary.

The TDMs also show differences between Domains A/C and B. Domains A and C show similar amplitudes and offsets for all TDMs, and the phases are consistently similar. Domain B by contrast does not yield sensible results for A_T . The relative phases are not the same either—when the phase of A_0 is rotated to be equal to that of Domains A and C (values in parentheses), we can see that the phases A_1^0 and A_1^n for domain B do not match those in domains A and C. It is clear that Domain B shows qualitatively different excited state dynamics than Domains A and C, particularly when it comes to singlet fission.

3.3 Physical Origins of Heterogeneity

The origins of the observed heterogeneity can be divided into structural heterogeneity inferred from variations in transient time constants, illustrated in Figure 3.5, and the primarily orientational heterogeneity inferred from variations in the normalized offset, illustrated in Figure 3.6.

domain		A	B	C
A_0	c	29.4 ± 0.2	35 ± 3	30.0 ± 0.2
	a	16.3 ± 0.4	24 ± 2	16.1 ± 0.6
	θ_0 (°)	2 ± 1	-58 ± 2 (2 ± 2)	2 ± 1
A_1^0	c	-4 ± 1	-31 ± 7	0 ± 2
	a	198 ± 3	212 ± 4	215 ± 3
	θ_0 (°)	41.1 ± 0.3	-32.3 ± 0.4 (27.7 ± 0.4)	46.2 ± 0.4
A_1^n	c	12 ± 1	11 ± 12	10 ± 1
	a	284 ± 1	250 ± 8	301 ± 1
	θ_0 (°)	37.2 ± 0.3	-37.3 ± 0.4 (22.7 ± 0.4)	45 ± 0.3
A_0	c	44 ± 1	-1470 ± 30	37.4 ± 0.7
	a	63 ± 1	3910 ± 10	71 ± 1
	θ_0 (°)	22 ± 1	-121 ± 1 (-61 ± 1)	29 ± 1

Table 3.5: Transition dipole moments, calculated using the kinetic model, for GSB/SE (A_0), ESA from the hot excited state (A_1^n), thermalized excited state (A_1^0), and from triplets (A_T). Values in parentheses are the phases for TDMs in domain B after being rotated by 60° . This is to make the A_0 TDM line up with those in domains A and C, and facilitate comparison with the other TDMs.

Heterogeneity in dynamics is indicative of structural disorder among crystalline domains. Namely, differences in the crystal lattice parameters should generate differences in intermolecular electronic coupling and in the resulting electronic structures of the crystalline domains, some examples of which are sketched in Figure 3.9. Our measurements are sensitive to these differences through the statistically significant variations that we observe in the transient time constants for diF-TEs-ADT (Figure 3.5a,b). We note that, despite the sensitivity of our measurements, we observe no similarly sized variation among TIPS-Pn domains (Figure 3.5c,d). We thus conclude that the casting of our diF-TEs-ADT film supported the nucleation of slightly different crystal structures (Figure 3.9a), whereas those in the TIPS-Pn films seem far more homogeneous (Figure 3.9b). Films of diF-TEs-ADT are known to exhibit two polymorphs, with a transition temperature near room-temperature [86]. It is therefore possible that both crystal forms exist in our films, although the distribution of structural parameters that we have obtained is not obviously bimodal. It is possible that the heterogeneity of crystal structures we infer in diF-TEs-ADT films causes the singlet-triplet energy gap to fluctuate, resulting in variable rates of singlet fission among domains.

Heterogeneity in the normalized offset is indicative of non-azimuthal orientational disorder. Here, we distinguish between rotations of the crystal structure about the normal to the substrate plane and all other rotations. The former do not affect the offset parameter since the projections of the GSB and ESA transition dipole moments relative to one another are preserved upon rotation about the normal (see Figure 2.7). On the other hand, all other rotations of a crystallite (Figure 3.9c) do not preserve the relative strength of these projections




Figure 3.9: Schematic illustration of possible variations in structure and non-azimuthal orientation of small molecule crystal domains. (a) Structural variations, as inferred for diF-TES-ADT, involve changing the intermolecular spacing and angle with respect to the substrate. (b) TIPS-Pn appears to maintain a uniform intermolecular spacing and angle with respect to the substrate. (c) The non-azimuthal orientation of diF-TES-ADT crystal grains appears to be variable even if independent of lattice spacing.

and are hence able to alter the normalized offset. Although differences in packing structure could also, in principle, affect this parameter by altering the electronic structure, we assume that they are a less significant contribution. In TIPS-Pn films, there are hardly any variations in the offset parameter (Figure 3.6b). Therefore, we interpret that the angle of the TIPS-Pn molecules with respect to the substrate is essentially preserved in all of the many measured domains (Figure 3.9b). By contrast, it appears that the diF-TES-ADT crystal grains can take on various orientations rotated about an axis in the substrate plane (Figure 3.9c), as evidenced by the much larger spread in offset parameters observed over many diF-TES-ADT domains (Figure 3.6 a). Although it is possible that there exists a positive correlation between the extent of structural and orientational variations in small-molecule organic semiconducting films, a specific form of interdependence in our time constant and GSB/ESA metrics is not obvious.

3.4 Future Directions

We have illustrated that TAM can be applied to study single domains of a thin film solution-processed organic semiconductor, yielding clearer dynamics and revealing striking inter-domain heterogeneity. Extensions of this work could include correlating dynamics with local crystal structure, comparing other small molecules, and exploring the implications for transistors.

Though we hypothesize that the interdomain heterogeneity we observe is due to two distinct polymorphs that are known to occur in diF-TES-ADT [86], a deeper understanding of the relationship between structure and dynamics requires correlative imaging. That is, imaging the same domain with one technique to measure the local crystal structure and another to measure the dynamics. The crystal structure can be measured using either X-ray or electron diffraction techniques. Microbeam-grazing incidence wide-angle X-ray scattering (μ GIWAXS) has been used to map the local structure in spherulites (see Chapter 5) of TES-ADT (the non-fluorinated version of diF-TES-ADT), using beams focused to $15 \mu\text{m} \times 15$

μm [89]. This is small enough to interrogate the single domains studied here. Transmission electron microscopy (TEM) has also been used, in one study looking at single crystals of both diF-*TES*-ADT and *TES*-ADT [17]. The authors were able to extract not only the crystal structure but also the amplitudes of side chain motions. Therefore either X-ray or electron techniques could be used to correlate our observed dynamical heterogeneity with structure.

The comparison of TIPS-P_n and diF-*TES*-ADT shows that similar molecules can display significantly different levels of heterogeneity. Studying more molecules of this class could reveal a full range of behaviors and help us understand what makes a certain molecule more prone to heterogeneity. Potential candidates include *TES*-ADT, rubrene (see chapter 5), *TES*-Pentacene [90], and TIPS-Tetracene [91].

In the context of organic thin film transistors, a film's degree of interdomain heterogeneity could correlate to its macroscopic mobility, and this degree of heterogeneity could determine device consistency. For solution-cast polycrystalline films with impressive intradomain mobilities along the π -stacking direction, charge transport isotropy is important for device consistency [92, 11]. For example, a randomized array of crystallites provides a higher multiplicity of charge carrier scattering trajectories from source to drain than a single crystal whose π -stacking direction is misaligned with the channel. It is possible that the interdomain orientational heterogeneity that we discovered promotes such charge transport isotropy and could be used as a predictor for consistent, high-performing devices. By the same token, too much structural heterogeneity could create energetic traps. In order to investigate this balance one could create diF-*TES*-ADT films under a variety of casting conditions, use them to fabricate transistors, and measure the mobility anisotropy, hence revealing the balance between isotropy and trapping.

Chapter 4

TAM Probes the Correlated Triplet Pair in Single Domains of TIPS-Pentacene

In this chapter I discuss the ultrafast electronic dynamics of the organic semiconductor TIPS-Pentacene, which we study in unprecedented detail thanks to TAM. In particular we focus on singlet fission, in which a singlet exciton splits into two spin-correlated triplet excitons (Section 1.1.2) that can eventually dissociate, allowing for potentially higher-efficiency photovoltaics. Correlated triplet pairs are known to form rapidly in TIPS-Pentacene, but due to the spectrally overlapping contributions from several signals it has been difficult to pin down the triplet pair formation timescale. TAM allows us to probe single crystalline domains in a polycrystalline sample, and taking polarization-resolved measurements allows us to track several transition dipole moments separately. In this way we deduce the kinetics of triplet pair formation. Furthermore, polarization resolution allows us to isolate a subtle change at later times that we attribute to triplet pair dissociation. Performing simulations, we convert the measured triplet dissociation timescale to a triplet-triplet binding energy, and discuss this binding in the context of improving device performance. An additional benefit of polarization resolution is that it enables us to determine the absolute strength of the absorption of each excited species. Using a broadband probe we measure excited state absorption spectra, and find that triplet-triplet interaction perturbs the absorption and hence electronic structure. We explain this perturbation with a model in which the charge transfer character of the correlated triplet pair allows it to mix with nearby singlet excited states. This model is substantiated with first-principles calculations of excitons in crystalline TIPS-Pentacene. The work here is published in [50].

4.1 Singlet Fission and the Correlated Triplet Pair

4.1.1 Current Understanding of Singlet Fission

As introduced in Section 1.4, singlet fission converts a singlet exciton into two triplet excitons via a multi-step process. The initially generated triplets are entangled into an overall spin-0 state, known as the correlated triplet pair [93, 94, 95], and denoted $(TT)^1$. This state eventually dissociates into a pair of non-interacting triplets, although long-range spin correlations may persist [93, 95]. The correlated triplet pair is thus the linchpin of singlet fission: its electronic structure affects the dynamics of triplet formation and triplet separation, both of which are crucial for triplet harvesting. A proper understanding of the correlated triplet pair could therefore improve the efficiency of singlet fission-based photovoltaic devices, since although efficiencies of up to 45% are possible in theory [94], and progress has been made [96, 97], overall efficiencies remain low. Little is known about the energetics or kinetics of the correlated triplet pair, in part because it is often short-lived, and may be nearly indistinguishable spectroscopically from dissociated triplets. Direct measurements of the correlated triplet pair will both improve our fundamental understanding of singlet fission and lead to better design principles for devices.

The basic mechanism of fission remains a heated topic of research [98, 99, 100, 101, 102, 103, 104]. The fission rate is thought to depend strongly on the nature of the excited singlet and triplet states and on the associated details of the electronic structure, in particular on the presence of charge transfer (CT) states [105, 106, 107], through orbital overlap [98, 99] and vibronic coupling [99, 108, 109]. As a result, fission is sensitive to crystal structure: two polymorphs of the same material can have fission rates that vary by an order of magnitude or more [110, 111, 112, 113, 114], and within a disordered film fission can preferentially occur at specific sites [115]. Such sensitivity makes it difficult to predict a priori whether or how rapidly fission will occur in a given material, but one design principle that people have used to guide the quest for high efficiency fission devices is to engineer π -stacking. π -stacking occurs when aromatic rings pack face to face, creating overlap between the π -orbitals on neighboring molecules [116]. π -stacking often leads to high carrier mobilities, strong coupling between electrons on neighboring molecules, and relatively delocalized excitons [117, 118, 119]. Due to this strong coupling, it is generally assumed that slip-stacked crystals with strong $\pi-\pi$ interactions can potentially yield high fission efficiencies and rates [94, 106, 120]. Two recent theoretical studies have, however, found that strong coupling between singlets [121] or between triplets [100] may have a detrimental effect on the fission rate. In addition, derivatives of the small-molecule semiconductor thiophene-capped diketopyrrolopyrrole (TDPP) display strong π -stacking, yet analysis of transient absorption data indicates that an intermediate state hypothesized to be $(TT)^1$ forms on a time scale between 1 and 16 ps and that the triplet pairs dissociate on a time scale between 20 and 1600 ps [122]. These time scales are slow compared to those for pentacene films (100 fs) [60], even though pentacenes herringbone structure leads to weaker electronic coupling [99]. Could strong orbital interaction hinder singlet fission? Is the correlation between strong coupling and slower triplet

generation/separation a general effect, or is it specific to TDPP?

Experimental knowledge of the correlated triplet pair is limited. In pentacene and hexacene, photoemission spectroscopy has been used to track the dynamics of the triplet pair and show that it can form concurrently with the singlet [123, 107]. Some of the most insightful experiments have been performed on tetracene and TIPS-tetracene, because they exhibit slow fission and long triplet lifetimes. These longer lifetimes make it possible to track oscillations between magnetic sublevels using delayed fluorescence [124, 125], or to measure triplet interactions and coherence using electron spin resonance [126]. These experiments reveal that triplet pairs can remain proximal and coherent for a long time (10s to 100s of ns), and reveal two types of triplet-triplet interaction: a weak magnetic dipolar regime [124], and a strong exchange-coupled regime [126]. Similarly, spin correlations and triplet dissociation have been observed in isolated pentacene dimers [127], although the electronic structure and dynamics of such dimers can differ from those in films, where wave functions extend beyond a dimer of molecules.

Ultrafast optical observations of the triplet pair have been limited to using global analysis on bulk transient absorption (TA) data to argue for the existence of an intermediate state between the initially generated singlets and long-lived triplets. Such experiments reveal the transient absorption spectrum of the triplet pair, and provide some insight into the kinetics, although the models used thus far have been simplified to:



This experimental process has been carried out for crystalline tetracene [128], TIPS-tetracene in solution [129] and crystalline form [91], colloidal aggregates of pentacene derivatives [130], and polycrystalline films of TDPP derivatives [122]. A recently published work studied amorphous films of several organic semiconductors, and, using global analysis of temperature dependent TA data, found evidence of a bound correlated triplet pair [88]. Overall, our understanding of the interactions between correlated triplets remains limited, particularly in systems in which the triplets form on ultrafast time scales.

4.1.2 Background on TIPS-Pentacene

One prototypical π -stacked system that could help shed light on these questions is crystalline 6,13-bis(triisopropylsilylethynyl)pentacene, or TIPS-pentacene (TIPS-Pn). The structure of molecular TIPS-Pn is shown in Figure 4.1a, and Figure 4.1c,d shows the crystalline form of a thin film from two angles: looking top down and from the side. The molecular structure is a derivative of pentacene with two carbons in the middle ring attached to triisopropylsilylethynyl (TIPS) groups—that is, two double bonded carbons (ethynyl) attached to a silicon atom, which itself has three isopropyl groups coming off of it. The TIPS groups make the molecule more soluble and less prone to damage than ordinary pentacene [131]. They also lead to the packing structure that is visible in Figure 4.1c,d—individual molecules are standing with their TIPS groups almost vertical, but the molecule is leaning slightly.

Looking from the top-down the molecules pack in a brickwork pattern, leading to significant π -stacking and charge transport along the (110) direction. Intermolecular interactions in the out-of-plane direction (c-axis), by contrast are negligible [132]. The $S_0 \rightarrow S_1$ transition is known to lie along the short axis of the pentacene core (red arrow in Figure 4.1c) [133], whereas several other transitions will be found in this chapter to lie along the long axis of the pentacene core (blue arrow).




Figure 4.1: Structure of TIPS-Pentacene. (a) Molecular structure [9], (b) optical image of a crystalline thin film with an $8 \mu\text{m}$ laser spot marked in red and a uniform area inside the blue dashed lines, and crystalline packing structure looking (c) top-down and (d) side-on. Crystal axes are indicated, as are TDMs along the short (red) and long (blue) axis the pentacene core.

Thanks to the TIPS side groups and the strong π -stacking interaction, making crystalline TIPS-Pn is relatively straight-forward. TIPS-Pn was purchased as powder from Sigma-Aldrich, dissolved in toluene to a concentration of 5 mg/mL , and passed through a $0.45 \mu\text{m}$ filter. Approximately 140 L of this solution was deposited onto a glass cover-slip that had been sonicated in both acetone and isopropyl alcohol and then treated with trichloro(phenethyl)silane to make the solution wet the substrate. The substrate was heated to 55 C and covered as the sample dried. Crystalline domains were located by inspection. An example is shown in Figure 4.1b. Notice that the thickness appears to be uniform over an area approximately $100 \mu\text{m}$ in extent. It should be noted that many other methods of depositing TIPS-Pn have been explored in order to produce more reliable films [90, 134].

TIPS-Pn is known to undergo singlet fission [135, 136], but there has been controversy surrounding the time scale: for amorphous films, it has been reported to occur on both a 100 fs time scale [137] and a 1 ps time scale [138]. Nanoparticles have been studied as well, but the situation becomes murkier. Heterogeneity and diffusion between different regions of a nanoparticle have been implicated as leading to multiple fission time scales: 220 fs and 1.37 ps in one work [139], and 2.9 and 169 ps in another [140]. More recent work asserts that the time scale is 100 fs in crystalline and 1.2 ps in amorphous nanoparticles [141]. Interestingly, it has been asserted that increasing crystallinity leads to both more [141] and less [140] efficient fission, a discrepancy that has been attributed to different sample preparation protocols. Finally, TAM-based diffusion measurements in a thin crystalline film of TIPS-Pn have concluded that the fission timescale is a relatively slow 5 ps. [80]. Clearly, further study is needed to understand the fission time scale and how it is controlled by crystallinity, which for TIPS-Pn goes hand in hand with π -stacking.

4.2 Kinetic Model of TAM Data

4.2.1 Overview of Data




Figure 4.2: Absorption spectrum of TIPS-Pn (dark blue), pump laser spectrum (red), and non-degenerate probe wavelengths used in this study (vertical bars).

In order to learn more about singlet fission in TIPS-Pn we use TAM to study three different single crystalline domains (most of the data shown are from “Domain 1”). The salient optical points are summarized in Figure 4.2. The absorption spectrum of TIPS-Pn is shown in dark blue and the pump laser spectrum is shown in red. The absorption spectrum shows vibronic structure, but we are primarily exciting the band-edge transition at 700 nm. For the probe we use a degenerate copy of the pump (stage 1 of the TA microscope discussed in Section 2.3), and also use a broadband pulse that is filtered to produce 12 probes ranging from 486 nm to 730 nm, indicated by rectangles in Figure 4.2 (stage 2 of the TA microscope).

For each probe wavelength λ we measure ΔT and T at a range of probe polarizations θ and delay times τ .

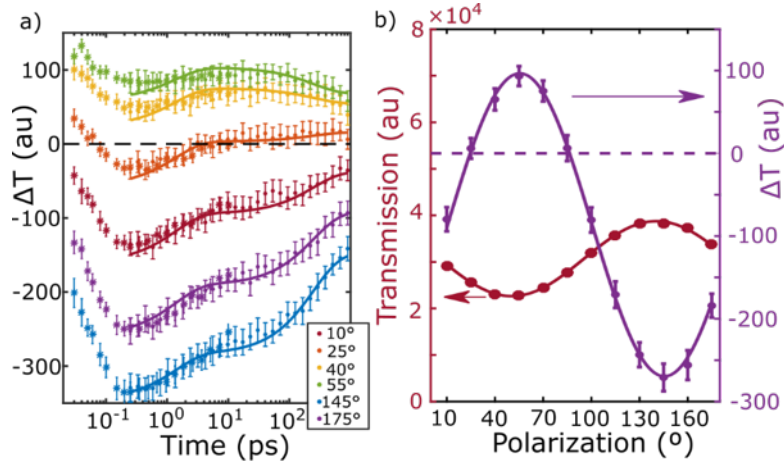


Figure 4.3: Degenerate polarization-resolved TA of a single crystalline domain of TIPS-Pn, represented by (a) fixing probe polarization and plotting the signal over time, and (b) fixing the delay time and plotting the signal vs. polarization. Curves are fits to the kinetic model. The polarization-resolved transmission is also shown, in red.

To begin, consider the TA data taken with a degenerate probe, $\lambda = 700$ nm. Data are visualized in Figure 4.3 using two techniques. In Figure 4.3a, each of the curves represents the excited state time evolution at a different, specific, fixed θ . Curves are fits to the kinetic model, which will be derived in this section. Due to stability issues the data were collected in two batches—one for short times (30 fs to 10 ps) and one for longer times (250 fs to 900 ps), but they can be seen to match up nicely. We find three clear time scales for all curves: the signal drops over 100 fs, rises over 1 ps, and tends toward zero over several hundred ps, consistent with previous measurements [138]. The polarization resolution of TAM permits us to observe both GSB and ESA, as evidenced by the fact that we see curves with both positive and negative values of ΔT . To more clearly see the balance between GSB and ESA, we fix τ at 10 ps and plot ΔT as a function of θ (Figure 4.3b, purple). The curve is fit to

$$\Delta T(\theta) = A \cos^2(\theta - \theta_0) + C, \quad (4.1)$$

which is the functional form for a sum of transition dipole moments (TDMs).

Figure 4.3b also shows the linear transmission as a function of polarization (red). Its minimum indicates the orientation of the GSB TDM, which is along the short axis of the pentacene core [133]. Because $\Delta T < 0$ when the transmission is high (absorption is low), we deduce that the GSB and ESA TDMs are not parallel (see Figure 4.1c). Both ΔT and T are reported in the same units, allowing one to see that the fractional change in optical density is small—in this case on the order of 0.01

4.2.2 Power Dependence




Figure 4.4: Pump power dependence of TA signal at several delay times for (a) Domain 1 and (c) Domain 3. Pump powers are chosen so that we remain in the linear regime, even at long delay times.

For our subsequent analysis, it is crucial that all data are collected in a regime free of higher order nonlinear processes. To ensure that this is the case, we measure $\Delta T/T$ as a function of pump fluence for several delay times, and plot the results in Figure 4.4a for Domain 1 (being discussed here) and Figure 4.4b for Domain 3. At higher fluence (higher exciton density), the signal saturates, particularly for $\tau \geq 300$ ps. This saturation indicates that excitons created in different parts of the sample are interacting due to diffusion. We therefore use a pump fluence around $270 \mu\text{J}/\text{cm}^2$ for Domain 1, indicated by the vertical dotted black line in Figure 4.4a, to remain in the linear regime for all of our measurements. For Domains 2 and 3, peak fluences of 273 and $286 \mu\text{J}/\text{cm}^2$ were used, as indicated by the dashed and dotted lines in Figure 4.4b.

These fluences may seem high compared to those used when working with samples of randomly oriented TIPS-Pn molecules, but in its crystalline form the $S_0 \rightarrow S_1$ TDM is oriented nearly normal to the substrate, as shown in Figure 4.1d and couples only weakly to the incoming light. We therefore find modest exciton densities. For example, a peak fluence of $270 \mu\text{J}/\text{cm}^2$ at a pump repetition rate of 2.5 kHz corresponds to a peak intensity of $0.007 \mu\text{W}/\mu\text{m}^2$. Converting intensity to exciton density requires a proportionality constant that measures how efficiently pump light is absorbed. This number is calculated in Section 4.2.7, but for now we use a typical result of $1 \mu\text{m}^2/\mu\text{W}$, leading to an exciton density of 0.7% .

Even at 0.7% exciton density one might still expect inter-exciton interaction within the first 900 ps. However there are two factors that limit the ability of excitons in TIPS-Pn

to interact. First, there is the fact that the exciton delocalization and transport are both quasi-one dimensional. The lattice spacing along the π -stacking direction is 8 Å[142], hence the average nearest-neighbor distance $\frac{1}{2} \cdot \frac{8 \text{ \AA}}{0.007} = 57 \text{ nm}$. Furthermore, as we shall see in Section 4.2.7 the exciton spends most of its time as a slowly-moving triplet. Therefore the pump fluences used put us in a linear regime, and we can consider each exciton as evolving independently of the others.

4.2.3 Normalized Offset

To better understand the underlying dynamics from the TAM data, we find it helpful to define the *normalized offset*, $\zeta = \frac{C + A/2}{|A|/2}$, similar to what we did in Chapter 3. For fixed τ , consider the plot of ΔT vs θ and the functional form that describes it in Equation 4.1. The numerator of ζ is the value of ΔT averaged over θ ; the denominator serves to normalize it by the half-amplitude of ΔT 's cosine-squared functional form. This procedure is shown intuitively in Figure 4.5a. Positive values of ζ mean that GSB is the dominant contribution to ΔT , while negative values indicate that ESA is the dominant contribution. The most important quality of the normalized offset is that it is insensitive to the absolute excited state population (assuming measurement in the linear regime). If all excited populations decay uniformly, then all GSB and ESA signals decay uniformly. Yet, because it is normalized, ζ will be invariant in time over the course of the uniform decay. A shift in ζ can therefore be used to infer a change in the makeup of the excited state population.

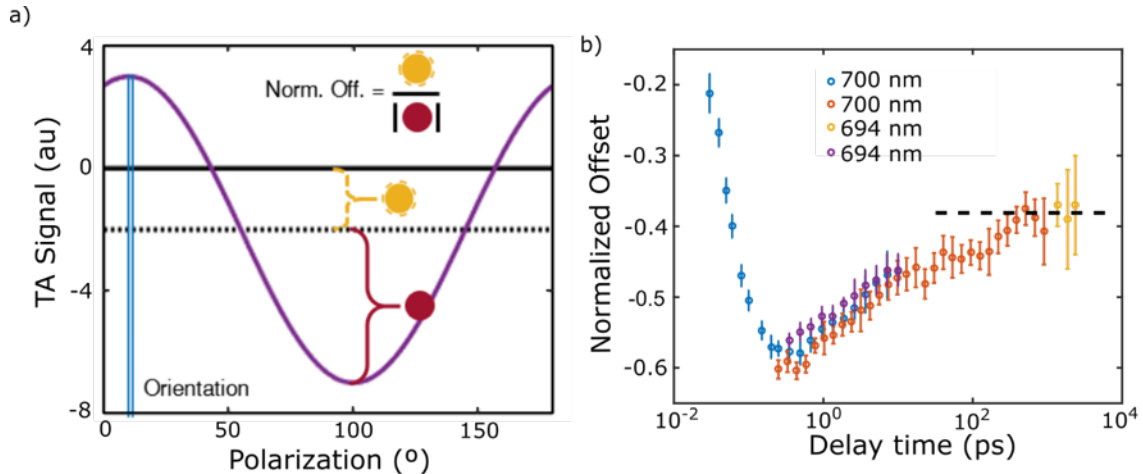


Figure 4.5: Explanation and example of normalized offset, ζ . (a) Cartoon depicting how ζ is calculated. (b) Measured ζ in TIPS-Pn for four data sets, two with a degenerate 700 nm probe and two with a 694 nm probe. The dashed line is a guide to indicate the plateau of ζ at long times.

As an example, we plot ζ as a function of τ in Figure 4.5b for four different scans: two degenerate ones, and two using white light filtered at 694 nm (which should closely reproduce the degenerate data). The scans agree within experimental error, indicating that normalized offset is a physically meaningful attribute, independent of changes in laser beam overlap or other experimental conditions. We observe three population shifts in the data. The first two, occurring over approximately 100 fs and 1 ps time scales, correspond to previously measured decays. At longer time delays, ζ is roughly constant until it undergoes a final shift with a 330 ps time scale, whereupon it is stable until at least 2.5 ns. The shift is statistically significant: the difference between the average signals for time points before (40–125 ps) and after (500 ps–2.5 ns) the shift is 3.8 standard deviations. Because long-lived triplets are known to be produced in TIPS-Pn films, we conclude that the excited state population beyond 500 ps delay consists almost exclusively of dissociated triplets. ΔT at long times is therefore due to two TDMs: one associated with GSB on the $S_0 \rightarrow S_1$ transition, indicative of how many triplet excitons exist, and another associated with the ESA from the triplet excitons. We plot ΔT vs θ at $\tau = 1.4$ ns, and fit to the form given by

$$\Delta T(\theta) = A_{GSB} \cos^2(\theta - \theta_{GSB}) + A_{T-ESA} \cos^2(\theta - \theta_{T-ESA}), \quad (4.2)$$

where θ_{GSB} is fixed based on the polarized transmission measurement (red curve in Figure 4.3b). The data and fit are shown in Figure 4.6. Based on the fit we extract quantities θ_{T-ESA} and A_{T-ESA}/A_{GSB} , both of which provide further information about the dissociated triplet absorption and will be crucial to the subsequent kinetic modeling.




Figure 4.6: TA signal vs. probe polarization at $\tau = 1.4$ ns and $\lambda = 694$ nm, fit to a sum of two TDMs (Equation 4.2).

4.2.4 Assigning Singlet Fission

Our central challenge in producing a kinetic model of the excited state dynamics in TIPS-Pentacene is to determine when singlet fission occurs. Recall from Figure 4.3a that there are two initial processes, one that causes the TA signal to drop over ~ 100 fs and another that causes it to rise over ~ 1 ps. One of these processes should be fission, but it has been unclear which one (see Section 4.1.2). We conclude that the 1 ps process corresponds to the formation of the correlated triplet pair, and the reasoning is explained below.

Several papers assert that the initial decay in the TA signal of TIPS-Pn is due to singlet fission, and that this process occurs with a yield approaching 200% [108, 98, 137]. If this were the case, then it is difficult to explain why our TA signal continues to evolve past 200 fs. We are in an annihilation-free regime for the entire experiment (Figure 4.4), but population loss is clearly occurring, and by studying the normalized offset we know that the makeup of the population is changing. Isolated triplets in a homogeneous region should live for much longer than 1 ns. Triplet pairs could undergo geminate annihilation, but this would not explain the shift in the normalized offset. The dissociation of the triplet pair could lead to a small change in the signal, but its hard to imagine what processes could be occurring so as to yield changes of this magnitude on multiple timescales and at all polarizations.

The clearest evidence that the 100 fs process cannot be singlet fission, however, comes from considering the early time increase in ESA. As evidenced by the signal when the probe polarization is orthogonal to the GSB TDM ($\theta = 145^\circ$), the ESA becomes much stronger during the first 100 fs. If it were due to singlet fission, this would imply that the triplet state has a relatively strong ESA at 700 nm. However, in Figure 4.6 we plot the TA signal at $\tau = 1.4$ ns, which is due entirely to dissociated triplets. We therefore extract the strength of the triplet ESA relative to GSB, as well as the orientation of the triplet ESA TDM, finding values of 2.19 ± 0.08 and $146.9^\circ \pm 0.7^\circ$, respectively. These values are not strong enough to explain the observed ultrafast drop at $\theta = 145^\circ$ as being due to triplet formation. To see this, we fit the TA data to a kinetic model in which triplet pair formation is the first process. Although the correlated triplet pair and dissociated triplets do not have identical optical properties, they should be similar enough for us to draw conclusions. The results are shown in Figure 4.7, and they are clear: the model does not fit the data well. Explaining the initial dynamics in TIPS-Pentacene as being due to singlet fission would require the triplet ESA to be much stronger than we have measured it to be. Note that the polarization dependence of TAM is crucial to reaching this conclusion.

It has also been suggested that both the 100 fs and the 1 ps process could correspond to singlet fission, particularly in disordered nanoparticles [139]. The idea is that disorder might lead to a distribution of intermolecular couplings, which show up as a distribution of fission rates, and presents as a biexponential. Yet we are studying a crystalline sample. Even if there is some disorder it should be much less than nanoparticles, yet we see the same two timescales in roughly the same proportions. Furthermore, thanks to the polarization-dependence of our measurement, we can see that the ~ 100 fs process and ~ 1 ps process cause the TA signal to change in opposite ways. Hence, they cannot both be due to singlet




Figure 4.7: Attempted fit of TA data to an ultrafast singlet fission model, fixing the relative strength and orientation of the triplet ESA TDM. The model is incapable of fitting the data.

fission.

Having established that 100 fs singlet fission is incompatible with our data, I now address the evidence that has led others to draw the opposite conclusion. Two theoretical studies [98, 143] have predicted that singlet fission in TIPS-Pn should occur on a 100 fs timescale. Both find that strong coupling between neighboring molecules (due to the π -stacking geometry) drives rapid fission. Recent studies have, however, indicated that making the excitonic coupling stronger [121] or increasing triplet delocalization [100] may both have detrimental effects, causing the fission rate to decrease and promoting triplet fusion. Given that another strongly π -stacked molecule has recently been found to undergo slower (ps scale) fission, [122] this is an intriguing possibility. One of the theoretical studies [98] cites a measurement of the fission rate. They used TA to observe a decay with a time constant of 110 fs. Yet, no argument is advanced as to why this decay is due to singlet fission. Two studies have tracked vibrational modes in TIPS-Pn on ultrafast timescales, finding that vibrational coherence persists throughout the initial ultrafast process [137], and that vibrations are crucial for coupling different electronic states [108]. Both papers assume singlet fission is behind their observations, but there does not seem to be anything about their results which proves that this is so. In fact, the authors of one paper note how their observations are reminiscent of nuclear wavepackets undergoing internal conversion between electronic states [108]. Their results are strongly indicative of a conical intersection (an intersection between two potential energy surfaces that allows non-adiabatic transitions between states to occur), but it is not necessarily an intersection between the singlet and triplet manifolds. Finally, test devices have been made in which PbSe acceptors tuned to the energy of TIPS-Pn triplets are able to harvest excitons with internal quantum efficiencies of up to $160 \pm 40\%$ [135]. This is perfectly consistent with our results. 1 ps fission is still faster than most decay processes, and is capable of producing very high triplet yields (see Section 4.5).

4.2.5 Assigning Other Timescales to Complete the Kinetic Model

If the fast initial process is not singlet fission, we must advance some alternative explanation. One possibility is that excess energy in the pump excites intramolecular vibrations, which relax over the first 100 fs. This process could change the singlet ESA, and it would cause a loss of stimulated emission at the pump wavelength, which would explain initial drop in signal for polarizations aligned with the $S_0 \rightarrow S_1$ TDM (such as $\theta = 55^\circ$ in Figure 4.3a). Intramolecular vibrations can have periods as short as 10 fs, and sub-100 fs intraband relaxation has been observed in tetracene and rubrene [44], so this is a plausible assignment. Furthermore, the lack of fluorescence observed from TIPS-Pn [45] would be explicable if the vibrational relaxation changes the molecular structure in such a way as to weaken optical coupling to the ground state.

For more evidence that the initial process in TIPS-Pn is due to thermalization, consider the fact that experiments done with shorter pump wavelengths tend to result in a longer measured timescale. Higher energy pump photons would mean the exciton had more excess energy to release before it is thermalized, which could take longer. Conversely, if the ultrafast process were singlet fission, we might expect higher energy pump photons to speed up the process, if anything. A literature survey is summarized in Table 4.1. The primary trend is for longer wavelengths to produce faster timescales, although there is a prominent outlier in Reference [137]. Examining the Supporting Information of Ref. [137], the value of 80 fs comes from probing at $\lambda = 810$ nm. The authors also show data from $\lambda = 525$ nm. Although the accompanying fit is not labeled, we extracted some data points and found it fits well to a biexponential with time constants 129 fs and 1.7 ps.

Pump wavelength (nm)	First timescale (fs)	Source
530 nm (possibly 590 nm)	220 fs	[139]
500 nm - 650 nm	80 fs (also 129 fs)	[137]
550 nm - 650 nm	110 fs	[98]
600 nm	120 fs	[144]
700 nm	50 fs	[71]

Table 4.1: Measured first timescale in TIPS-Pn, arranged in decreasing order of pump photon energy.

Having established that triplet pairs and singlets equilibrate over the 1 ps timescale, we consider the longer-time dynamics. The TA signal continues to decay (Figure 4.3a) and the normalized offset undergoes a later shift (Figure 4.5). Hence we deduce that unidirectional triplet pair formation is *not* occurring. Instead we find that correlated triplet pairs are being formed, and that these correlated triplet pairs have the ability to re-form singlets. An equilibrium will therefore be established between the vibrationally relaxed singlet and the correlated triplet pair. The loss of TA signal strength over the subsequent time indicates

the existence of a decay channel. Because we work in a linear regime with respect to pump power, this decay cannot be diffusion-based recombination. Fluorescence is weak in TIPS-Pn [45], so we conclude that this decay is non-radiative internal conversion originating from either the singlet or correlated triplet. Despite this decay pathway we observed long-lived triplets, meaning that the correlated triplet pair must be able to dissociate. This dissociation is what gives rise to the shift in normalized offset around 330 ps. We start with a population of singlets and triplets in equilibrium, but the ability of the triplets to dissociate means that the population eventually shifts to all triplets. Singlets and triplets have different ESA, so the GSB-to-ESA balance shifts, and hence ζ shifts as well. The fact that this shift and hence dissociation takes place over hundreds of picoseconds implies that the correlated triplet pair is bound.




Figure 4.8: A Jablonski diagram of the kinetic model for TIPS-Pn

Putting it all together, our kinetic model is shown in Figure 4.8. The pump (red arrow) excites from the ground state S_0 to the vibrationally excited singlet S_1^n . It thermalizes to the vibrationally relaxed singlet S_1^0 with rate k_{th} . The singlet and correlated triplet pair $(TT)^1$ interconvert via singlet fission and triplet fusion, k_{sf} and k_{tf} . Either of these states internally convert with rate k_{int} (we cannot distinguish between the two because the equilibration timescale is much faster than the internal conversion timescale). The correlated triplet pair dissociates to spatially separated but still spin correlated triplets $T+T$ with rate k_{dis} .

Before going on, I will address and rule out some other explanations for our observation that the TA signal strength and normalized offset continue to change after singlet fission has occurred. With these pathways non-viable, the kinetic model put forth above becomes the most sensible option. First, consider the possibility that singlet fission only occurs at certain sites. If that were the case then the remaining singlets could evolve in a way that produced the observed shifts. However this is unlikely because we are probing a highly crystalline and uniform sample, as evidence by the fact that the TA signal has an extremum of $\Delta T(\theta) = 0$ at several different wavelengths (see next Section). In a uniform sample, it is unclear why certain spots would be incapable of promoting fission. Even if such spots did exist, the singlet excitons would rapidly diffuse away from them given the high diffusivity of singlets [80], we

can expect one to explore a region of roughly 50 nm radius within 10 ps. Since the decay of ΔT does not start occurring until about 100 ps, this should be enough time for a given exciton to find a spot suitable for fission, and we would not observe significant population loss. It is possible that the no-fission zones are large (order of hundreds of nm) or consist of entire planes of TIPS-Pn (since out-of-plane diffusion is slow). However if this were the case, we would expect significant heterogeneity among our micron-sized sample spots. Our previous study of TIPS-Pn investigated many domains and found them to all be similar [71]. Therefore we conclude that triplet formation occurs uniformly.

Another possibility is that triplet pairs form and dissociate within a few picoseconds, and then diffuse to undergo non-geminate recombination. However non-geminate recombination would be a non-linear process whose rate increases for increasing pump fluence. Because we are in a linear regime (Figure 4.4), this cannot be the case.

Other processes to consider include intersystem crossing, diffusion out of the pump volume, triplet quenching due to oxygen, and hopping to trap states. Intersystem crossing in TIPS-Pn would be much slower than 1 ns, and was not observed in solution [136], so we do not include it. Diffusion is insignificant for our pump size (about 8 μm) and timescale considered, given the diffusivity of excitons in TIPS-Pn is around 0.1 nm²/ns [80]. Oxygen is irrelevant because the TIPS side groups protect the molecule from oxidation [131], and we have found that films up to 3 years old are still very optically active. Finally, we would expect that trap states display TDMs that do not necessarily align with those of the bulk crystal, and hence if significant trapping were occurring, the normalized offset would shift. No such behavior is observed.

In conclusion, we have argued that the 100 fs timescale is incompatible with triplet formation, and is probably thermalization. The 1 ps timescale must therefore be triplet formation. The fact that the TA signal continues to decay past 100 ps indicates a decay pathway. After ruling out several other options, the most likely candidate is internal conversion. The normalized offset shifts over about 330 ps, indicating that there was originally an equilibrium between singlet and correlated triplets, but triplet dissociation disrupts this equilibrium. This leads to the kinetic model shown in Figure 4.8, which we will subsequently see fits our data at all wavelengths quite nicely. Note that polarization-resolved measurement and the ensuing normalized offset were crucial to drawing these conclusions.

4.2.6 Non-degenerate Probes

Before fitting the kinetic model, consider data collected with non-degenerate probe wavelengths. Figure 4.9 shows data at $\lambda = 520$ nm, 550 nm, and 730 nm, along with fits. Because we are interested in fission, only the first 10 ps of dynamics are shown; each trace therefore shows a single decay, though the amplitude of the decay and its polarization dependence vary depending on how the probe couples to each TDM. We observe that triplet formation tends to make ΔT more negative as a function of time at bluer probe wavelengths (such as 520 nm) and more positive in time at redder wavelengths (such as 730 nm). At 550 nm we observe a mixture of both behaviors. Unlike at redder wavelengths (such as in Figure 4.3), in

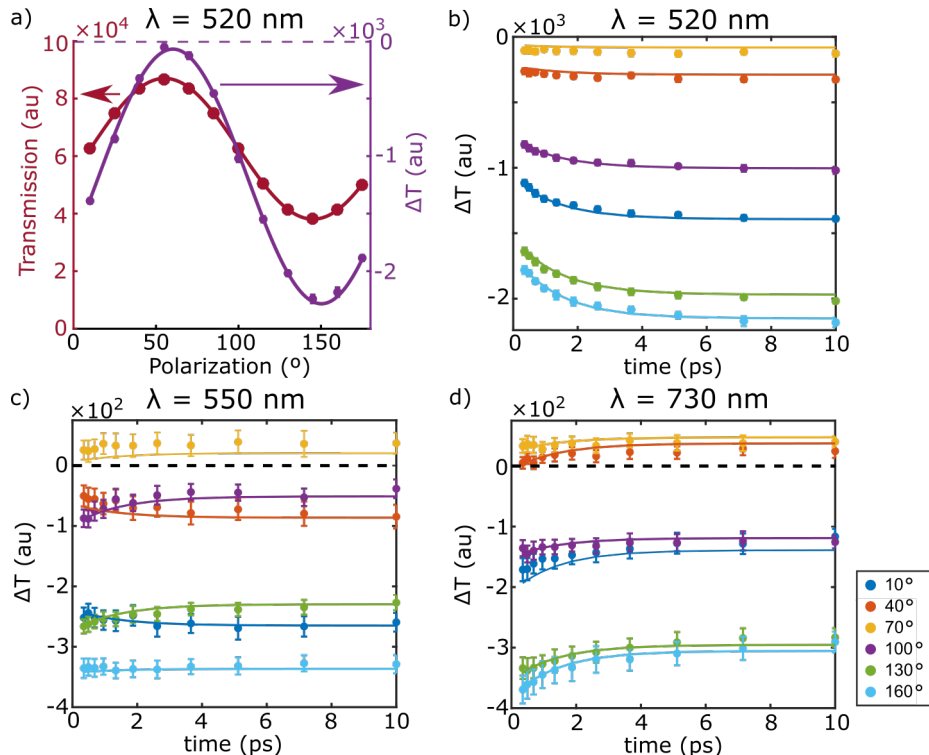


Figure 4.9: Polarization-resolved TAM and fits with selected non-degenerate probes

the bluer region of the spectrum ($\lambda \leq 550$ nm) the GSB and ESA TDMs are nearly aligned. This can be seen in Figure 4.9a, where we plot ΔT vs θ for $\lambda = 520$ nm and compare it to the linear transmission. The highest transmission (weak or possibly no GSB) occurs around $\theta = 55^\circ$, where the TA signal returns to 0 (weak or possibly no ESA). This fact indicates that all of the relevant TDMs in the sample are aligned, and the probe couples to none of them, which is strong corroborating evidence for the uniform crystallinity of the sample. Parallel TDMs, however, make it more challenging to separate the contributions of GSB and ESA, but we develop analytic techniques in the next section to surmount the difficulty.

The remaining probe wavelengths are shown in Figure 4.10. For $\lambda = 600$ nm, the weak signal makes time slices difficult to visually interpret, so the data are shown as polarization slices. Figure 4.11 shows the TA signal out to $\tau = 900$ ps for a fixed probe polarization. These data and fit will be crucial to extracting the change in ESA strength upon triplet dissociation.

4.2.7 Fitting the Kinetic Model

Since we have established the kinetic model in Figure 4.8, we can fit it to the TAM data, $\Delta T(\tau, \theta, \lambda)$, at the non-degenerate probe wavelengths in Figure 4.9 and 4.10. For simplicity we start the model at $\tau = 250$ fs, and assume that thermalization is complete. The initial

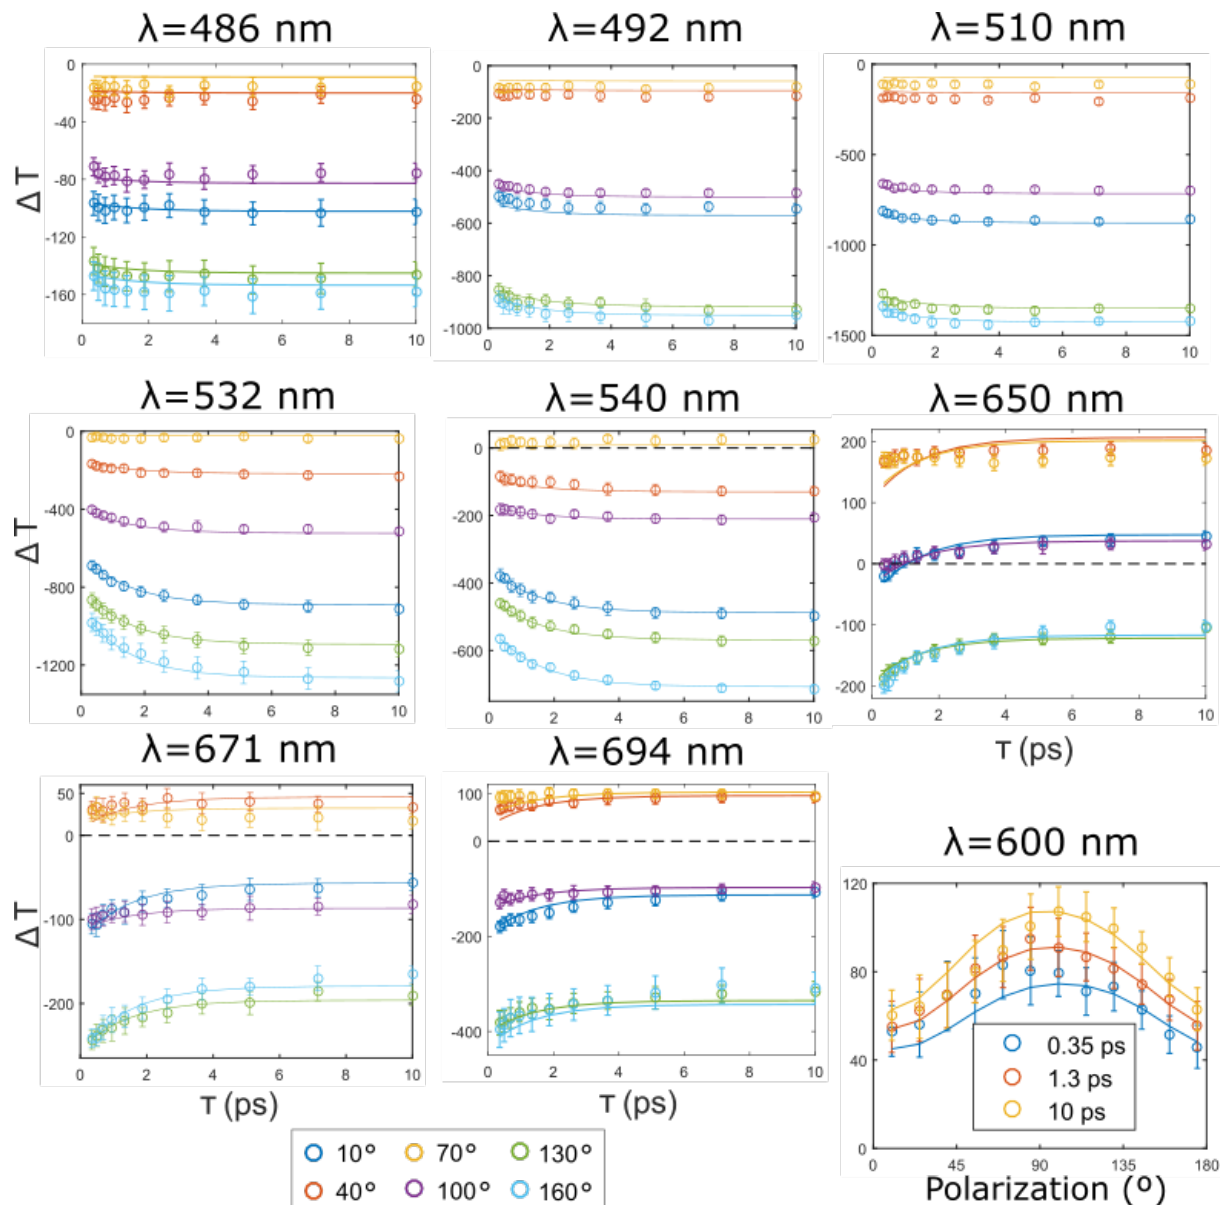


Figure 4.10: Polarization-resolved TAM and fits with all non-degenerate probes

state is therefore all excitons in the thermalized singlet excited state, or $P_S = 1$. Other variables are the population of the correlated triplet, P_{TT} , and population of the isolated triplet, P_T . The equations of the model are therefore given by:




Figure 4.11: Long time TA signal and fits for two selected probe wavelengths

$$\begin{aligned}
 \frac{dP_S}{dt} &= -k_{sf}P_S - k_{int}P_S + k_{tf}P_{TT} \\
 \frac{dP_{TT}}{dt} &= k_{sf}P_S - k_{tf}P_{TT} - k_{dis}P_{TT} \\
 \frac{dP_T}{dt} &= 2k_{dis}P_{TT}.
 \end{aligned} \tag{4.3}$$

These equations can be put in matrix form and solved as in Section 3.2.3. Once we determine the populations, we calculate the transient absorption signal. Let $A_G(\theta, \lambda)$ be the amplitude for absorption from the ground state for a given probe polarization and wavelength. The absorptions from the excited states are $A_S(\theta, \lambda)$, $A_{TT}(\theta, \lambda)$, and $A_T(\theta, \lambda)$. The ESA signals are proportional to the population of each excited state, while the GSB is proportional to the number of ground state vacancies. The overall transient absorption signal, ΔT , is given by

$$\begin{aligned}
 \Delta T(\tau, \theta, \lambda) &= A_G(\theta, \lambda)(P_S(\tau) + 2P_{TT}(\tau) + P_T(\tau)) \\
 &\quad - A_S(\theta, \lambda)P_S(\tau) - 2A_{TT}(\theta, \lambda)P_{TT}(\tau) - A_T(\theta, \lambda)P_T(\tau).
 \end{aligned} \tag{4.4}$$

Each absorption is generally considered to be due to one TDM, and hence can be modeled with

$$A_X(\theta, \lambda) = A_X(\lambda) \cos^2(\theta - \theta_X(\lambda)), \tag{4.5}$$

where X equals S , TT , or T .

For a given wavelength, the data are fit to this model in order to extract all free parameters. The value of $\theta_G(\lambda)$ is fixed using the probes linear transmission. All rates are found

initially at $\lambda = 700$ nm, but because the dynamics are independent of the probe pulse, we use the same rate constants for subsequent fits at other probe wavelengths. The ground state absorption strength, $A_G(\lambda)$ is fixed to a value calculated using linear absorption data. The ESA signals are slightly modified by a shape correction factor. At $\lambda = 600$ nm, 650 nm and 671 nm, we also allow the ground state bleach to have a constant offset. All of these modifications are discussed later in this section.

Parameterization of Singlet Fission

Although the kinetic model returns values for rate constants, it is more intuitive to report the following quantities. First, we report the internal conversion and triplet dissociation timescales, which are the inverses of their respective rates: $\tau_{int} = 1/k_{int}$, and $\tau_{dis} = 1/k_{dis}$. Next consider singlet fission and triplet fusion. Ignoring the longer processes, equilibrium forms between excited singlets and correlated triplet pairs with rate constant $k_{te} = k_{sf} + k_{tf}$. The timescale for triplet equilibration is therefore $\tau_{te} = \frac{1}{k_{sf} + k_{tf}}$. This is the timescale on which the TA signal changes.

At equilibrium, detailed balance requires $P_{TT}k_{tf} = P_Sk_{sf}$. If f is the fraction of excitons that form triplet pairs at equilibrium, then $P_{TT} = f$, $P_S = 1 - f$, and we can solve to find $f = \frac{k_{sf}}{k_{tf} + k_{sf}}$. This is a measure of how strongly triplet pairs are favored at equilibrium, but is not to be confused with the triplet generation internal quantum efficiency, η , which is the ultimate number of dissociated triplets produced at long times for each initial singlet exciton. In this model, $\eta = P_T(\tau \rightarrow \infty)$.

We solve for η by making the approximation that the singlet and correlated triplet pair are always in equilibrium. Because internal conversion and dissociation are much slower than fission/fusion, this approximation is valid. Let $p(t)$ be the total singlet/triplet pair population. Then $P_{TT}(t) = fp(t)$ and $P_S(t) = (1 - f)p(t)$. In a short time Δt , a small amount of triplet pairs dissociate and a small amount of singlets undergo internal conversion. The change in triplet population is $\Delta P_T = 2k_{dis}fp(t)\Delta t$, while the change in ground state population is $\Delta P_{S0} = k_{int}(1 - f)p(t)\Delta t$. Taking the ratio of these two, we obtain a time-independent quantity:

$$\frac{\Delta P_T}{\Delta P_{S0}} = \frac{2k_{dis}f}{k_{int}(1 - f)}.$$

At long times, every initially generated singlet exciton will either return to the ground state or form two triplets. Hence, $p(0) = \Delta P_{S0}(t \rightarrow \infty) + \frac{1}{2}\Delta P_T(t \rightarrow \infty)$. Combining this and the above equation we have, for the triplet generation internal quantum efficiency,

$$\eta = \frac{2}{1 + \tau_{dis}(1 - f)/\tau_{int}f}. \quad (4.6)$$

Results!

domain	τ_{te} (ps)	f	τ_{int} (ps)	τ_{dis} (ps)	η (%)
1	1.4 ± 0.2	0.74 ± 0.03	149 ± 24	330 ± 53	112 ± 14
2	2.0 ± 0.3	0.69 ± 0.04	106 ± 16	360 ± 53	79 ± 13
3	1.6 ± 0.3	0.60 ± 0.04	196 ± 32	240 ± 60	110 ± 17

Table 4.2: Results of kinetic model fit to degenerate probe data (time scales).

Results of fitting the kinetic model to degenerate probe data are shown in Table 4.2. Data shown so far have been for Domain 1 of the three that were studied in depth. Data and fits for Domains 2 and 3 are shown in Figure 4.12. We find that fission/fusion come to equilibrium in 1-2 ps, and that equilibrium favors triplet pairs. Internal conversion and triplet dissociation both occur over hundreds of picoseconds, and their competition results in an efficiency that can be above 100%.

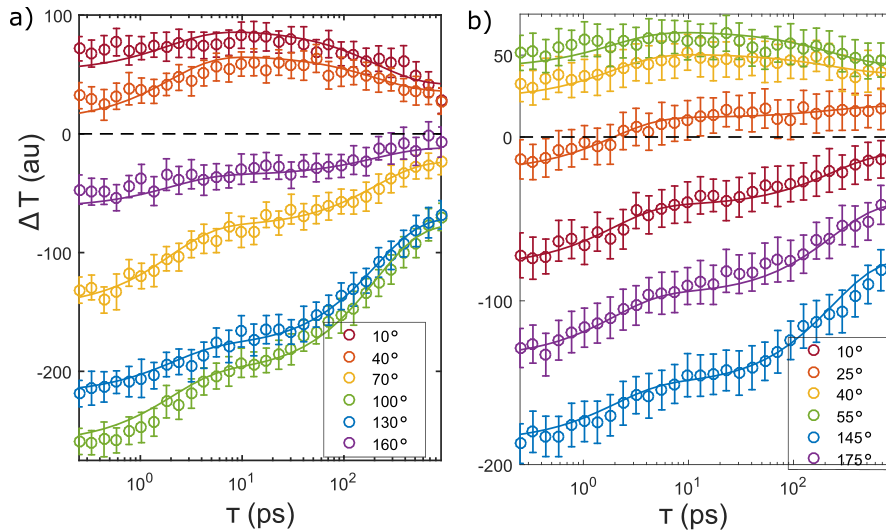


Figure 4.12: Degenerate TA data and kinetic model fits for Domains 2 (a) and 3 (b)

All domains are found to be quite similar, though there are some observed differences, especially in f and τ_{int} , which lead to differences in η . These could be due to differing degrees of lattice strain, as discussed in Chapter 3. Fit curves for Domains 2 and 3 are shown in Figure 4.12.

Table 4.3 shows more kinetic model results—these are the TDMs for GSB and ESA from S , $(TT)^1$, and T . The TDM strengths are reported relative to the GSB strength. Note that the ESAs are similar for Domains 2 and 3, but stronger for Domain 1, indicating non-azimuthal rotation of the crystal structure (rotation that is not purely in the plane of the

Domain	S_1 $\frac{\text{GSB}}{\text{ESA}}$	TT $\frac{\text{GSB}}{\text{ESA}}$	T $\frac{\text{GSB}}{\text{ESA}}$	GSB θ ($^\circ$)	S_1 ESA θ ($^\circ$)	TT ESA θ ($^\circ$)	T ESA θ ($^\circ$)
1	5.47 ± 0.27	2.18 ± 0.14	2.2 ± 0.08	50.5 ± 0.3	149 ± 1	146 ± 1	147 ± 1
2	4.38 ± 0.11	1.72 ± 0.09	1.6 ± 0.05	28.7 ± 0.8	105 ± 1	108 ± 3	108 ± 1
3	4.52 ± 0.18	1.67 ± 0.15	1.6 ± 0.07	46.4 ± 0.3	150 ± 1	147 ± 2	145 ± 3

Table 4.3: Results of kinetic model fit to degenerate probe data (transition dipole moments).

substrate). The correlated and dissociated triplets have similar TDMs, as expected. Finally the ESA TDMs are always about orthogonal to the GSB, although we can see that Domain 2 is flipped upside-down relative to the other domains.

Fixing the kinetic rates found above, we fit the kinetic model to the non-degenerate probe data. The fits are solid lines in Figures 4.9, and 4.10. We find excellent agreement at all probe wavelengths studied, further substantiating the validity of our model. We also extract the orientation and absorbance (projected onto the sample plane) of the ESA TDMs due to both singlets and bound correlated triplets. We measure the ESA TDM due to dissociated triplets by performing TAM with longer time delays and fitting the kinetic model to that data (Figure 4.11). The GSB TDM comes from transmission measurements (see Section 4.2.8). In this way, we independently measure the TDM for absorbance from four states: ground state, singlet exciton, bound triplet exciton, and dissociated triplet exciton.

The absorbance and orientation of the TDMs considered here are plotted in Figure 4.13. Considering the absorbance spectra, we note that the excited states share a prominent feature around 525 nm. We also note that the bound and dissociated triplets display similar, but not entirely identical absorption profiles—the difference between the two is magnified in the inset. As for the orientations, the band edge absorption TDM is known to be oriented along the short axis of the pentacene core [133], which for the purposes of Figure 4.13b we define to be 0° . But this appears to be the exception—most TDMs are oriented roughly along the long axis of the pentacene core. Similar TDM orientations are observed in pentacene [145].

4.2.8 Fine Details of the Kinetic Model

The following section goes into detail regarding all the little pieces of the kinetic model. If you are not trying to re-create our results, feel free to skip to Section 4.3.

Calculating the Bleach Amplitude for Non-degenerate Probes

Because the dynamics depend only on the pump and the pump does not change between different probe wavelengths, we can use the rate constants extracted for the degenerate experiment and apply them to all probe wavelengths. The only free parameters are the TDMs, both their amplitudes and orientations. However we still cannot perform the fit, because for many of the bluer wavelengths the TDMs are all oriented the same way (see the Section 4.2.6 on Non-Degenerate Probes and also Figure 4.1c). This means that we do not obtain as much information as we did when the GSB and ESA were oriented orthogonally

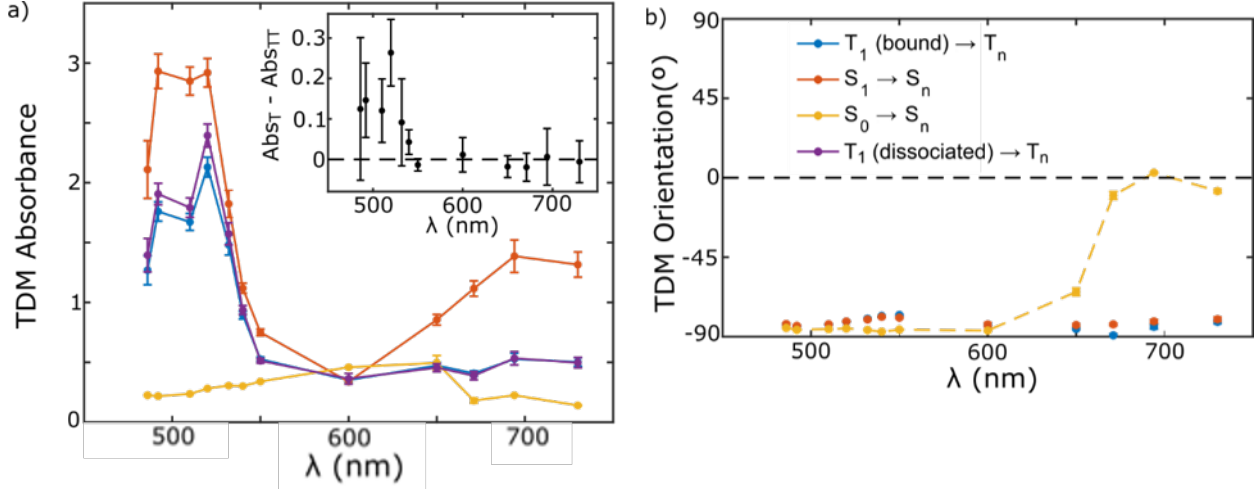


Figure 4.13: Transition dipole moments of TIPS-Pn: (a) absorbance (for unpolarized light) and (b) orientation. Dashed line in (b) is a guide to the eye. Orientations are relative to the absorption dipole at 700 nm.

and we could track them separately. However, we can use the polarization dependence of the probe transmission to extract the probe absorption at each wavelength. That absorption can be used to calculate the amplitude of the GSB TDM. The derivation follows.

Let the probe absorption as measured on our detector be $Abs(\theta)$. If the pump causes some fraction, δ , of the molecules to be excited, then the new probe absorption is $(1 - \delta)Abs(\theta)$, and so the change in transmission is $\Delta T = \delta Abs(\theta)$. Both of these quantities are spatially non-uniform, so we must integrate this product over the sample plane to get the full value of ΔT .

Assume that the pump pulse has a Gaussian mode of characteristic size σ_{pu} and total power P , centered at $r = 0$. The pump power density is given by

$$\frac{P}{\pi\sigma_{pu}^2} e^{-r^2/\sigma_{pu}^2}. \quad (4.7)$$

One can easily verify that integrating this expression over all space yields the total power, P . We assume we are in the linear regime, where the likelihood of creating an exciton is proportional to the pump power. Define a proportionality factor, ξ , which converts from pump power density, Expression 4.7, to fractional exciton density. This number, ξ , depends on the sample thickness and orientation, so it will vary among domains. It is, however, independent of the probe beam, and should be consistent across all experiments within a single domain.

Assume that the probe also has a Gaussian mode, concentric with the pump and with characteristic size σ_{pr} . In the linear regime, the absorption as a function of position is given by

$$\frac{Abs(\theta)}{\pi\sigma_{pr}^2} e^{-r^2/\sigma_{pr}^2}. \quad (4.8)$$

Multiplying Expressions 4.7, 4.8, and ξ , and integrating over all space yields the TA signal due to GSB,

$$\frac{\xi P Abs(\theta)}{\pi(\sigma_{pu}^2 + \sigma_{pr}^2)}. \quad (4.9)$$

What is the form of the linear absorption, $Abs(\theta)$? In simple cases it can be modeled with Equation 4.5, but at some wavelengths there may be multiple non-aligned TDMs contributing to the absorption. Conveniently the sum of cosine squared terms also has the form of a cosine squared, with an offset, so we model the absorption as $A \cos^2(\theta - \theta_0) + C$. However we do not measure absorption directly. We measure transmission, which can be written as

$$T = T_0 - \gamma T_0 - \alpha T_0 \cos^2(\theta - \theta_{GSB}), \quad (4.10)$$

where $\alpha \equiv A/T_0$ and $\gamma \equiv C/T_0$ are the apparent amplitude and offset, normalized to the total transmission.




Figure 4.14: Idealized probe transmission vs. polarization

An idealized transmission curve is shown in Figure 4.14 (compare to real data in Figure 4.3b). Although we may not know T_0 , we directly measure both αT_0 and θ_{GSB} . Plugging the absorption into Equation 4.10, we can write the GSB signal as

$$\Delta T_{GSB} = \frac{\xi P * (\alpha T_0)}{\pi(\sigma_{pr}^2 + \sigma_{pu}^2)} \cos^2(\theta - \theta_{GSB}) + \frac{\xi P * (\gamma T_0)}{\pi(\sigma_{pr}^2 + \sigma_{pu}^2)}. \quad (4.11)$$

The second term is a constant offset to the GSB—it has no θ –dependence. For most wavelengths it is not relevant because $\gamma = 0$, as is discussed later in this section. For wavelengths for which γ is significant it is included as a fitting parameter.

For $\lambda = 694$ nm, 700 nm, and 730 nm, we perform a kinetic model fit to find ΔT_{GSB} . The exponential time constants are fixed to the values obtained from fitting the kinetic model to the degenerate probe polarized TA data, but this fit enables us to determine the parameters associated with the various TDM absorbances at each of these different probe wavelengths. After performing the fit we know the value of almost every term in Equation 4.11: pump power and beam diameters can be directly measured, αT_0 and θ_{GSB} are obtained from the linear transmission, and $\gamma = 0$. This leaves ξ as the only unknown, meaning it can be calculated. **Once ξ is calculated, that value can be used to predict the strength of the GSB at other probe wavelengths.** Fixing the GSB strength and orientation allows the fit to return unique values for the ESA from the singlet and triplet pair.

We calculate ξ using the data taken with probe wavelengths $\lambda = 694$ nm, 700 nm, and 730 nm, and show the results in Table 4.4. We find that there is some variation between domains, but values within the same domain are very similar. This check provides evidence that ξ is a physically meaningful value. Furthermore, we can use ξ to compute the peak fractional excitation density as $P\xi/\pi\sigma_{pu}^2$, and in all three domains the fraction of excitons is around 1%, as asserted earlier.

Probe λ (nm)	ξ , Domain 1 ($\mu\text{m}^2/\mu\text{W}$)	ξ , Domain 2 ($\mu\text{m}^2/\mu\text{W}$)	ξ , Domain 3 ($\mu\text{m}^2/\mu\text{W}$)
694 nm	1.10 ± 0.03	1.64 ± 0.05	1.89 ± 0.12
700 nm	1.00 ± 0.03	–	1.69 ± 0.05
730 nm	0.87 ± 0.05	1.50 ± 0.12	2.0 ± 0.2

Table 4.4: Exciton creation efficiency, ξ , measured for three domains and three probe wavelengths. We do not have reliable data for Domain 2 at 700 nm.

Shape Correction Factor for Excited State Absorption

Next, we calculate a form for the polarization-resolved transient absorption signal due to ESA. For notational simplicity, define $\epsilon \equiv \frac{\xi P}{\pi(\sigma_{pu}^2 + \sigma_{pr}^2)}$, which is an effective fractional exciton density. Assume the excited molecules absorb the probe via a TDM that has absorbance β and orientation θ_{ESA} . The ESA signal will be proportional to ϵ , β , and to the probe transmission. We must, however, remember that the probe intensity is changing as it passes through the sample due to linear absorption. Therefore we cannot use the transmission we measure on the detector, but must consider an average transmission, \bar{T} . the ESA is:

$$\Delta T_{ESA} = -\epsilon \bar{T} \beta \cos^2(\theta - \theta_{ESA}). \quad (4.12)$$

The probe enters the sample at full strength, T_0 , and exits with intensity given by Equation 4.10. As long as we are not in a strongly absorbing regime, the average probe intensity will be the mean of these two values. We combine Equations 4.12 and 4.11 to write the complete form of the TA signal,

$$\Delta T = \epsilon\gamma T_0 + \epsilon\alpha T_0 \cos^2(\theta - \theta_{GSB}) - \epsilon\beta T_0 \left(1 - \frac{\gamma}{2} - \frac{\alpha}{2} \cos^2(\theta - \theta_{GSB})\right) \cos^2(\theta - \theta_{ESA}). \quad (4.13)$$

There are two things we learn from Equation 4.13. First, the ESA term does not have a simple cosine squared dependence, but it is a more complicated function of θ , hence the need for a shape factor to modify the ESA. In the kinetic model the ESA terms are fit to

$$\Delta T_{ESA} = \epsilon\beta T_0 \left(1 - \frac{\gamma}{2}\right) (1 - \alpha' \cos^2(\theta - \theta_{GSB})) \cos^2(\theta - \theta_{ESA}), \quad (4.14)$$

where α' is a free parameter. In principle it could be calculated from the linear absorption, but letting it vary accounts for two experimental non-idealities: excitons are created non-uniformly throughout the sample (due to pump absorption), and the two beams are crossing at an angle, leading to non-uniform interaction.

Based on the above analysis, one can see why we report ΔT in this work, instead of the more conventional $\Delta T/T$. Because ΔT (Equation 4.13) and T (Equation 4.10) have different polarization dependencies, dividing them does not serve to normalize the transient absorption signal. As fluctuations in probe power are not the most significant source of error, we choose to report ΔT for simplicity. Figure 4.15 demonstrates the skewed shape of $\Delta T/T$, juxtaposed with ΔT , which is easier to fit. Including the shape factor, α' , leads to a slight improvement.

When to Include an Additional Offset Term in the Bleach

It is simplest to model absorption features using Equation 4.5, which would be the result for a single TDM. In most cases we find that this is justified. For λ between 486 nm and 550 nm, the fact that the signal peaks at $\Delta T \approx 0$ (seen clearly in Figure 4.9a, purple curve) implies that the probe polarization can be simultaneously orthogonal to all TDMs. Hence the TDMs are all aligned with each other, and we can use Equation 4.5 to describe them.

At $\lambda = 600$ nm, however, the TAM data are qualitatively different (Figure 4.10). ΔT is positive at all polarizations, which is only possible if there are multiple non-aligned GSB TDMs. We therefore use Equation 4.11 for the GSB. This is supported by theoretical calculations, which show that the second lowest energy linear absorption peak contains several contributions with different polarization dependencies [133]. Practically, including γ as a free parameter drastically improves the fit.

We would also expect the second absorption peak is be significant at 650 nm and 671 nm. We find ΔT to be more positive at these wavelengths than we would expect given the linear absorption, so we model the GSB at these wavelengths using Equation 4.9 as well, which improves the fit.




Figure 4.15: Comparison of $\Delta T/T$ (blue) with ΔT (orange) showing the simplicity of the latter. $\Delta T/T$ does not follow a \cos^2 shape (dotted line). ΔT does fit well to \cos^2 (dashed line), and including the shape correction factor (solid line, Equation 4.14) leads to a slight improvement

The linear absorption at $\lambda = 694$ nm, 700 nm and 730 nm is primarily due to the lowest energy peak, which is caused by a single optical transition [133]. Therefore, we model the GSB at these wavelengths using Equation 4.5. We are able to fit the data without introducing any more free parameters.

4.3 Triplet Pair Separation and Binding Energy

The kinetic model returns a dissociation timescale of 330 ps for Domain 1, which is surprisingly long—freely diffusing triplets in TIPS-Pn would remain adjacent to one another for only a few picoseconds. We therefore conclude that we are observing the correlated triplet pair, and that interactions between the triplets produce a bound state. This is not unprecedented. In films of TDPP derivatives the intermediate state forms rapidly, yet triplets take up to 1.6 ns to dissociate [122].

In order to quantify the triplet-triplet binding energy we perform a simple simulation of triplet diffusion. Each TIPS-Pn molecule is treated as an individual site, and the simulation is initialized with two adjacent triplets. Because the interactions in TIPS-Pn are much stronger along the a -axis than any other direction [132], we consider a one-dimensional chain of molecules. The triplets each hop with rate constant k , as shown in Figure 4.16.

Because the diffusion is mediated by Dexter transfer, the triplets only hop one site at a time. We do not allow for annihilation, and the triplets cannot hop through each other. Diffusion is primarily controlled by the triplet-triplet interaction, which we model with

$$E(\Delta i) = \begin{cases} -E_1 & \Delta i = 1 \\ -E_2 & \Delta i = 2, \\ 0 & \Delta i > 2 \end{cases}, \quad (4.15)$$

where Δi is the site separation between the two triplet excitons (in Figure 4.16 $\Delta i = 2$), and $E_1 \geq E_2 \geq 0$.




Figure 4.16: Diagram of the hopping model used to estimate triplet binding energy.

This form was chosen based on examining the wave function of the lowest triplet exciton, which was computed by solving the Bethe-Salpeter equation (BSE) (calculations done by Neaton group, see Section 4.4) and is visualized in Figure 4.17. The red star denotes the fixed hole position, and the corresponding electron isosurface is plotted in yellow. The triplet is seen to be somewhat delocalized—triplets can be separated by up to two lattice vectors along the π -stacking direction and still have significant wave function overlap. Hence we allow for triplet-triplet binding when $\Delta i \leq 2$.

The simulation proceeds via the Gillespie algorithm: for any hop that involves increasing the energy by an amount E , the rate is decreased by a Boltzmann factor to $ke^{-E/k_B T}$. The diffusivity of triplets along the fast axis (100) of TIPS-Pn has been measured as $D = 0.006 \text{ cm}^2/\text{s}$ [80]. The separation between TIPS-Pn molecules along the 100 direction is $\Delta x = 7.7 \text{ \AA}$ [142]. The hopping rate is therefore $k = \frac{D}{(\Delta x)^2} = 1.01 \text{ ps}^{-1}$.

The simulation is run for many energetic landscapes (choices of E_1 and E_2). For a given energetic landscape, we run many trajectories to get $\Delta i(t)$. In order to relate our measured dynamics to a timescale, we recall from Figure 4.13a that triplet-triplet interactions perturb the absorbance, an issue that is explored in more detail in Section 4.4. To lowest order, the change in the absorbance for a given separation is proportional to the binding energy. Therefore, for two attractive sites, we calculate the following weighted probability that the triplets are interacting:

$$\text{Signal}(t) = p(\text{sep} = 1; t) + p(\text{sep} = 2; t) \times \frac{E_2}{E_1}. \quad (4.16)$$




Figure 4.17: A visualization of the lowest energy triplet exciton wave function in TIPS-Pn, calculated by solving the BSE.

This function, $\text{Signal}(t)$, decays from a value of 1 at $t = 0$ to a value of 0 at long times. Because there are infinitely many dissociated sites, all triplet pairs are dissociated at steady state, despite the presence of a non-negligible binding energy. The decay is not exactly exponential, but is similar enough that a characteristic dissociation timescale, τ_{dis} , can be reliably determined by searching for the delay at which $\text{Signal}(t)$ falls to e^{-1} .

A contour plot of the dissociation timescale for a variety of energetic landscapes is shown in Figure 4.18. For E_2 values greater than the thermal energy the contours bend, signifying that the presence of a binding energy between next nearest neighboring sites is able to significantly slow down the dissociation kinetics for a given binding energy E_1 between nearest neighbors. Examining Table 4.2, we see that there is some uncertainty and variability in the measured dissociation timescales, but they fall roughly within the range from 200 ps to 400 ps. We highlight these contours in red and yellow colors, and we use them to determine the range of plausible binding energies. As indicated by the purple vertical lines, these contours are contained within the region $50 \text{ meV} \leq E_1 \leq 80 \text{ meV}$. Hence, the triplet-triplet binding energy should be within this range as well.

There is little literature with which to compare our measured 50 - 80 meV value for the binding energy, but a report on several acene films, including amorphous TIPS-Pn, also found evidence of an attractive triplet-triplet interaction. They report energies of about 30 meV [88]. The non-negligible binding energy is consistent with prior reports that suggest that the triplets are coupled via an exchange interaction [126], and suggests that we must consider the possibility that the singlet-mediated triplet exciton diffusion mechanism advanced by Huang et al. for tetracene [81] and rubrene [80] is also significant in TIPS-Pn. They hypothesize that triplets are able to reform singlets, which effectively boosts the apparent triplet diffusion length because singlet transport is much more rapid than triplet transport. They also investigate TIPS-Pn with TAM [80], but presumably trade off the quantity of distinct time delays measured for exquisite spatial resolution, which could explain why they do not




Figure 4.18: Contour plot of τ_{dis} for various binding energy landscapes. Vertical bars indicate the range of binding energies that can produce dissociation timescales within our measured range of 200-400 ps.

resolve the singlet-triplet equilibration and long bound triplet lifetime that we observe.

4.4 Effects of Charge Transfer Character on the Electronic Structure of Bound Triplets

4.4.1 Experimentally Observed Optical Signatures of Bound Triplets

Examining the excited state absorbance spectra in Figure 4.13a, we see that the singlet, correlated triplet, and dissociated triplet all share a peak around 525 nm, in agreement with previous observations for TIPS-Pn both in solid state [138, 144] and solution [136] (blue-shifted). Those TA spectra, however, included contributions from multiple ESA and also GSB TDMs, which are clearly seen to overlap in Figure 4.13a. In contrast, here the *polarized* white light TAM allows us to cleanly measure the absorbance spectra of individual species in quantitative terms, and compare them absolutely. We can therefore say that the bound and dissociated triplets have nearly identical absorbance spectra, but the interacting triplets show slightly diminished absorbance for $\lambda < 550$ nm, as highlighted in the inset to Figure 4.13a.

Although this reduction in absorbance is small and close to the error of our measurement, it is intriguing in light of the results of Pensack et al. [130], who studied five pentacene derivatives and found that the TA spectra of interacting triplets tend to be similar to dissociated triplets, but with a few systematic discrepancies. They find evidence that the interacting triplet tends to have slightly less positive ΔOD than dissociated triplets around 530 nm, which could indicate weaker ESA in agreement with our results. They also find that the interacting triplet has slightly less ESA around 1000 nm, but displays the opposite behavior

around 1300 nm. In that region, the interacting triplet has stronger ESA than the dissociated triplet. Is this behavior general? It is difficult to draw unequivocal conclusions from the limited data, but in the following sections I present evidence that the effect of triplet-triplet interactions should be to generally decrease the absorbance.

4.4.2 A Simple Model of Triplet-Triplet Interactions

Following previous work on Pentacene [104], we define the two-particle excitonic wave function for state X_N as follows:

$$|X_N\rangle = \sum_{kcv} A_{kcv}^{X_N} |kc\rangle |kv\rangle \quad (4.17)$$

Here, $|kc\rangle$ and $|kv\rangle$ are one-particle Bloch wave functions associated with states in the conduction and valence bands, c and v , respectively. The exciton is described as a sum over the tensor product of these two states weighted by the exciton coefficient $A_{kcv}^{X_N}$. The energy of the state is E_{X_N} .

In the subsequent analysis, we will be interested in computing the absorption spectra for various exciton states. The simplest transitions to examine are the $S_1 \rightarrow S_N$ transitions. First we write the imaginary part of the dielectric function, which is related to the absorption:

$$\epsilon_2^{S_1}(\omega) \propto \frac{1}{\omega^2} \sum_N |\hat{e} \cdot \langle S_N | \vec{r} | S_1 \rangle|^2 \delta(\omega - (E_{S_N} - E_{S_1})), \quad (4.18)$$

where \vec{r} is the position operator and \hat{e} is the polarization vector of the light field. Here, the matrix element $\langle S_N | \vec{r} | S_1 \rangle$ is expressed in the electron-hole basis as:

$$\langle S_N | \vec{r} | S_1 \rangle = \sum_k \left(\sum_{c,v,v'} A_{kcv}^{S_1} A_{kcv}^{S_N*} \langle kv' | \vec{r} | kv \rangle - \sum_{c,c',v} A_{kc'v}^{S_1} A_{kcv}^{S_N*} \langle kc | \vec{r} | kc' \rangle \right). \quad (4.19)$$

With Equation 4.18 one can compute the excited state absorption of a single exciton, but we are interested in biexcitons. Define a biexciton state, $|B_N\rangle$, using the exciton states from Equation 4.17:

$$|B_N\rangle = \sum_{nm} C_{nm}^{B_N} |X_n\rangle |X_m\rangle. \quad (4.20)$$

This expression is analogous to Equation 4.17, with the Bloch states having been replaced by exciton states at zero center of mass momentum. To formally construct a complete basis for the biexciton state, we should also include finite momentum excitons in the above sum. Here we note that in a periodic crystal the computed exciton band structure for TIPS-Pn exhibits little dispersion, and therefore we neglect any variations of the exciton character over the Brillouin zone.

Using this framework, we consider both correlated and dissociated triplet pairs. Both are spin-correlated and are overall singlets. In the case of the dissociated triplet pair, the

interaction between triplet excitons is negligible, and, as a result, Equation 4.20 simplifies and we may express the triplet pair, consisting of triplet excitons T_n and T_m , as the tensor product $|T_n\rangle|T_m\rangle$. The energy of this biexciton state is simply $E_{T_n} + E_{T_m}$, while the absorption spectrum associated with the transition $|T_1T_1\rangle \rightarrow |T_nT_m\rangle$ can be computed using Equation 4.18 with the matrix element modified to be $\langle T_nT_m|\vec{r}|T_1T_1\rangle$. If we assume that the incident photon interacts primarily with an individual triplet exciton, it can be shown that this matrix element reduces to:

$$\langle T_nT_m|\vec{r}|T_1T_1\rangle \rightarrow \langle T_n|\vec{r}|T_1\rangle\delta_{1m} + \langle T_m|\vec{r}|T_1\rangle\delta_{1n}. \quad (4.21)$$

We expect this to be a good approximation for the matrix element related to triplet absorption, in the dissociated limit.

To approximate the correlated case, we expect that these uncorrelated bi-triplet states can mix with other states that are nearby in energy, which also lends them both charge-transfer and singlet character [88, 146, 129]. Indeed it has been shown that in the solid state, two-point interactions as simple as the bare Coulomb interaction are sufficient to mix bi-triplet and singlet states [104].

Here, we hypothesize that the singlet character of the correlated triplet pair may be responsible for the slightly diminished absorption observed in experiment. To test this hypothesis, we approximately express the correlated triplet pair as

$$|(T_1T_1)^1\rangle = N_1(|T_1\rangle|T_1\rangle + \lambda_1|S_1\rangle), \quad (4.22)$$

where λ_1 is taken to be a small mixing coefficient and $N_1 = (1 + \lambda_1^2)^{-1/2} \approx 1 - \lambda_1^2/2$ is the appropriate normalization factor. It should be noted that $|S_1\rangle$ is similar in energy to $|T_1\rangle|T_1\rangle$, and has the same overall spin, and thus it is certainly plausible that a perturbation could mix these two. Other triplet pairs, such as $|T_1\rangle|T_2\rangle$, are at least 800 meV higher in energy than $|T_1\rangle|T_1\rangle$ (see Figure 4.19), and hence are not expected to mix significantly.

In a similar manner, we model the excited correlated triplet pair as

$$|(T_1T_n)^1\rangle = N_n(|T_1\rangle|T_n\rangle + \lambda_e|S_e\rangle). \quad (4.23)$$

Here, $|S_e\rangle$ is an excited singlet state with energy comparable to $|T_n\rangle|T_1\rangle$. We note it is an oversimplification to express the correlated pair in terms of a single bi-triplet and singlet. Equations 4.22 and 4.23 can be generalized to include multiple singlet and bi-triplet states. For the moment we have restricted ourselves to this simplified form for illustrative purposes.

The squared dipole matrix element associated with absorption of the correlated triplet pair can be expressed as:

$$|\langle (T_nT_1)^1|\vec{r}|(T_1T_1)^1\rangle_x|^2 \approx |\langle T_n|\vec{r}|T_1\rangle_x|^2 * \left(1 - \lambda_1^2 - \lambda_e^2 + 2\lambda_1\lambda_e \frac{\langle T_n|\vec{r}|T_1\rangle_x \langle S_e|\vec{r}|S_1\rangle_x}{|\langle T_n|\vec{r}|T_1\rangle_x|^2}\right), \quad (4.24)$$

where we have chosen to focus on the Cartesian x -component (i.e., x -polarized incoming light) for concreteness, and have kept terms up to second order in λ , assuming that all

mixing coefficients are small compared to 1. Note that due to the inversion symmetry of the system, all TDMs and mixing coefficients are real numbers.

What can we learn from Equation 4.24? We see that the absorption of the correlated triplet pair is similar to that of the dissociated triplet, but there are several second order correction terms. The first two are negative definite, indicating why the correlated triplet might *generally* absorb less strongly than the dissociated triplet. Intuitively, as singlets mix into the wave function for the correlated triplet pair, the system is less likely to be found in the $|T_1T_1\rangle$ state, and absorption to populate $|T_1T_n\rangle$ becomes less likely, because matrix elements of the form $\langle T_1T_n|\vec{r}|S_1\rangle$ are equal to 0 from parity arguments. A more strongly π -stacked system will likely display larger mixing coefficients, and hence a larger discrepancy between bound and dissociated triplet absorption.

The above arguments can explain why triplet-triplet interactions generally cause the absorbance to decrease, but the opposite has also been known to occur. Pensack et al. [130] found a near-IR region in which the bound triplet absorbs much more strongly than the dissociated triplet. To understand why, we consider the last term of Equation 4.24, which can be either positive or negative. The first part depends on mixing coefficients, and the second part depends on TDMs. Define the second part as $C_x = \langle T_n|\vec{r}|T_1\rangle_x \langle S_m|\vec{r}|S_1\rangle_x / |\langle T_n|\vec{r}|T_1\rangle_x|^2$. We can calculate the TDMs for absorption from the S_1 and T_1 states, and hence C_x , which can be thought of as a measure of how strong and correlated the absorption from the singlet and triplet excitons are. If C_x is small, then we expect the negative definite terms to dominate. If C_x is large, then the last term of Equation 4.24 has the potential to override the others, and we could observe stronger absorption from the correlated triplet.

We now turn to the general case. The correlated triplet pair that would dissociate to $|T_i\rangle|T_j\rangle$ is expressed as

$$|(T_iT_j)^1\rangle = N_{ij} \left(|T_i\rangle|T_j\rangle + \sum_l \lambda_{ij,l} |S_l\rangle + \sum_{lm} \eta_{ij,lm} |T_l\rangle|T_m\rangle \right), \quad (4.25)$$

where the normalization factor is

$$N_{ij} = \left(1 + \sum_i \lambda_{ij,l}^2 + \sum_{lm} \eta_{ij,lm}^2 \right)^{-1/2}. \quad (4.26)$$

We introduce a broadening σ by replacing $\delta(\omega - E) \rightarrow \frac{1}{\sigma\sqrt{2\pi}} e^{-(\omega-E)^2/2\sigma^2}$. The correlated triplets can mix with many states, but we assume that the mixing is small, and limited to states within a narrow energy range relative to σ . Summing over all possible final states, we find for x -polarized light:

$$\epsilon_{2,x}^{(T_1T_1)^1}(\omega) \approx \epsilon_{2,x}^{T_1T_1}(\omega) * \left[1 - \left(\sum_i \lambda_{11,i}^2 \right) - \overline{\Lambda_{11,i}^2} - \overline{H_{\omega,x}^2} + 2 \frac{\overline{\langle S_\omega|\vec{r}|S_1\rangle_x} * \overline{\langle T_\omega|\vec{r}|T_1\rangle_x}}{|\overline{\langle T_\omega|\vec{r}|T_1\rangle_x}|^2} \left(\sum_i \lambda_{11,i} \right) * \overline{\Lambda_{\omega,x}} \right], \quad (4.27)$$

where the full definition of each term can be found in Equation 4.28. Similar equations can be written for y - and z -polarized light as well. As before, we find that the second order correction




Figure 4.19: Energy levels for low-lying singlet and triplet excitons, computed with BerkeleyGW.

includes several negative definite terms, and a term of indeterminate sign. In Section 4.4.3 we compute the value of this term, and draw conclusions about the circumstances in which the correlated triplet pair is likely to absorb less than dissociated triplets.

$$\begin{aligned}
 \overline{\langle S_\omega | \vec{r} | S_1 \rangle_x} &= \sum_n \frac{e^{-(E_{S_n} - \omega)^2 / 2\sigma^2}}{\sigma \sqrt{2\pi}} \langle S_n | \vec{r} | S_1 \rangle_x \\
 \overline{\langle T_\omega | \vec{r} | T_1 \rangle_x} &= \sum_n \frac{e^{-(E_{T_n} - \omega)^2 / 2\sigma^2}}{\sigma \sqrt{2\pi}} \langle T_n | \vec{r} | T_1 \rangle_x \\
 \overline{|\langle T_\omega | \vec{r} | T_1 \rangle_x|^2} &= \sum_n \frac{e^{-(E_{T_n} - \omega)^2 / 2\sigma^2}}{\sigma \sqrt{2\pi}} |\langle T_n | \vec{r} | T_1 \rangle_x|^2 \\
 \overline{\Lambda_{\omega,x}^2} &= \frac{\sum_n e^{-(E_{T_n} - \omega)^2 / 2\sigma^2} |\langle T_n | \vec{r} | T_1 \rangle_x|^2 \sum_i \lambda_{1n,i}^2}{\sum_n e^{-(E_{T_n} - \omega)^2 / 2\sigma^2} |\langle T_n | \vec{r} | T_1 \rangle_x|^2} \\
 \overline{\Lambda_{\omega,x}} &= \frac{\sum_n e^{-(E_{T_n} - \omega)^2 / 2\sigma^2} \langle T_n | \vec{r} | T_1 \rangle_x \sum_i \lambda_{1n,i}}{\sum_n e^{-(E_{T_n} - \omega)^2 / 2\sigma^2} \langle T_n | \vec{r} | T_1 \rangle_x} \\
 \overline{H_{\omega,x}^2} &= \frac{\sum_n e^{-(E_{T_n} - \omega)^2 / 2\sigma^2} |\langle T_n | \vec{r} | T_1 \rangle_x|^2 \sum_{jl} \eta_{1n,jl}^2}{\sum_n e^{-(E_{T_n} - \omega)^2 / 2\sigma^2} |\langle T_n | \vec{r} | T_1 \rangle_x|^2}
 \end{aligned} \tag{4.28}$$

4.4.3 Calculating Properties of the Correlated Triplet Pair

These calculations were performed by Jonah Haber using the BerkeleyGW package [147], following previous work on TIPS-Pn [133]. Basis states $|kc\rangle$ and $|kv\rangle$ are computed via density functional theory [147], while the exciton coefficients and energies in Equation 4.17 are computed from many-body perturbation theory by solving the BSE in the quasiparticle basis, using the GW approximation [148]. The computed exciton energies are shown in Figure 4.19 and agree well with our observations and the literature [108]. Knowing the exciton wave functions enables us to compute the TDMs such as $\langle S_n | \vec{r} | S_1 \rangle$. The calculated linear absorption spectrum is shown in Figure 4.20 with 40 meV broadening, and shows good agreement with experiment.




Figure 4.20: Comparison of the polarization-averaged linear absorption spectrum calculated with BerkeleyGW and measured on a disordered polycrystalline film of TIPS-Pn.




Figure 4.21: Absorption spectra of excited singlet and triplet states S_1 and T_1 computed with BerkeleyGW.

We also calculate the absorption from lowest excited singlet and triplet S_1 and T_1 . These are shown in Figure 4.21 with 50 meV broadening. These calculations are highly novel and not yet at the point where a quantitative comparison with experiment can be done for crystalline solids, due to the high energies of the states involved. Nevertheless, we do see that the essential physical features are captured, particularly the large peak around 550 nm.

We now turn to the central question of what effect triplet-triplet interactions has on the absorption. The crucial term is the last term in Equation 4.24, which can be either positive or negative and appears in a more complicated form in Equation 4.27. The first part of this term is determined by the excited state absorption spectra, which is calculated and displayed




Figure 4.22: Coefficient of the uncertain-sign term in Equation 4.27, for unpolarized light incident along the z -axis, and broadening $\sigma = 100$ meV.

in Figure 4.21. The term also depends on the mixing coefficients, but estimating those values is beyond the scope of this work, and we can learn something without them.

To facilitate comparison with Figure 4.13a, we consider the case of unpolarized light incident normal to the substrate, along the z -axis. That is, we are interested in the following quantity:

$$\frac{\overline{\langle S_\omega | \vec{r} | S_1 \rangle_x} * \overline{\langle T_\omega | \vec{r} | T_1 \rangle_x} + \overline{\langle S_\omega | \vec{r} | S_1 \rangle_y} * \overline{\langle T_\omega | \vec{r} | T_1 \rangle_y}}{|\overline{\langle T_\omega | \vec{r} | T_1 \rangle_x}|^2 + |\overline{\langle T_\omega | \vec{r} | T_1 \rangle_y}|^2}. \quad (4.29)$$

As discussed earlier for the simpler case, if this term is small then the negative definite terms in Equation 4.27 are likely to dominate. If this term is much bigger than 1 in magnitude, then it is possible the last term of Equation 4.27 could dominate all of the others, and if it is positive result in a correlated triplet that absorbs *more* strongly than the dissociated triplets. Expression 4.29 is plotted in Figure 4.22, and we see that throughout the visible region it is small and oscillatory. This explains why when correlated triplet absorbance differs from that of dissociated triplets in the visible region, it tends to be weaker. We do, however, observe a large, positive feature around 1400 nm, which opens the door to enhanced absorption of the correlated triplet pair due to triplet-triplet interactions, possibly explaining the strong correlated triplet pair absorption observed by Pensack et al. [130] in this spectral region.

4.4.4 Future Directions

Our calculation above of the correlated triplet pair can only provide limited, qualitative information, in part because current ab initio many body perturbation theory methods cannot treat the four particle Greens function required to describe a bi-exciton from first principles in the system explored. Work is currently underway in the Neaton group to develop a theory using the four particle Green's function, which will enable calculation of the binding energy and absorption spectrum of the correlated triplet pair.

4.5 Implications for Devices

Given the above results, we consider the effect of π -stacking on singlet fission and what it means for devices. Though it is generally assumed that π -stacking will lead to strong interactions, rapid fission, and rapid triplet dissociation [139, 149, 150, 151], our findings suggest that this is not uniformly the case. The slip-stacked structure of TIPS-Pn makes it a model π -stacking system, but fission is slower than in pentacene [60] (which has a herringbone structure), internal conversion is rapid enough to cause significant losses, and correlated triplets dissociate slowly. These findings may be more generally true of strongly π -stacked systems: similar time scales were found for TDPP derivatives [122], and a theoretical study has indicated that strong coupling between neighboring molecules can slow fission [121]. Fluorinated pentacene undergoes efficient fission in the herringbone polymorph, but when deposited in the π -stacked brickwork polymorph, increased photoluminescence is observed, implying that some triplet pairs fuse to form singlets [152].

We hypothesize that π -stacking might hinder singlet fission because of concomitant increased orbital overlap relative to other packing motifs. Though triplet wave functions in pentacene are primarily confined to a single molecule [153], those in TIPS-Pn are comparatively delocalized, extending about 1.2 nm in the π -stacking direction [154]. Hence non-negligible overlap is expected between triplets on neighboring TIPS-Pn molecules. More overlap could increase the coupling between excitons and cause more significant charge-transfer state mixing, which in turn could increase the binding of the correlated triplet pair. Strong coupling could also hinder the fission yield by promoting internal conversion: in acene dimers, increased overlap has been shown to speed up $(TT)^1$ relaxation [155]. The crystal structures commonly found in π -stacked systems could also play a role. It has recently been shown that a direct Coulomb coupling between the singlet and two distinct correlated triplet excitons, arising from the herringbone crystal structure, can drive fast fission in crystalline pentacene [104]. Following Ref. [104], however, the low-lying triplet pair can only couple via this direct mechanism to dark singlet states due to inversion symmetry in TIPS-Pn, indicating that fast triplet formation via a purely Coulomb process would be significantly hindered. For all of these reasons, we propose that our observations on TIPS-pentacene may be generalizable across multiple π -stacked systems.

The greater understanding of the kinetics and energetics of the correlated triplet pair that

is emerging could lead to more efficient singlet-fission-based photovoltaics. We stress that the relevant quantity to optimize for photovoltaics is the triplet formation internal quantum efficiency, η : the number of triplets harvested per exciton created, derived in Equation 4.6. This efficiency depends not only on the fission dynamics, but also on the extent and ease of diffusion and on the interplay between triplet dissociation and other decay pathways. We observe triplet binding that is comparable to $k_B T$ at room temperature, which makes singlet-mediated diffusion viable, and hence could allow more triplet pairs to be extracted before they decay, thus boosting η .

From Table 4.2, η in the domains that we have studied is approximately 100 %. A real device, however, has a triplet extraction layer, which effectively shortens τ_{dis} and therefore boosts η . As a concrete example, lowering τ_{dis} to 50 ps in Domain 1 would yield $\eta = 179\%$. This is in line with the results of Yang et al. who achieved internal quantum efficiencies of $160 \pm 40\%$ using a PbSe acceptor [135]. With a TIPS-Pn layer only 16 ± 4 nm thick, even if the diffusivity were $0.1 \text{ cm}^2/\text{s}$ (ten times lower that of a single crystal [80]), one would still expect most excitons to find an interface within 50 ps. Hence, π -stacking remains a promising design principle for singlet fission, so long as the strength of the resulting triplet-triplet interaction is tailored to maximize triplet extraction while minimizing non-radiative decay.

Chapter 5

TAM (and friends) Study Self-Assembly of Rubrene Spherulites

There are many techniques for creating solid samples of organic semiconductors, and in general there is a trade-off between ease of processing and device performance. At one extreme, physical vapor transport heats a material to sublimation and slowly deposits the vapor, which can form a large, single crystal that displays ideal transport properties [156]. At the other extreme are solution processing techniques, which can be as simple as making a solution of the material and dropcasting it on to a substrate to dry, or spincoating to make a thin, uniform, but generally non-crystalline film. Using industrial-scale techniques such as roll-to-roll processing can make the final result more repeatable [157], and careful control of the deposition conditions can induce crystalline structure [158], but in general solution-processed films are disordered. One manifestation of this disorder is the existence of grain boundaries, which are often the primary impediment to charge transport [159].

In order to improve the performance of a solution-processed film, one can use post-deposition processing to increase uniformity and crystallinity. Two common techniques are solvent-vapor annealing, in which the film sits in a saturated environment of solvent vapor, and thermal annealing, in which the film is heated above room temperature [11]. In both cases, the molecules can re-orient themselves into a more ordered and hence lower energy state.

In this chapter I describe work on amorphous films of the organic semiconductor rubrene, which are thermally annealed to produce polycrystalline spherulites. Spherulites are circular objects that are locally crystalline, but have a smoothly-varying crystalline orientation as a function of azimuthal angle. They have kinetically trapped structure over many orders of magnitude in size, and there are open questions about the forces driving their formation. For example, rubrene spherulites become more polycrystalline as the annealing temperature is increased, behavior that runs counter to other spherulite-forming materials. We have studied spherulite formation with a variety of temporally- and spatially-resolved techniques, including real-time X-ray scattering, X-ray microdiffraction, and atomic force microscopy. Those studies show that spherulite nuclei template regions of strong misorientation (the

gradient of the crystalline orientation varies substantially over a small change in azimuthal angle) which, depending on the temperature, can either propagate lines of strain or relax into separated crystalline domains. These ideas are included in a Monte Carlo simulation of rubrene crystallization in order to explain the observed temperature dependence.

We and our collaborators are currently attempting several complementary measurements techniques to shed additional light on rubrene self-assembly, and those are discussed here as well. With scanning transmission electron microscopy our collaborators in the Minor group at UC Berkeley have measured the crystal structure with nanometer-scale spatial resolution, which will enable us to localize strain within individual crystallites. Our collaborators in the Raschke group at CU Boulder have made attempts at infrared scanning-scattering near-field optical microscopy, which can measure the local phonon structure and orientation of the crystal. Lastly, we have done preliminary TAM measurements in order to determine whether local spherulite structure has an influence on ultrafast electronic dynamics.

This chapter describes the early stages of an ambitious project to study self-assembly. Using rubrene spherulites as a model organic system, we are correlating multiple imaging techniques to gain a complete picture of the relevant forces, how those forces produce both transient and final structures, and how those structures influence the local excited state behavior.

5.1 Fabrication and Basic Characterization of Rubrene Spherulites

5.1.1 Properties of Rubrene

Rubrene, or 5,6,11,12-Tetraphenylnaphthacene, is the small organic molecule shown in Figure 5.1a. It consists of a tetracene core and four phenyl groups attached to the central rings. Rubrene has been extensively studied over the years for its potential optoelectronic applications. Strong luminescence makes it promising for OLEDs [160], the ability to undergo singlet fission makes it promising for OPVs [161, 162] and carrier mobilities of above 40 $\text{cm}^2/\text{V} \cdot \text{s}$ [163, 164] are a record for organic transistors [165].

Rubrene's performance is due to its crystalline structure. It forms three stable polymorphs: orthorhombic, triclinic, and monoclinic [166, 167, 168]. The orthorhombic one is the lowest in energy [166], and is also the polymorph that exhibits the highest mobility [167]. Its crystal structure is shown in Figure 5.1b,c from both the side-on and top-down views. When formed on a substrate, thin films of rubrene grow with their a_3 -axis oriented orthogonal to the substrate [169]. The a_1 - and a_2 -axes are in the plane of the substrate, with the a_2 -axis corresponding to the π -stacking direction, as can be seen clearly in Figure 5.1c. The length of the unit cell along the a_2 -axis is 7 Å, and the distance between parallel tetracene cores is merely 3.5 Å. Intermolecular coupling is so strong that single crystals of rubrene have been shown to display bandlike transport and delocalized excitons similar to those found in inorganic semiconductors [170].




Figure 5.1: Structure of rubrene: (a) single molecule [9] and (b) orthorhombic crystal looking from the side and (c) top-down.

5.1.2 Structure and Detection of Spherulites

Spherulites are structures that are locally crystalline, but the crystalline orientation gradually varies as a function of space. They are often circular and the crystalline orientation varies azimuthally so that some fast growth direction is always directed radially outwards from a central point. Spherulites are most frequently observed and studied in polymers, but they can also be found in metals, minerals, and small organic molecules [171]. Diagrams are shown in Figure 5.2. Figure 5.2a is a generic spherulite, in which the crystalline orientation of some axis, indicated by the green arrows, always points radially outwards. Figure 5.2b more specifically represents rubrene, with each blue oval standing in for one molecule. Zooming in on a small region, the structure is crystalline, but it continuously varies so that the a_2 -axis is always oriented radially outwards [172].

Variation of the crystalline axes implies that there are innumerable interfaces within a spherulite between micro-crystalline domains. If two adjacent crystallites have nearly the same crystalline orientation, then there is a *low-angle grain boundary* between them. Many low-angle grain boundaries accumulate to create the spherulite structure. On the other hand there can also be *high-angle grain boundaries* in which neighboring crystallites have significantly different crystalline orientations. Either of these boundaries can occur within a single spherulite or at the interface where two spherulites have grown together.

Both intra- and inter-spherulite grain boundaries have a strong affect on charge transport. Most studies of charge transport in small molecule organic semiconductor spherulites have been done on TES ADT, which is the non-fluorinated version of the molecule shown in Figure 3.1a. Transistors fabricated entirely within a single spherulite were found to have an orientation-independent mobility of about $0.4 \text{ cm}^2/\text{V} \cdot \text{s}$ [92]. This lack of orientational




Figure 5.2: Structure of (a) generic spherulite and (b) rubrene spherulite.

anisotropy implies that hopping across low-angle grain boundaries is the primary bottleneck, not intragrain transport. Indeed, further studies found that low-angle ($0 \pm 20^\circ$) interspherulite boundaries do not affect current flow, but high-angle ($90 \pm 20^\circ$) boundaries do [173]. Furthermore high-angle grain boundaries have significantly more carrier traps than low-angle boundaries [174]. It is hypothesized that molecules kinetically trapped at the inter-spherulite interface can help bridge the gap, promoting charge transport [173].




Figure 5.3: Cartoon schematic of crossed polarizers.

Spherulites are typically viewed optically through crossed polarizers, diagrammed in Figure 5.3. Incident light passes through polarizer #1, the sample, polarizer #2, and then enters the detector. Polarizers #1 and #2 are oriented orthogonally to each other so that no light reaches the detector unless the optical polarization is rotated by the sample. This can happen if the sample is birefringent, which is common for crystalline materials. Incident light that is polarized somewhere in between the fast and slow optical axes of the crystal

will be polarization-rotated, leading to a brighter region in the image. If the incident light is polarized parallel to either the fast or slow axis there will be no rotation and hence a dark spot. Therefore a spherulite with a full 360° crystal rotation should display four bright and four dark wedges.




Figure 5.4: Rubrene spherulites viewed through crossed polarizers in (a) our microscope, (b) our collaborator's microscope in reflection mode, (c) and in transmission mode (courtesy of Jordan Dull, Princeton University).

Images of spherulites viewed through crossed polarizers are shown in Figure 5.4. The procedure for making this sample is discussed in Section 5.1.3. Notice that when viewed through our microscope (Figure 5.4a) there are two bright and dark wedges, whereas when rubrene spherulites are viewed through our collaborators' microscope (Figure 5.4b) there are four bright and four dark wedges. However when our collaborators adjust their microscope from operating in reflection to transmission mode (as ours does), they observe two bright and two dark wedges.

We have a mystery on our hands—why does transmission mode seem to result in two-fold symmetry and reflection mode in four-fold symmetry? Interestingly, in one work two-fold symmetry was intentionally induced in order to unambiguously determine the crystalline axes [172]. This was done with a Nomarski prism in the polarized microscope, however no such optic is in our setup. Furthermore one batch of low-quality spherulites did appear in our microscope with four-fold symmetry, as shown in Figure 5.6, but it is the exception. Scouring the literature reveals several other examples of spherulites that appear to have two-fold symmetry when viewed through crossed-polarizers, such as Ref. [175] (Figure 10) and Ref. [89] (Figure 2c). In neither case is this fact remarked upon.

My hunch is that scattering plays a big role in how the spherulites appear when imaged. The spherulite shown in Figure 5.4a has very different structure and hence scattering properties than the one in Figure 5.6. Another fact to notice is that the dark regions in Figure 5.4a,c are *darker* than the background, whereas the traditional birefringence explanation predicts that the darkest spherulite regions should be equal to the background level, as seen in Figure 5.4b. This implies that the dark regions in Figure 5.4a are dark at least in part because light is being scattered out of that region. It is plausible that scattered light could be imaged differently in transmission than in reflection geometry. Although it would be

preferable to have a complete explanation, between Figure 5.4 and the AFM measurements (see Section 5.1.5), I am confident that the crystal structure in our spherulites undergoes a full 2π rotation.

5.1.3 Procedures for Creating Spherulites




Figure 5.5: Attempted rubrene spherulites made with solution processing techniques.

We initially focused on solution-processed spherulites. Rubrene was deposited from solution and annealed with solvent vapor [176] or heat [169, 177]. These references were used as inspiration to try many different protocols. Ultimately we were not able to make reliable, high-quality rubrene spherulites. Four examples are shown in Figure 5.5, and the protocols used for those specific samples are described below.

Figure 5.5a: Dissolve both rubrene and polystyrene (PS) (280 kDa molecular weight) separately in chloroform, mix so that the total solution is 0.60 wt%, with a 4:1 rubrene:PS

ratio. Filter with $0.45\ \mu\text{m}$ mesh. Clean a cover slip for 1 minute with oxygen plasma in the reactive ion etch (RIE). Spin coat $40\ \mu\text{L}$ of rubrene/PS solution on the cover slip rotating at 1250 revolutions per minute (rpm) for 30 seconds. Place the cover slip in a solvent vapor annealing chamber with a pool of dichloromethane at the bottom. The cover slip was placed at 50% of the chamber height (about 9 inches above the solvent). Leave it for 10 hours. Rod structures form, some of which are very bright when viewed under crossed polarizers.

Figure 5.5b: Dissolve both rubrene and 280 kDa PS separately in chloroform, mix so that the total solution is 0.60 wt%, with a 4:1 rubrene:PS ratio. Filter with $0.45\ \mu\text{m}$ mesh. Clean a cover slip for 1 minute with oxygen plasma in the RIE. Spin coat $40\ \mu\text{L}$ of rubrene/PS solution on the cover slip rotating at 1250 rpm for 30 seconds. Place on a hot plate at $170\ ^\circ\text{C}$ for 5 minutes. Circular structures form, and the upper-right corner of one is shown here. There appear to be lots of little crystals, some of which appear very bright under crossed polarizers.

Figure 5.5c: Dissolve both rubrene and 280 kDa PS separately in chloroform, mix so that the total solution is 0.53 wt%, with a 5:1 rubrene:PS ratio. Filter with $0.45\ \mu\text{m}$ mesh. Clean a cover slip for 20 minutes in piranha solution, a 7:3 ratio of sulfuric acid to hydrogen peroxide. Spin coat $40\ \mu\text{L}$ of rubrene/PS solution at 1250 rpm for 30 seconds. Anneal for 140 minutes in a solvent vapor annealing chamber, perched 3.5 inches above a pool of dichloromethane. A circular structure forms. Here we see the dark center spot and small structures growing outwards. There are also rod-like structures threaded throughout. This is an ordinary transmission image – when viewed through crossed polarizers there is no contrast.

Figure 5.5d: Clean a cover slip by sonicating for 20 minutes in acetone, then isopropyl alcohol, then millipore water. Dissolve 2.5 kDa PS in toluene to 0.625 wt%, and rubrene in chloroform to 1.0 wt%. Filter both solutions through $0.45\ \mu\text{m}$ mesh. Spin down $40\ \mu\text{L}$ of PS solution at 2500 rpm for 60 seconds, place it on a hot plate at $90\ ^\circ\text{C}$ for 2 minutes to dry. Then spin $40\ \mu\text{L}$ of rubrene solution at 2500 rpm for 30 seconds, place on a hot plate at $190\ ^\circ\text{C}$ for 2 minutes to anneal. There is a weak spherulite structure when viewed through crossed polarizers—two bright and two dark wedges can be seen. Lots of small crystal make up the overall structure, and the contrast is weak.

The above are a sampling of the many procedures that we tried. Not represented above are the facts that we also performed solvent vapor annealing with carbon disulfide, and that we dissolved rubrene in other solvents besides chloroform, including dichlorobenzene. Two important things we learned are that the hot plate reading is not reliable (an external temperature probe is crucial), and that one can achieve more uniform films in a spin coater by depositing solution after the substrate is already spinning.

On the advice of Anna Hailey, who was a graduate student in Professor Yueh-Lin Loo's group at Princeton, we switched to thermally evaporated rubrene for greater control. The substrates were cleaned by sonicating at least 15 minutes in acetone, isopropyl alcohol, millipore water, and then a second container of clean millipore water. A 0.625 wt% solution of PS in toluene was dissolved and filtered through a $0.45\ \mu\text{m}$ mesh. We used both 2.5 kDa and 280 kDa PS and found similar results. Hailey found that low molecular weight




Figure 5.6: Rubrene spherulite made with thermal evaporation.

PS (around 2.5 kDa) works better [177]. 20-60 nm of rubrene was then deposited in the MBRAUN thermal evaporator at the Molecular Foundry, using a rate of about 1 \AA/s and a pressure below 6×10^{-6} mbar. Films were then annealed at temperatures between $160 \text{ }^\circ\text{C}$ and $200 \text{ }^\circ\text{C}$ for times ranging between 30 seconds and 5 minutes.

A rubrene spherulite made on the first attempt is shown in Figure 5.6. It was annealed at $175 \text{ }^\circ\text{C}$ for 90 seconds. It consists of many visible grains, but the four-fold spherulite symmetry is clear. Unfortunately several subsequent runs failed to produce anything reliably as good. We eventually discovered that the problem was rubrene purity. Any amount of oxygen in the rubrene will form rubrene endoperoxide when exposed to light [178, 179], which inhibits crystallization [180]. This most likely doomed our attempts to make solution processed spherulites, as the solvents used contain small amounts of oxygen.

In order to create good spherulites, extremely pure rubrene is needed. Not even the 99.99% purity rubrene sold by Sigma-Aldrich will do. When creating the sample shown in Figure 5.6, we used a boat full of already-used rubrene. That is, somebody else had thermally evaporated some of the rubrene and left some behind. Thermal evaporation burns off impurities, yielding a better spherulite. Our next few attempts used fresh rubrene, and yielded poor results. We eventually found that doing two sacrificial thermal evaporation runs to burn off impurities and then thermally evaporating the remaining rubrene yielded decent films, an example of which is shown in Figure 5.8. However, a better way to purify rubrene is to use thermal gradient sublimation [181]. This is done by Professor Barry Rand's group, at Princeton, with whom we eventually partnered. They also use a variety of organic underlayers instead of polystyrene, in particular TPTPA (Tris(4-(5-phenylthiophen-2-yl)phenyl)amine) [182], but otherwise the procedure is the same as described above. All subsequent films were evaporated by Jordan Dull, in the Rand group. In some cases he annealed them as well and in some cases they were shipped to Berkeley for annealing, always being careful to avoid oxygen exposure.

Even when using purified rubrene, precise temperature control is important to achieving



Figure 5.7: Picture of thermal annealing setup.

reproducible results. We eventually realized that our hot plate (Thermo Scientific) was not stable. We experimented with temperature-controlled heating elements, but they tend to be small, which makes it difficult to heat the entire film evenly. Ultimately we decided to use the hot plate but monitor its temperature. The setup inside our glovebox is shown in Figure 5.7. We used a ThorLabs TC-200 temperature controller with two leads attached to a PT1000 temperature sensor. The sensor is attached to a 1/16" thick sheet of aluminum and the aluminum is attached to the hot plate with Kapton tape. In general, we found that setting the hot plate to a specific temperature, T_{set} caused the measured temperature to rapidly approach T_{set} , stay there for a few minutes, and then over the course of about 15 minutes slowly rise to roughly $T_{set} + 10^{\circ}\text{C}$. Further changing the set temperature causes comparable changes to the measured temperature, though it responds faster to temperature increases than decreases.




Figure 5.8: Rubrene spherulites evaporated and annealed at Berkeley (170°C), using purified rubrene.




Figure 5.9: Rubrene spherulites evaporated at Princeton and annealed with a range of temperatures. Nuclei, spherulites, and platelets are noted. Discontinuities are places where the crystalline orientation abruptly changes.

Crossed polarizer images of samples made by Jordan and annealed at Berkeley are shown in Figure 5.9 for three annealing temperatures: 140 °C, 165 °C, and 170 °C. These images are chosen to illustrate a trend – lower annealing temperatures produce structures comprising a countable number of single-crystal domains, known as platelets, whereas higher temperatures produce more smoothly-varying spherulites [182]. Both structures originate from nuclei, which will be seen in Section 5.1.5 to have similar morphologies. As will be discussed in Section 5.5, this trend of increasing fragmentation at higher annealing temperatures is opposite what is predicted by theory and what is observed for most spherulite materials. Understanding the physical distinction between spherulites and platelets, and why they have the observed temperature dependence, will be a primary focus of the remainder of this chapter. Discontinuities, or places where the crystalline orientation abruptly changes, are noted as they will be significant.

5.1.4 Characterization of Optical Properties

Detailed optical measurements of crystalline rubrene have shown that both absorption and PL are much more likely to occur for light polarized parallel to the a_3 -axis [183]. The absorption spectrum is characterized by a vibronic progression starting at 2.32 eV (535 nm) and continuing in steps of 170 meV: subsequent peaks are at 2.49 eV (498 nm) and 2.66 eV (466 nm). For light polarized along the a_1 - or a_2 -axes the 535 nm peak is essentially nonexistent and the other peaks are about 1/5 the strength they are for light polarized along the a_3 -axis [183]. The PL spectrum also shows a vibronic progression, with the first peak at 2.2 eV (564 nm) and the second at 2.05 eV (605 nm). Emission polarized along the a_3 -axis is about 20 times stronger than emission polarized along the a_1 - or a_2 -axes [183].

Absorption and emission spectra of rubrene in chloroform solution are shown in Figure 5.10a. They match what is expected from the literature [183]. If the solution is exposed simultaneously to oxygen and light, rubrene endoperoxide will form. As shown in Figure

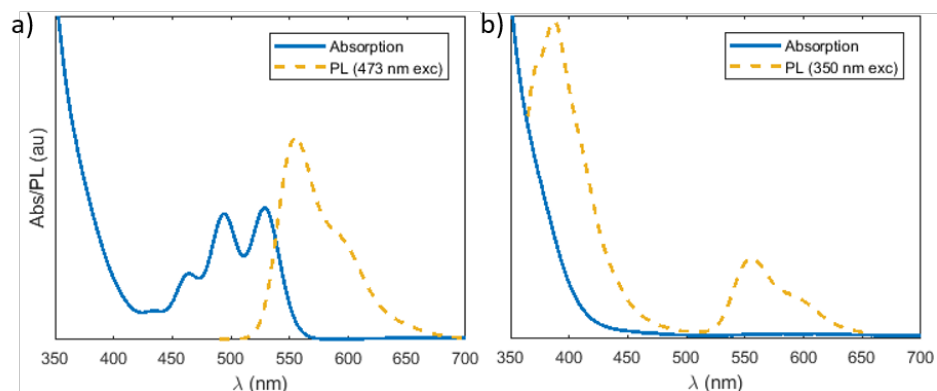


Figure 5.10: Absorption and photoluminescence of (a) rubrene and (b) rubrene endoperoxide.

5.10b the vibronic progression disappears, leaving only a high-energy absorption tail [180]. Exciting the oxidized rubrene at 350 nm, we observe higher energy emission. We can use this fact to check for oxidation of our films.




Figure 5.11: Photoluminescence (a) image and (b) spectra of a proper rubrene spherulite. 405 nm excitation light preferentially couples to oxidized rubrene.

Unfortunately because rubrene has its strongest transitions along the a_3 -axis, it is difficult to study crystalline rubrene in which the a_3 -axis is oriented normal to the substrate. Both absorption and PL are low, as shown in Figure 5.11a. The dark patches are spherulites and the bright areas are amorphous rubrene. PL spectra after 473 nm excitation (Figure 5.11b) show that the amorphous region (orange curve) displays the expected PL spectrum (compare to Figure 5.10b), but the emission of the spherulite (red curve) is skewed towards lower energy. This is an artifact of imaging emission polarized in the $a_1 - a_2$ plane [183]. The emission is also very weak, so the excitation beam scattering is visible around 500

nm. Exciting at 405 nm (purple curve) shows high-energy emission, evidence of oxidation. Oxidation was found to get worse over time in amorphous regions, but to remain consistently small in well-crystallized regions. This is because a thin oxide layer forms on the surface of crystalline rubrene, but the tight crystal packing prevents oxygen from penetrating further into the bulk [184].




Figure 5.12: Photoluminescence (a) image and (b) spectra of a thermally evaporated rubrene spherulite excited at 473 nm.

With less well-oriented spherulites, the a_3 -axis develops an in-plane component and PL becomes a more useful structural probe. Figure 5.12a shows the confocal PL image of a thermally evaporated spherulite film, similar to the one shown in Figure 5.6. Individual grains are visible at higher spatial resolution, but even looking at this bulk structure we can see a slight shift between the edge (oval) and interior (square) of the spherulite. The PL spectra are shown in Figure 5.12b. This may provide insight into the structure of the spherulite edge, but because we transitioned to higher-quality films after these observations, this line of inquiry was abandoned.

Finally, optical techniques can be used to test for oxidation. Many solution processed films, such as the one shown in Figure 5.13a display dispersed small regions that fluoresce more than their surroundings when excited at 405 nm. An example is circled and labeled # 1. As shown in Figure 5.13b (purple), it emits broadly at high energies when excited at 405 nm. The dip at 473 nm is due to a dichroic mirror. Other regions of the film do not display this high-energy emission, implying that the rubrene endoperoxide clumps together for some reason. Exciting at 473 nm yields the blue spectrum in Figure 5.13b. This sample was made by spinning a 0.625 wt% solution of 280 kDa PS in toluene, letting it dry, spinning 1.0 wt% rubrene in chloroform, then annealing on a temperature-controlled hot plate at 175 °C for 60 seconds. The bubbles are places where the rubrene has peeled back during annealing.

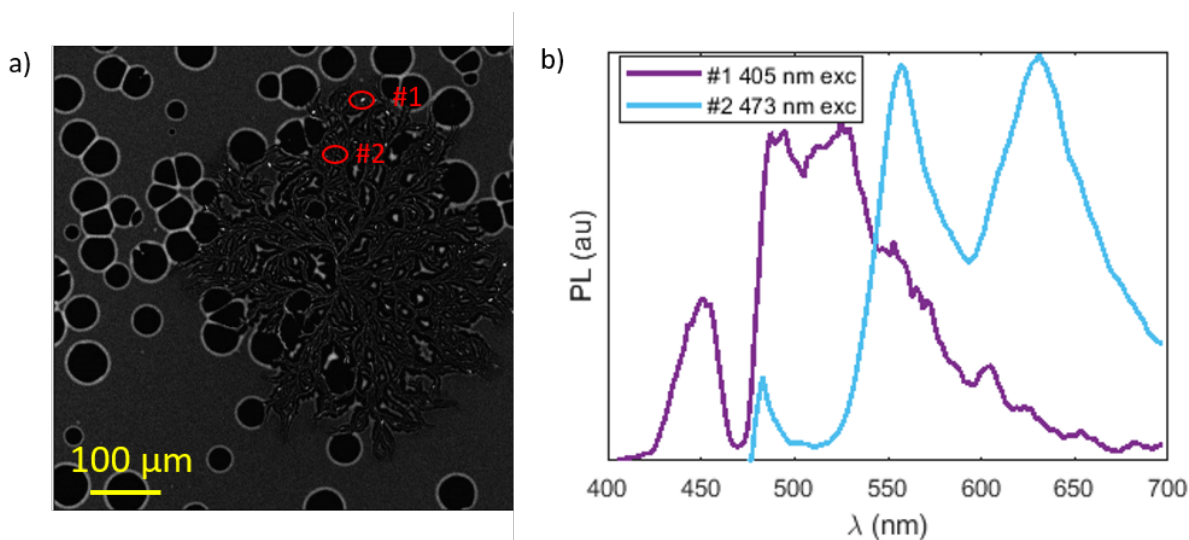


Figure 5.13: Photoluminescence (a) image and (b) spectra of a solvent vapor annealed rubrene spherulite.




Figure 5.14: AFM image of TPTPA on substrates: (a) ITO and (b) bare glass.

5.1.5 Characterization of Surface Roughness and Interfaces

Atomic force microscopy (AFM) can be a good way to check the morphology and growth direction of a spherulite. To start with we consider the morphology of the substrates, which for many samples fabricated by the Rand group are glass with two strips of indium tin oxide (ITO), a conductive transparent material. The ITO is rougher—its RMS roughness is 3.8 nm, as compared with 1.0 nm for glass. Depositing TPTPA does not change this fact, as shown in Figure 5.14. The roughness on ITO is essentially unchanged, while the RMS roughness of the TPTPA on glass is about 2.0 nm. The next step, thermally evaporating

rubrene, results in a very flat film. The RMS roughness of amorphous thermally evaporated rubrene was measured to be less than 0.5 nm, which is comparable to the dimensions of a single rubrene molecule. Despite this flatness, the roughness of the underlayer continues to exert itself. Spherulite formation was found to proceed differently depending on whether the spherulite was above glass or ITO (see Section 5.2.3).

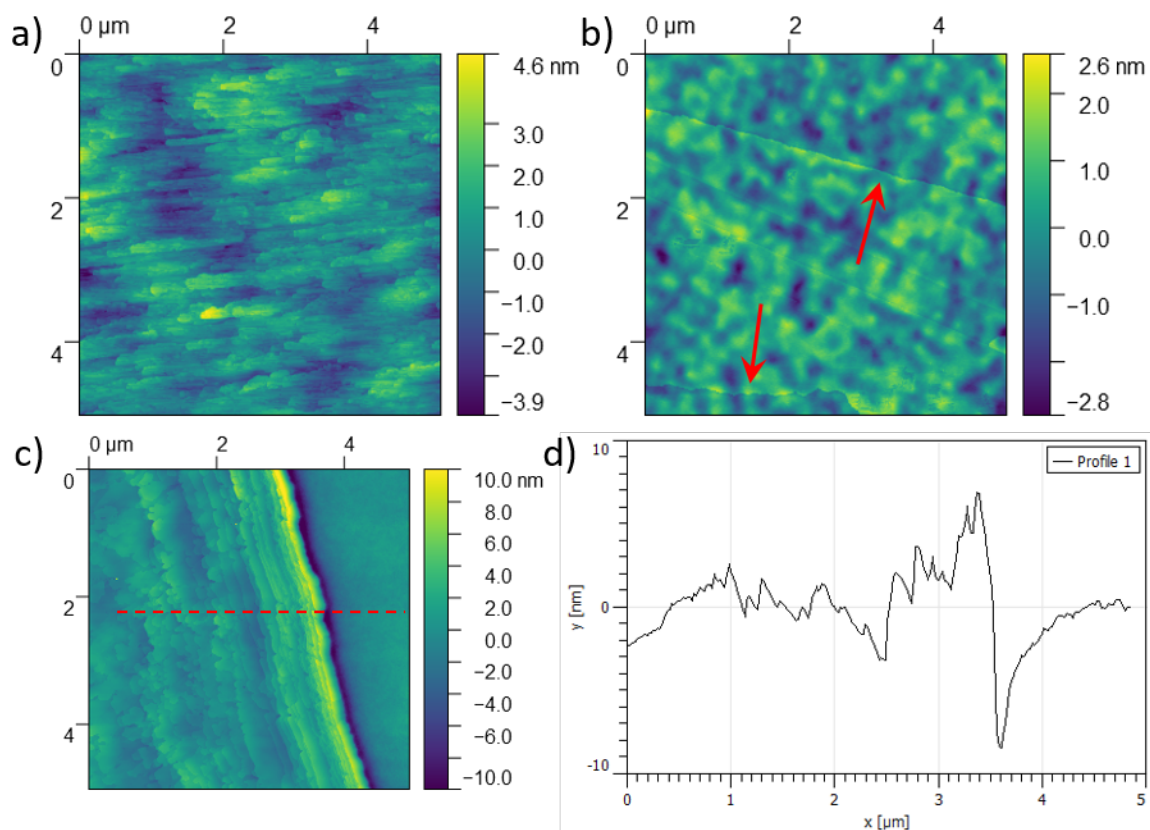


Figure 5.15: AFM image of (a) rubrene spherulite showing growth direction, (b) rubrene platelet showing interfaces (red arrows), (c) spherulite edge showing lip, and (d) linecut across spherulite edge.

Figure 5.15a shows an AFM scan of a small portion of a rubrene spherulite that was annealed at 170 °C. It appears to be made up of rod-like crystallites that extend from right to left. That direction (right to left) happens to point radially outward from the center of the spherulite. In fact looking azimuthally around the spherulite, the apparent crystalline pointing direction varies so as to always be roughly outwards, indicating that the rod-like structure arises because growth is faster along the a_2 -axis and that we can use pointing direction as a proxy for crystalline orientation. It is unknown whether the rods extend into the film or only exist on the surface.

The structure shown in Figure 5.15a is generic. It seems to appear in all films and throughout most of the spherulite. The RMS roughness is generally between 1 and 2 nm. For comparison, the height of a single rubrene molecule along the vertical a_3 -axis is 1.4 nm [169]. Where the structure does differ is at the spherulite nucleus, to be discussed next, and at the edge of the spherulite. An AFM image of the spherulite edge is shown in Figure 5.15c, with a linecut in 5.15d corresponding to the red dashed line in 5.15c. Radially-oriented rods give way to a wave-like pattern, culminating in a tall lip and a deep trough. This feature forms soon after annealing starts—the film shown in Figure 5.15c was annealed for 20 seconds and the spherulites were only about 100-200 μm in diameter. With longer annealing times and larger spherulites, this lip can grow to be 30 nm tall. We believe that the lip is due to densification of rubrene during crystallization, but do not fully understand this process.

Rubrene films annealed at lower temperatures show a different morphology. The film in Figure 5.15b was annealed at 140 $^\circ\text{C}$ to produce platelets. No rod-like structures are visible, perhaps because at lower temperatures the difference in growth rate between the fast- and slow-axes is not so stark [185]. Instead we see a very flat film, with a RMS roughness of less than 1 nm. We also see clear cracks/steps between different crystalline platelets. These lines in the AFM images correspond to the interfaces seen in Figure 5.9, indicating the existence of a physical crack/step where we see crystalline misorientation.




Figure 5.16: AFM images of a (a) rubrene spherulite and (b) a discontinuity emanating from the nucleus. (c) Polarized optical microscopy image of the spherulite.

Intra-spherulite boundaries with large crystalline misorientations (referred to as “discontinuities”) are optically imaged in Figure 5.9, and turn out to be important for understanding spherulite growth. In Figure 5.15b we saw that a discontinuity manifest as a line in the AFM image. We find that the same is true for spherulite films, and that these discontinuities can be traced back to the nucleus of the crystalline structure. The nucleus of a spherulite annealed at 170 $^\circ\text{C}$ is shown in Figure 5.16a. It is a tall, disordered structure on the order of 5 μm^2 in area. Artifacts due to the AFM tip bouncing prevent us from seeing the left and right sides of the nucleus, but on the top and bottom we can see lines emanating from the nucleus, as indicated with dashed lines. Following these lines out into the spherulite, we see that they persist as boundaries between differently-oriented domains. In some cases

the line indicates a step edge between the two regions, and in others it indicates a small gap between the two regions. In Figure 5.16b, a dashed line marks the boundary between the two regions—the growth directions are indicated by arrows and they are seen to be different. Zooming out and using polarized optical microscopy, as shown in Figure 5.16c, we see that the intra-spherulite interface that we could observe in AFM is optically visible as a discontinuity. Hence we learn that discontinuities emanate from the nucleus and that both the crystalline orientation and morphology differ on either side of the discontinuity.




Figure 5.17: AFM images of (a) small and (b) large rubrene platelet nuclei. (c) Polarized optical microscopy image of the structure shown in panel b.

If the nuclei template the discontinuities that will give structure to the rubrene film, they are worth studying in more detail. The nucleus of a platelet annealed at 140 °C is shown in Figure 5.17a. It is very similar to the spherulite nucleus shown in Figure 5.17b. It is messy, and we can see interface lines emanating from it. Examining several nuclei did not reveal any patterns—spherulite and platelet nuclei appear similar. In either case the nucleus can be much more complicated, but it does not have a strong effect on the ultimate structure. An example of a large nucleus is shown in Figure 5.17b—this one from a platelet. It is about $200 \mu\text{m}^2$, and many crystals can be seen growing out at odd directions. Yet, those structures that are not growing radially outwards tend to get pinched off, so that when zooming out, as shown in Figure 5.17c, an ordered, circular structure is observed.

5.2 In-Situ X-Ray Scattering of Self-Assembling Rubrene Spherulites

In order to better understand the self-assembly of rubrene spherulites, we undertook *in situ* grazing-incidence wide-angle X-ray scattering (GIWAXS). GIWAXS probes the average crystal structure of a thin film, and by rapidly acquiring many GIWAXS images during the thermal annealing process we were able to learn about crystal growth.

5.2.1 GIWAXS Applied to Rubrene




Figure 5.18: Illustration of grazing-incidence wide-angle X-ray scattering (GIWAXS)

GIWAXS is an X-ray scattering experiment that is designed to probe the crystalline structure of thin films. A cartoon GIWAXS experiment is shown in Figure 5.18. The incoming X-ray beam, \vec{k}_i , is incident upon the sample at some shallow angle α , and scatters with some wave vector \vec{k}_f . Thanks to the crystalline structure of the sample, X-rays reflected at certain angles experience constructive interference and produce bright spots on the detector, shown here with orange spots. This diffraction pattern is similar to what one would see in other X-ray diffraction experiments, and how to calculate it is discussed in the next paragraph. Grazing-incidence is particularly useful for rubrene thin films because the angle α can be tuned to be between the critical angles of the rubrene-air and the rubrene-substrate, leading to total internal reflection off of the substrate [177]. For rubrene on TPTPA and glass/ITO, $\alpha \approx 0.15^\circ$, which has the added benefit of producing a long path length through the sample and hence a stronger signal. We performed this work at beamline 7.3.3 of the Advanced Light Source (ALS) [186], where the geometry resulted in a path length through the sample of about $700 \mu\text{m}$.

Let the sample be crystalline with lattice vectors \vec{a}_i , as shown in Figure 5.18 for rubrene. The probability of an X-ray scattering off of a certain point is proportional to the electron

density at that point, which is quite complicated, but we know it must have a periodicity defined by the unit cell. This insight can be shown [3] to imply that strong reflection can only occur if $\vec{k}_f - \vec{k}_i = \vec{G}$, where $\vec{G} = h\vec{b}_1 + k\vec{b}_2 + l\vec{b}_3$ for some integers h, k, l and \vec{b}_i have the form given by

$$\vec{b}_i = 2\pi \frac{\vec{a}_j \times \vec{a}_k}{\vec{a}_i \cdot (\vec{a}_j \times \vec{a}_k)} \epsilon_{ijk}, \quad (5.1)$$

where ϵ_{ijk} is the Levi-Civita symbol. The reciprocal lattice vector \vec{b}_i has a length that is inversely proportional to \vec{a}_i , and it points in a direction orthogonal to \vec{a}_j and \vec{a}_k . The grid of points \vec{G} define all possible diffraction peaks, but we must also enforce conservation of energy, which requires $|\vec{k}_i + \vec{G}|^2 = |\vec{k}_i|^2 \Rightarrow 2\vec{k}_i \cdot \vec{G} = G^2$ (making use of the fact that $-\vec{G}$ is also a reciprocal lattice vector). This is a fairly exacting condition. If the set of \vec{G} form a grid in reciprocal space, it is equivalent to requiring that a circle of radius $|\vec{k}_i|$, known as the *Ewald Sphere*, intersect a grid point in order for it to produce a scattering peak. Furthermore some peaks that are otherwise valid might not appear due to symmetry considerations [3]. The above condition can be shown to reduce to the Bragg condition, $2d \sin \theta = n\lambda$, where d is an effective lattice spacing and n is the order of diffraction [3]. Periodic features with smaller length scales show up as larger scattering angles, hence wide-angle scattering probes structure on the molecular scale.

The strictness of the Ewald sphere construction means that we only see a small slice of reciprocal space on the detector. This can be seen in Figure 5.19, the GIWAXS pattern for a thin film of rubrene spherulites annealed at 175 °C. It shows a grid of peaks, meaning that the rubrene is uniformly crystallized and ordered. The axes show two components of \vec{q} , which is defined as $\Delta\vec{k} = \vec{k}_f - \vec{k}_i$. Yet reciprocal space is three-dimensional. We see a two-dimensional segment, and in particular we are blind to order in x , the direction of beam propagation. What we do measure are the q_z component of each diffraction peak, corresponding to order in the out-of-plane z -direction, and the q_y component, corresponding to order in the in-plane y -direction. Because our rubrene films do not have a preferred in-plane orientation and we probe a large area of the sample, we can take the q_y axis as being representative of all in-plane order (note: it's a bit more complicated than I make it out to be here and in reality one must calculate $|\vec{q}_r| = \sqrt{|\vec{q}_x|^2 + |\vec{q}_y|^2}$ [187], but those corrections are not necessary for the subsequent discussion).

The real space lattice for orthorhombic rubrene is shown in Figure 5.1 and 5.18. The lattice vectors are orthogonal and have lengths $a_1 = 14.14 \text{ \AA}$, $a_2 = 6.86 \text{ \AA}$ (π -stacking direction), and $a_3 = 25.91 \text{ \AA}$ (ideal out-of-plane stacking direction) [169]. Some studies use different notation [188, 189, 168] or find slightly different values (in particular the lattice vectors in Ref. [169] are about 3% shorter than in other studies, perhaps because [169] is investigating spherulites and others look at large single crystals), but by and large this is a well-established structure. Because it is orthorhombic, the reciprocal lattice vectors are simply $\vec{b}_i = \frac{2\pi}{|\vec{a}_i|} \hat{a}_i$: $b_1 = 0.44 \text{ \AA}^{-1}$, $b_2 = 0.92 \text{ \AA}^{-1}$, and $b_3 = 0.24 \text{ \AA}^{-1}$. Given these values we




Figure 5.19: GIWAXS pattern of oriented orthorhombic rubrene

can connect each diffraction peak to a reciprocal lattice vector \vec{G} , which is denoted using the coefficients: (hkl) .

The structure factor and expected powder diffraction pattern for rubrene have been studied by others [168, 190, 191], and we expect the most prominent peaks to be (002) ($|q| = 0.485 \text{ \AA}^{-1}$), (113) ($|q| = 1.25 \text{ \AA}^{-1}$), (211) ($|q| = 1.30 \text{ \AA}^{-1}$), and (313) ($|q| = 1.76 \text{ \AA}^{-1}$). All of these peaks are bright in Figure 5.19. A peak's position along the q_z axis is given by the b_3 coordinate (i.e. the (002) peak is at $q_z = 2b_3 = 0.48 \text{ \AA}^{-1}$), indicating that the crystals are aligned with the a_3 -axis normal to the substrate. The broad smears at larger $|q|$ values are due to the ITO underlayer.

When crystals are not well-aligned the GIWAXS pattern produces rings instead of dots, as is seen in powder X-ray diffraction. These can be seen in Figure 5.20, which were measured on (a) an early thermally evaporated film (same one seen in Figure 5.6) and (b) a solution processed film. The structure has well-defined lattice planes, leading to peaks at well-defined $|\vec{q}|$, but they are oriented every which way, causing the diffraction signal to spread out over all angles, ϕ . The peaks in Figure 5.20a indicate that both orthorhombic and triclinic rubrene are present. The peaks in Figure 5.20 indicate monoclinic rubrene. Neither monoclinic nor triclinic exhibit good π -stacking. A description of these crystal structures can be found in




Figure 5.20: GIWAXS patterns of disordered (a) orthorhombic and (b) triclinic rubrene

Ref. [166] (theoretical) and Ref. [168] (experimental).

5.2.2 Description of *In Situ* GIWAXS Experiment

Because grazing-incidence X-ray scattering can rapidly measure the diffraction pattern of a thin film, it has frequently been used for *in situ* measurements of self-assembly. Selected examples include observing the formation of nanocrystal superlattices [192, 193], organic thin films [194], and functioning OPV cells [195, 196]. In extending this technique to rubrene spherulites there are two primary challenges: avoiding oxygen exposure and ensuring rapid heating to the annealing temperature.

As discussed in Section 5.1.4, the combination of oxygen and light causes amorphous rubrene to form rubrene endoperoxide, which inhibits spherulite growth. Getting the amorphous rubrene films from the nitrogen environment of our collaborators' glovebox at Princeton University to the experimental hutch of the GIWAXS beamline therefore presented a challenge. The amorphous films were first sealed into a vacuum nipple and shipped to Berkeley, where the seal was opened in our glovebox. Each sample to be brought to the ALS was put in a container that was wrapped in aluminum foil (to keep out light) and sealed in a Ziploc bag (to keep out oxygen). These were then sealed up once again in the same vacuum nipple and brought to the ALS, where the seal was opened and the samples stored in a nitrogen dry box. Although now technically exposed to oxygen and light, the combination of the Ziploc bag and aluminum foil kept the films stable for at least a few hours.

Inside the experimental hutch, all windows and light sources were covered with aluminum

foil, and the lights turned off. The sample itself sat inside a box that was flooded with helium. A helium box is generally used to improve signal, because X-rays travel through helium without scattering, but in this case it had the added benefit of limiting oxygen exposure. Once the oxygen level in the box had dropped below 1%, a sample would be taken out of the dry box and rushed to the hutch. Once in the dark, the Ziploc bag was opened and the aluminum foil removed. Using only a faint red bicycle light for illumination (red light is below the band gap of rubrene and hence should not oxidize it), the helium box was cracked open, the sample inserted, and the box quickly sealed shut. This protocol allowed us to produce films of rubrene spherulites, but as discussed in Section 5.2.3 some oxygen exposure still occurred, and this slightly changed the morphology of the resulting films.




Figure 5.21: Temperature rise of rapid heater stage, measured with the stage controller.

In order to ensure rapid heating we worked with beamline scientists Chenhui Zhu and Alex Liebman-Pelaez to purchase the most rapid heating stage we could find—the CER-1-01-00105 ceramic heater from Watlow with an EZ-ZONE controller. Rapid heating is important both to mimic the common situation in which a film is placed onto an already heated stage for annealing, and also to avoid the growth of other crystal structures. Between 110 °C and 140 °C triclinic rubrene is known to form, whereas orthorhombic rubrene forms above 140 °C [185]. The rise of the heater stage temperature when set to 170 °C is shown in Figure 5.21. The stage takes about 10 seconds to reach 110 °C, which is defined as $t = 0$ because that is when crystallization begins. It takes another 10 seconds to get within 5 °C of the desired temperature and a further 5 seconds to reach 170 °C, at which point the temperature is stable to $\pm 1^\circ\text{C}$. As most of the crystallization was found to occur between $t = 20$ and $t = 40$ seconds, we are reasonably confident that our *in situ* annealing studies mimic the conditions that generally occur when making rubrene spherulite films in a glovebox.




Figure 5.22: Rubrene spherulites annealed in situ at the GIWAXS beamline. Optical cross-polarized transmission Images taken above ITO for temperatures (a) 160 °C, (b) 165 °C, (c) 170 °C, (d) 175 °C, (e) 180 °C, and (f) above glass at 170 °C.

5.2.3 Results

Films were annealed for 90 seconds at the following temperatures: 160, 165, 170, 175, and 180 °C. Images of samples annealed at each temperature can be seen in Figure 5.22. As expected the films go from more platelet-like to more spherulite-like, with more frequent branching at higher annealing temperature. At higher temperatures we see that the branching becomes more severe near the edge of the spherulite, with branched dendrites forming. These dendrites are even more prominent in the region above glass, as shown in Figure 5.22f (images a-e are taken above ITO). AFM scans show that these dendrites can be more than to 100 nm tall and are surrounded by areas devoid of rubrene. Presumably this is due to small amounts of rubrene peroxide which inhibits the advancing growth front and causes rubrene to branch and pile up. The fact that the behavior depends on the substrate indicates that it involves interactions with the TPTPA, which has a substrate-dependent roughness (see Figure 5.14), however the exact mechanism is unclear.

Although branched dendrites make the resulting films look somewhat different from those annealed in a glovebox (Figure 5.9), X-ray scattering shows that they have the same crystal structure. Figure 5.23 shows the GIWAXS pattern of rubrene spherulites annealed at the beamline at 170 °C, after the annealing has finished. Although the signal/noise ratio is a bit lower, the pattern is the same as that in Figure 5.19. Growth dynamics can be seen




Figure 5.23: GIWAXS pattern of rubrene spherulites annealed *in situ* at 170 °C (5 second exposure).

in Figure 5.24, which shows the GIWAXS pattern at five time points during the annealing process. Crystalline structure begins to form about 15 seconds after the stage passes 110 °C, with most of the formation happening between 20 and 40 seconds. No intermediate states are seen.




Figure 5.24: Time-resolved GIWAXS of rubrene spherulites during *in situ* annealing at 170.

The dynamics of rubrene spherulite formation are clearly subtle. To understand them better, we zoom in on the (002) peak, which yields the strongest signals and exemplifies behavior we see in all peaks. The (002) peak is shown in Figure 5.25a. It is clearly structured—the signal appears smeared out both horizontally and vertically. A plot of the signal along a q_z




Figure 5.25: (a) GIWAXS image of (002) peak of rubrene spherulite film after being annealed *in situ* at 170 °C. Slices of the peak along the (b) q_z and (c) q_y direction.

slice (vertical line) is shown in Figure 5.25b. There are at least three sub-peaks: one at 0.47 \AA^{-1} (corresponding to the known structure of rubrene single crystals [188]), one at 0.485 \AA^{-1} (corresponding to the known structure of rubrene spherulites [169]) and one at 0.50 \AA^{-1} (origin unknown).

Smearing along the horizontal direction is a bit more complicated. It's not entirely clear if the signal is smeared horizontally or along an arc. Two guiding lines are drawn on the left half of Figure 5.25a: one horizontally at $q_z = 0.485 \text{\AA}^{-1}$ and one along an arc of radius $|q| = 0.485 \text{\AA}^{-1}$. The “whiskers” extending out of the $q_z = 0.47 \text{\AA}^{-1}$ peak appear to follow the arc, but the signal extending from the higher-level peaks appears to be more horizontal. There is important distinction between these behaviors, because as discussed in Section 5.2.4 they each imply a different structural deformation. In Figure 5.25c the signal is plotted along a horizontal line at $q_z = 0.485 \text{\AA}^{-1}$. It is symmetric, with broad wings.




Figure 5.26: Time-resolved slices of the (002) peak of rubrene spherulites annealed *in situ* at 170 °C.

To better understand the growth of this scattering peak, we visualize the dynamics with heat maps, as shown in Figure 5.26. Slices along the q_z and q_y directions (such as those shown in Figure 5.25b,c) are taken at each time point and combined with a pseudocolor plot. From the q_z slices we see that the structure of the sub-peaks evolves in time. The upper peak seems to drift to higher values of q_z (black dashed line) while the lower peak doesn't seem to appear until $t \approx 40$ seconds. By contrast, the q_y slices show no temporal evolution. The peak initially grows in magnitude but its shape and width do not change.

The dynamics described here for the (002) peak are also seen with other peaks and other annealing conditions. We observe sub-peak structure along the radial $|\vec{q}|$ direction that evolves in time, whereas the peaks are smeared along the angular direction but show no dynamics. The first fact can be seen quite clearly in Figure 5.19. Note that the entire diffraction pattern appears to be radially stretched, with multiple sub-peaks at each point. That film was deposited and annealed on the same day in the same glovebox, so the results we find for *in situ* annealing are not artifacts of the films being transported or exposed to oxygen. The meaning of these sub-peaks and the smearing are discussed in the next section.

5.2.4 Interpretation of Peak Structure

So far we have seen that there are two types of heterogeneity in the crystal structure of rubrene spherulites—its diffraction pattern is stretched along the radial $|\vec{q}|$ direction, with multiple sub-peaks, and it is also smeared along the direction orthogonal to \hat{q} . Time-resolved GIWAXS of *in situ* annealed films showed that the stretching evolves in time while the smearing is static. I now discuss the physical interpretations of these observations.

The stretching of the diffraction pattern along the radial \vec{q} direction has a simple potential explanation: the entire crystal structure could in some places be dilated (expanded or contracted). Consider Bragg's law, $2d \sin \theta = n\lambda$ [3]. An overall dilation of the crystal structure is akin to changing all lattice vectors d , which must be compensated for by changing the scattering angle, θ . Expanding the crystal structure shifts all diffraction peaks towards the origin and contracting the crystal structure shifts them away from the origin.

Using the (002) peak as an example, Figure 5.26a shows that the 0.485 \AA^{-1} peak, which has been seen previously for spherulites [169], forms initially. The peak at 0.47 \AA^{-1} , which has been seen for single crystals [188], only becomes prominent around $t = 40$ seconds, implying that single-crystalline platelet regions form later in the process. Looking at the AFM of the nucleus region (Figure 5.17), the interior of the structure does seem to be more chaotic, but growth eventually produces an ordered structure in which the a_2 -axis always points more-or-less radially at every point, and a mixture of spherulites and platelets are present (see Figure 5.22). It would seem as if the many discontinuities templated by the nucleus produce a more strained, spherulite-like structure, but as the spherulite grows there is more space between discontinuities and single crystalline platelets can form.

The third peak has an even higher value of q ($q_z \approx 0.50 \text{ \AA}^{-1}$), implying even more compressive strain. The fact that the peak shifts to higher $|\vec{q}|$ (more strain) as time goes on

could indicate that this signal is coming from the edge of the crystalline region. We have already seen that the edge of the spherulite appears qualitatively different in AFM than the body, with radially-oriented crystallites giving way to a wave-like pattern and eventually a tall lip (Figure 5.15c,d). Studying several films for which the annealing has been arrested at progressively longer times shows that the size of this lip grows from about 5 to 15 nm as annealing progresses, which could correspond to increasing strain as seen in Figure 5.26. Though plausible, this hypothesis will remain speculative unless we can directly measure the crystal structure of the spherulite edge.

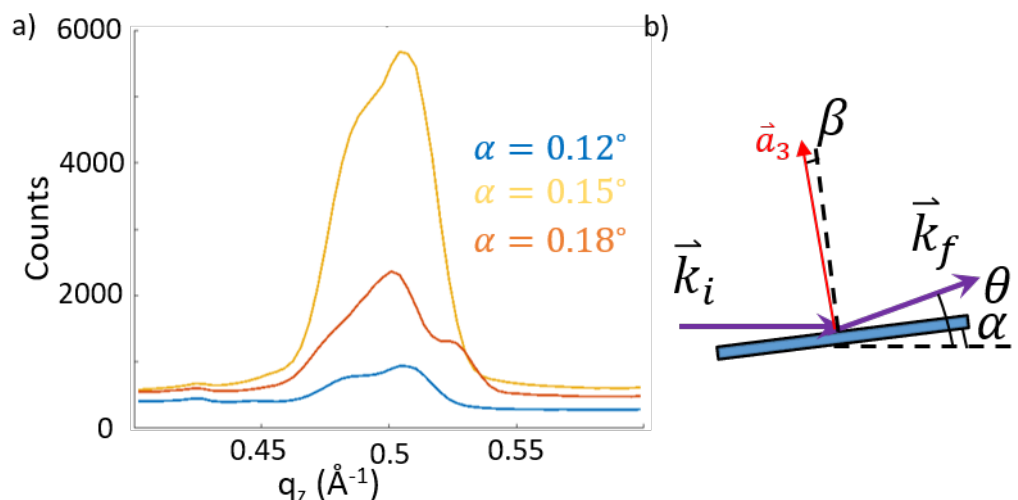


Figure 5.27: (a) Linecuts of the (002) diffraction peak for different values of the sample orientation angle, α . (b) Schematic showing incoming and outgoing X-rays, sample tilt, and crystalline orientation.

Although the existence of multiple sub-peaks indicates that different parts of the film have different morphologies, it is difficult to make quantitative statements about exactly what percentage of the film has a specific morphology. In part this is because we only measure a small subset of all crystallites, and that subset depends on the angle of the beam relative to the film, α . Figure 5.27a shows the q_z slice of the (002) peak for three diffraction patterns that were collected on the same spot of the same film, but for three values of α . Most data presented in this chapter were collected at $\alpha = 0.15^\circ$ because at that angle the X-rays transmit into the rubrene and reflect off of the TPTPA underlayer, producing a strong signal. At other values of α the signal is not as strong, but crucially the profile also looks different. For example at $\alpha = 0.18^\circ$ there is a fairly strong peak at $q_z = 0.53 \text{ \AA}^{-1}$ which does not appear to exist for the other values of α .

Why do the peak positions and strengths depend so strongly on α ? Consider the Laue conditions, which are a way of visualizing the Ewald sphere [3]:

$$\vec{a}_1 \cdot \Delta\vec{k} = 2\pi h \quad (5.2a)$$

$$\vec{a}_2 \cdot \Delta\vec{k} = 2\pi k \quad (5.2b)$$

$$\vec{a}_3 \cdot \Delta\vec{k} = 2\pi l \quad (5.2c)$$

Using the notation in Figure 5.27b, let's try to enforce these conditions for the (002) peak. We have $\vec{a}_1 \cdot \Delta\vec{k} = 0$, $\vec{a}_2 \cdot \Delta\vec{k} = 0$, and $\vec{a}_3 \cdot \Delta\vec{k} = 4\pi$. Because this is an orthorhombic system, Equations 5.2a-b are equivalent to requiring that \vec{a}_3 and $\Delta\vec{k}$ are parallel. We define $\vec{k}_i = |\vec{k}|\hat{x}$ as being along the x -axis and $\vec{k}_f = |\vec{k}|(\cos\theta\hat{x} + \sin\theta\hat{z})$ as being at an angle θ in the $x - z$ plane. Let the angle of deviation between \vec{a}_3 and the substrate normal be β . If \vec{a}_3 and $\Delta\vec{k}$ are to be parallel, then a bit of inspection shows that \vec{a}_3 must bisect the oblique angle between \vec{k}_i and \vec{k}_f , hence $\alpha + \beta = \theta/2$.

Writing $\vec{a}_3 = a_3(-\sin(\alpha + \beta)\hat{x} + \cos(\alpha + \beta)\hat{z})$, Equation 5.2c reduces to $4\pi/a_3k = \sin\theta\cos(\alpha + \beta) + (1 - \cos\theta)\sin(\alpha + \beta)$. Assuming small angles, we have $\theta = 4\pi/a_3k$, which for $\lambda = 1.24 \text{ \AA}$ as is the case at beamline 7.3.3, yields $\theta = 0.096 = 5.5^\circ$. θ and α are fixed for a given experiment, meaning that only one specific value of β leads to the Laue conditions being satisfied. As a result we do not measure all rubrene crystallites in the illuminated volume, but instead only a small subset whose a_3 -axes are oriented at a specific angle β that is slightly off of the substrate normal. β is around 2.6° in the grazing-incidence experiments shown in Figure 5.23-5.25, but the exact value for the crystallites being selected depends on α .




Figure 5.28: Linecuts of (002) peak of rubrene film for several values of incident angle, α . To the left of the jagged line the horizontal axis is stretched for visibility. The dashed line corresponds to studying vertically-oriented crystallites.

To get a better sense of the diffraction pattern due to crystallites with different orientations, we measured the diffraction pattern for a wide variety of α , going up to (and past)

2.75° , which measures crystallites that are oriented with their a_3 -axis perfectly normal to the substrate. Linecuts of the (002) peak for different α are shown in Figure 5.28. Note that the low- α side of the figure is stretched horizontally for visibility. We see that the lattice dilation discussed above is only present for the more tilted crystallites probed (smaller α corresponds to larger tilting, β). For $\alpha > 0.3^\circ$ ($|\beta| < 2.45^\circ$), the diffraction pattern simplifies to a single peak at $q_z = 0.47 \text{ \AA}^{-1}$. It would be interesting to use these data to estimate what fraction of the film is tilted, however doing so is difficult because we are close to the critical angle so reflection is also a strong function of α .




Figure 5.29: One possible type of shear strain on rubrene’s crystal structure and how it would affect both the real and reciprocal space lattices.

As seen in Figure 5.25a the diffraction peaks are smeared both along an arc and along a straight line. Smearing along an arc implies an overall rotation of the crystal structure about an axis in the substrate plane—as seen in Figure 5.20, crystallites that are completely randomly oriented produce diffraction rings instead of peaks. Smearing along the q_y axis is less obvious, but a possibility is shown in Figure 5.29. If shear strain pulls the rubrene molecules up and off of the substrate, then it would be as if one of the in-plane crystal axes, such as \vec{a}_2 , were rotated out of plane by some small angle θ . Calculating the new reciprocal lattice vectors (Equation 5.1) we see that \vec{b}_3 changes from $\frac{2\pi}{a_3}\hat{z}$ to $\frac{2\pi}{a_3}(\hat{z} - \theta\hat{y})$, whereas the other reciprocal lattice vectors do not change to first order in θ . As a result the (002) peak would remain at the same value of q_z , but gain a q_y component, which is what we observe.

5.3 Microdiffraction—Measuring Spherulite Crystal Structure with Fine Spatial Resolution

In order to answer some of the questions raised by the time-resolved GIWAXS data, we turned to microdiffraction at ALS beamline 12.3.2 [197]. Microdiffraction is an X-ray scattering measurement similar to GIWAXS, but the beam is focused to about $5\ \mu\text{m}$ in width by $3\ \mu\text{m}$ in height (as opposed to about $500\ \mu\text{m}$ at beamline 7.3.3). The exact beam spot size on the sample depends on the tilt angle, α , which is generally too large to be considered “grazing incidence.” We used $\alpha = 5^\circ$ to be as similar to the GIWAXS experiment as possible, which lengthened the projection of the beam to 5 by $35\ \mu\text{m}$ on the sample. We used a beam energy of 6 keV. Using the equations derived in the previous section, that means that the (002) peak was diffracted at an angle of $\theta = 0.16 = 9.2^\circ$, hence $\beta = \theta/2 - \alpha = -0.4^\circ$. Therefore the crystallites probed in microdiffraction are oriented more closely to normal than those probed in GIWAXS, where the grazing-incidence angle of $\alpha = 0.15^\circ$ corresponded to $\beta = 2.6^\circ$.

5.3.1 Strain at Spherulite Interfaces




Figure 5.30: (a) Microdiffraction image of (002) peak of rubrene spherulite annealed at $170\ ^\circ\text{C}$. Linecuts of the peak along the (b) q_z and (c) q_y direction.

A microdiffraction image collected from a rubrene spherulite annealed at $170\ ^\circ\text{C}$ is shown in Figure 5.30a. Because the beam is smaller and we are not using a grazing-incidence geometry, the signal is weak and only the (002) peak is visible. Slices along the q_z and q_y directions are shown in Figure 5.30b,c. This is similar to Figure 5.25, but the fact that the experiments probe crystallites at different tilts relative to the substrate means we cannot make a direct comparison. The q_z slice shows a single peak at $q_z = 0.47\ \text{\AA}^{-1}$ and faint sub-peaks on either side. Looking at Figure 5.28, we have agreement between microdiffraction and the GIWAXS with $\alpha = 2.35^\circ$ (when both are probing $|\beta| = 0.4^\circ$).

Comparing q_y slices we see that the peak in microdiffraction is much smaller than the peak in GIWAXS (note: the fact that the peak in Figure 5.30 is not centered at 0 is due




Figure 5.31: Microdiffraction on a rubrene spherulite. The beam was rastered from left to right, and the (002) peak at each point used to produce maps of the q_y and q_z slices. The q_z slice is consistent, but the q_y slice varies around the discontinuity, indicating shear strain.

to poor calibration). This discrepancy remains when measuring GIWAXS at $\alpha = 2.3^\circ$, but is cleared up when we scan around a spherulite and collection diffraction patterns at many spots. An example is visualized in Figure 5.31. The center of a spherulite is shown, with two clear orientational discontinuities emerging. The X-ray spot is roughly the size of the small white rectangle. It is rastered from left to right, tracing the orange path and collecting a diffraction pattern at each point. Taking those (002) peaks and slicing them along both q_y and q_z produces the heat maps. The diffraction pattern is consistent along q_z , but varies strongly along q_y , on both sides of the discontinuity. The pattern is repeated throughout this spherulite and for all spherulites studied across multiple films: at discontinuities, the

(002) diffraction peak varies along the q_y direction. We also investigated spherulite edges, but were not able to resolve any significant variation.

What does this variation at discontinuities mean? First of all it explains why the GIWAXS pattern is smeared along q_y : it is a bulk measurement and hence sees the discontinuities. As discussed in Section 5.2.4, it could indicate shear strain of the crystal lattice. The crystal structure on either side of a discontinuity could be pulled off of the substrate, as shown in Figure 5.29. Because the smearing along q_y was found to have no unusual time evolution in the *in-situ* annealing GIWAXS study (Figure 5.24), this interpretation would imply that this form of shear strain is templated by the nuclei at discontinuities and propagates outward without relaxing. We do see from AFM that discontinuities can be traced back to the nucleus (Figure 5.16). Therefore our understanding of strain and growth dynamics, while incomplete, is thus far consistent—orientational discontinuities are templated by the nucleus, propagate radially, and in spherulites carry strain whereas in platelets that strain is somehow relaxed.

5.3.2 A Lack of Strain at Platelet Interfaces




Figure 5.32: Microdiffraction on a rubrene platelet. The beam was rastered from left to right, and the (002) peak at each point used to produce maps of the q_y and q_z (not shown) slices. The peak does not vary as a function of spatial position.

The same experiment was carried out on a platelet, as shown in Figure 5.32, and yielded drastically different results: the (002) diffraction peak is stable. This is despite the fact

that there is a significant misorientation between the two adjacent crystalline regions investigated, and AFM revealed that large misorientations manifest as visible interfaces for both spherulites and platelets (Section 5.1.5).

Why should misorientations between single crystal platelets be strain-free while misorientations within spherulites carry significant strain? The answer might lie in previous work done on solvent-vapor annealed spherulites of TES-ADT. The authors studied charge transport both within spherulites and across inter-spherulite boundaries for various misorientation angles. They found evidence that trapped molecules at the interface can facilitate interspherulite transport, but that this transport is much slower for high-angle as compared to low-angle grain boundaries [174, 92]. We therefore hypothesize that the platelet interfaces are much “cleaner”—there is a small gap between the two crystals, and no strained molecules trapped at the interface. Within the spherulite, by contrast, many molecules are trapped at the interface by the rapid growth process. They experience shear strain, but also might facilitate charge transport across the interface.

5.4 Local Structure and Dynamics of Rubrene Spherulites Probed Independently Via Several Imaging Techniques

5.4.1 Transient Absorption Microscopy

As discussed in Chapters 3 and 4, polarization-resolved TAM can provide fresh insight into the ultrafast electronic dynamics of a crystalline material. TA measurements have been performed on crystalline rubrene, but polarization-resolution may disentangle competing signals, and the spatial resolution of TAM may reveal differences in electronic dynamics between spherulites and platelets or between different regions of a spherulite.

TA measurements on crystalline rubrene result in an initial large excited state absorption (ESA) feature around 500 nm. The fact that no ground state bleach (GSB) is visible implies that the ESA is obscuring the GSB, meaning that polarization-resolution could be a valuable tool. This ESA changes shape and grows in strength over several picoseconds, remaining strong for at least several nanoseconds [161]. It is broadly accepted that singlet fission is the dominant early-time process, but details remain disputed. One paper suggests that both rapid coherent and slower incoherent fission are observed [198]. Another suggests that the correlated triplet pair forms coherently, and that the slower process corresponds to triplet pair dissociation [162]. Finally, there is a small feature that grows in around 800 nm and one group attributes to triplet absorption [198] while another attributes to polaron absorption [161].

Whatever the ultrafast electronic dynamics are, if TAM is to observe heterogeneity between different regions of the film then the first step is to measure a signal. This is challenging, first because the films are less than 100 nm thick and second because the strongly-




Figure 5.33: TAM spectra of a rubrene spherulite at three time delays. Figure courtesy Jenna Tan.

absorbing a_3 -axis [183] is oriented parallel to the incident beams. However, some preliminary data has been taken and is shown in Figure 5.33. We see the expected growth of the 515 nm peak over about 10 ps, followed by a decay. More, polarization-resolved data is necessary for further analysis.

5.4.2 Infrared Scattering-Type Scanning Near-Field Optical Microscopy

Infrared scattering-type scanning near-field optical microscopy (IR s-SNOM) combines the spatial selectivity of a nanoscale AFM tip with the chemical selectivity of infrared spectroscopy. An infrared beam is focused onto an AFM tip above the sample. The AFM tip acts like an antenna, localizing the light to a sub-diffraction near-field excitation spot. The scattered light is sent through an interferometer, recording the amplitude and phase of the material's response, which correspond to phonon modes. The IR light can either come from a laser, which is more powerful, or from the ALS, which is weaker but extremely broadband (this technique is called synchrotron IR nanospectroscopy, or SINS) [199].

IR s-SNOM and SINS have previously been used to map the local orientation of a polycrystalline organic semiconductor [199], and we hoped that it could do the same for our rubrene spherulites samples. Unfortunately, the SINS signal proved to be too weak to take spatially-resolved measurements. An example is shown in Figure 5.34 for the region near the edge of a messy spherulite. Spatial maps of the total SINS amplitude are shown on the right, and the amplitude and phase of the signal are plotted vs. frequency. By averaging over the entire region, a few peaks are visible, but the signal is weak. An example is indicated with red arrows—resonance manifests as a Lorentzian peak in the phase and a derivative line shape in the amplitude. Work is underway to collect IR scattering spectra using a laser source,




Figure 5.34: Infrared scanning-scattering near-field optical microscopy with synchrotron source. Infrared spectra (amplitude and phase) are shown after averaging over two crystalline regions of a rubrene film on gold. Maps of the two regions are shown on the right, after summing the signal amplitude over all frequencies. Figure courtesy of Omar Khatib, Raschke lab, CU Boulder.

which achieves higher power and hence more signal at the expense of smaller bandwidth.

5.4.3 Transmission Electron Microscopy

In 4-dimensional scanning transmission electron microscopy (4D-STEM), a focused electron beam transmits through a sample, and the diffraction pattern is recorded on a detector. Scanning the electron beam over the sample produces a two-dimensional spatial map of two-dimensional diffraction patterns, hence a four-dimensional data set that contains the spatially-resolved crystal structure of the sample. This technique has been used to map the local orientation of a crystalline organic material [200], and we believe it can map local orientation and strain in rubrene spherulites.

4D-STEM has proven difficult, in large part because we have not been able to transfer the rubrene samples to a TEM grid without breaking them. With platelets we have been able to transfer some small flakes (a few microns in extent), but the spherulites tend to fracture completely. The rubrene, TPTPA, and glass/ITO adhere to each other too well. Furthermore, what does get transferred tends to not lie entirely flat, making data collection difficult. Despite these setbacks, some preliminary data have been collected and are shown in Figure 5.35 and 5.36.

In Figure 5.35, the diffraction pattern at a single point is collapsed down to one dimension. Several peaks are labeled, and in some cases are shifted away from their expected positions, perhaps indicating strain. Note that electron diffraction convention differs from X-ray diffraction convention by a factor of 2π , but the labels match those used in Figure 5.23. Because the electron beam is incident along the a_3 -axis, we should not be able to see the (002) peak, but perhaps the film is tilted.




Figure 5.35: Diffraction pattern at one point on a rubrene spherulite as measured with TEM. Figure courtesy of Steven Zeltmann, Minor lab, UC Berkeley.




Figure 5.36: Strain map of rubrene spherulite, calculated by measuring the shift in several diffraction peaks. An example diffraction image is shown in the upper-left. Figure courtesy Steven Zeltmann.

Figure 5.36 combines many such diffraction patterns to create a strain map of a $1\ \mu\text{m}$ by $1\ \mu\text{m}$ region. An example diffraction pattern from one point is shown in the upper-left, and the displacement of a peak from its ideal position is mapped to strain. The clearest signals come from the (202) and (111) peaks, which both indicate expansion in the same area. More measurements are needed if we are to draw any conclusions. We have recently begun fabricating samples directly on free-standing SiN windows, which are transmissive to electrons and allow for 4D-STEM measurements without transferring the samples.

5.5 Rubrene Spherulite Growth is Potentially Driven by Kinetically Trapped Strain

As discussed above, rubrene’s temperature dependence is anomalous. We observe platelets at lower temperatures and spherulites at higher temperatures, whereas in most materials the opposite occurs: single crystals form at temperatures close to the melting point whereas spherulites emerge at high supercoolings. In this section I discuss the theories of crystallization and spherulite growth, and what makes rubrene unique.

5.5.1 Crystallization Theory Applied to Spherulites

Materials in the process of crystallizing are generally supercooled liquids, and rubrene is no exception. With a melting point of 315 °C [201], a film of amorphous rubrene at room temperature is highly supercooled, and thermal annealing provides enough energy to overcome the reorganization barrier and crystallize. The growth rate at a given temperature, $R(T)$, can be written as

$$R(T) = R_0(T)(1 - e^{-\Delta G/k_B T}) \approx R_0(T)(1 - e^{-\Delta S_m(T_m - T)/k_B T}), \quad (5.3)$$

where $R_0(T)$ is the diffusion-based kinetic growth rate and the rest is a thermodynamic term for crystallization with free energy difference between the melt and crystal ΔG . ΔS_m is the entropy difference between the melt and crystal, and T_m is the melting temperature [202, 203]. At $T = T_m$ there is no crystal growth on average, but as T decreases the thermodynamic factor, and hence the growth rate, increases. For some T far below T_m , $\Delta G \gg k_B T$ and the thermodynamic factor approaches 1. At this point the growth is dominated by kinetics instead of thermodynamics, and further lowering the temperature decreases the growth rate because $R_0(T)$ (to be discussed in the next paragraph) decreases [203]. Rubrene in the temperature range of interest ($T < 180^\circ\text{C}$) is well within the kinetics regime, so we will focus on the term $R_0(T)$. Note that some discussions also include a nucleation term [204], but during rubrene spherulite formation nucleation is fairly rare (that’s how the spherulites are able to grow so large), so we can neglect that term and focus on growth after a crystalline region has nucleated.

In many materials, especially films of supercooled organic molecules, the kinetic growth rate is written as

$$R_0(T) = R_0 e^{-C/k_B(T - T_\infty)}, \quad (5.4)$$

a modified Arrhenius behavior with some divergence temperature, T_∞ , which is below the glass transition temperature [205, 204]. This is the Vogel-Fulcher-Tammann equation, and it can be derived by considering the difference in entropy between crystalline and melt states [206]. However, in practice multiple growth regimes are usually observed, and we can come to a better understanding of the growth rate by considering the microscopic mechanism.

Crystal growth in the kinetic regime is controlled by the ability of atoms or molecules to diffuse. Both rotational and translational diffusion are important for molecules to rearrange and align themselves with the crystal growth front, and according to the Stokes-Einstein equations, both rotational and translational diffusion coefficients are inversely proportional to viscosity, η . Therefore growth can be understood as due to either rotation or translation. One can say that the growth rate is proportional to the translational diffusivity, $R(T) \propto D(T)$, or that it is proportional to the rotation rate, which goes as the inverse viscosity: $R(T) \propto 1/\eta(T)$ [204]. This link is important, because past attempts to explain the temperature-dependent morphology of rubrene has hinged on the fact that the growth rate is faster at higher temperatures, meaning that there is less time for molecules to reorient and that spherulite misorientations are therefore kinetically trapped by the expanding crystal growth front [182]. But if the molecular rotation rate rises in tandem with the crystal growth rate, then a given molecule will always on average explore the same number of orientations before becoming trapped, independent of temperature. The increase in growth rate is therefore incapable of explaining the transition from platelet to spherulite morphology at higher temperatures.

If translational diffusivity and inverse viscosity vary in tandem, how can we ever explain the emergence of spherulites? The answer lies in the fact that at high degrees of supercooling (lower temperature), viscosity and diffusivity decouple—viscosity begins to increase more rapidly than diffusivity decreases, molecules rotate more slowly than they diffuse, and misorientations are frozen in, creating spherulites [204, 207]. This decoupling has been observed in many glass-forming solids and is understood to be a result of dynamical heterogeneity in the supercooled state. As supercooling increases, some regions of the melt become well-aligned and effectively crystalline, whereas other regions retain some mobility. In these limited regions, translational diffusion is facilitated by the motion of nearby molecules. One molecule moves into an open space, leaving an open space behind it for another molecule to move into, et cetera [208, 209, 210]. These pathways do not facilitate rotational diffusion as effectively, hence the decoupling.

Although the decoupling of diffusivity and viscosity at low temperatures has been seen with many organic molecules [204, 202, 207, 211], we need to ask whether or not it is relevant for rubrene at the temperature range we are investigating (140 to 180 °C). The answer appears to be “no.” Interestingly, rubrene has been used as an optical probe to track the rotational diffusion time and to show decoupling of rotation and translation in other small molecule organic semiconductors [207], but those results don’t tell us about rubrene itself. Directly measuring the rotational diffusion time or viscosity of supercooled rubrene is difficult, because as we have seen crystallization occurs in just a few seconds. However, the crystallization rate can be measured by real-time observation, and it displays Arrhenius behavior over a wide temperature range [185]. A decoupling of diffusivity and viscosity would cause the growth rate to deviate from Arrhenius behavior, hence we conclude it is not occurring in rubrene.

Because rubrene displays Arrhenius growth its viscosity is most likely Arrhenius as well, meaning it has low *fragility*. Many floppy organic molecules deviate from Arrhenius be-

havior, which is the definition of high fragility, and has been shown to lead to decoupling between viscosity and diffusivity at fairly high temperatures [202]. Rubrene, however, which is relatively rigid for an organic molecule, appears to behave more like an inorganic material, at least in this temperature range. If the annealing temperature is lowered below 100 °C, rubrene spherulites are observed to grow over the course of several hours [212, 213], a process that is most likely driven by the decoupling of viscosity from diffusivity. But this is not the case between 140 and 180 °C. Hence, we understand why lowering the temperature in this range does not promote spherulite growth, but we are left with the opposite question: why does increasing the temperature should promote spherulite formation? Clearly, a different mechanism is at play.

Having discussed the microscopic interactions that produce crystallization, I will now discuss coarse-grained theories that attempt to predict the resulting structure while still being computationally tractable. The dominant theory is *phase field theory*, which has been successful in understanding the emergence of spherulites in most materials, [214, 171, 215] but does not predict the high-temperature spherulites we observe in rubrene. Basic phase field theory defines two values at each point in a (for our purposes two dimensional) grid: the phase, ϕ , and the orientation, θ . The phase is a measure of how crystalline the film is at a given point—0 corresponds to entirely liquid and 1 corresponds to perfectly crystalline. The orientation is the average orientation of the molecules within that grid point. These parameters are used to calculate the free energy, F , which might have a form similar to

$$F = \int_V d\vec{r} \left[\frac{\epsilon_\phi^2}{2} |\nabla\phi|^2 + f(\phi) + g(|\nabla\theta|) \right], \quad (5.5)$$

where the $|\nabla\phi|^2$ acts to de-mix crystalline and amorphous regions, $f(\phi)$ has a double-well form with local minima at $\phi = 0$ and $\phi = 1$, and $g(|\nabla\theta|)$ penalizes crystalline misorientations. The form of g is important. To leading order in $|\nabla\theta|$, g must be linear or else grain boundaries do not localize [216]. A quadratic term may also be included to reduce curvature at the grain boundaries, so one possible form of g is:

$$g(|\nabla\theta|) = m(\phi)|\nabla\phi| + h(\phi)|\nabla\theta|^2, \quad (5.6)$$

where both $m(\phi)$ and $h(\phi)$ tend to 0 as $\phi \rightarrow 0$ (the liquid state). Extensions of this model have included generalizing F to three dimensions [214], allowing for a multi-component mixture [171], and including “preferred misorientations” for which new crystals preferentially nucleate [171]; however, none of these are relevant for rubrene.

With the free energy defined, crystallization is governed by relaxation dynamics with noise terms added to model thermal fluctuations [171]. In the resulting equations there is a mobility term that controls how strongly gradients in θ or ϕ are translated into changes in those quantities. Those mobilities can be related to the rotational and translational diffusion coefficients. As discussed above, strong supercooling leads to a decoupling of rotational and translational diffusivity. This insight is incorporated into phase field theory by having the

θ mobility to ϕ mobility ratio decrease as temperature decreases. In this way, phase field theory predicts more spherulite structures at lower temperatures [171].

Finally, I will mention some ideas that have been used to describe the growth front, especially when crystallization occurs with surprising speed. Although these are not necessarily relevant for rubrene, we have seen that rubrene molecules pass through a “ridge” structure at the growth front before crystallizing (see Figure 5.15). This structure and the orientation of the near-interface melt may be important for understanding rubrene’s growth dynamics, and should be kept in mind.

In one study, the crystal growth of metals was found to be “barrierless”—phonon modes produced spontaneous ordering in the melt near the crystal growth front, meaning no thermodynamic barrier had to be overcome to reorient molecules [203]. It is unclear if the results of an atomic system can be applied to molecules, but another study that is certainly relevant looked at o-Terphenyl, which consists of three phenyl rings linked together. The authors used fluorescence polarization anisotropy to measure the rotational diffusion time as a function of distance from the crystal growth front. They found a $\sim 10 \mu\text{m}$ strip of enhanced rotational mobility right next to the growth front. This is because the crystalline form is denser than the amorphous form (as is the case for rubrene as well). Crystallization therefore leaves behind a lower-density region that is under tension and in which molecules rotate more easily. The tension repeatedly builds and is relieved by fracture or cavitation, leading to voids [217] and perhaps misorientations. More thought is needed to determine how relevant this mechanism might be in rubrene, and if it can promote spherulite morphology at higher temperatures.

5.5.2 Simulation Details

In order to understand rubrene crystallization as a function of annealing temperature, we created a Monte Carlo simulation of rubrene crystallization starting from an initial nucleus. We use many of the ideas from phase field theory, but simplify the calculation by only allowing pixels to be either crystalline or amorphous. That is, either $\phi = 1$ or $\phi = 0$. The simulation is described below, and complete code can be found in Appendix C. The code was developed and run in a Jupyter notebook [218], and Cython [219] was used to make the computation faster.

Initialization

The simulation is initialized with two parameters: `side_length` and `nucleus_radius`. The simulation space is a square grid `side_length` pixels wide by `side_length` pixels high. Each pixel has four attributes—whether or not it is crystalline (initially all are set to `False`), the crystalline orientation (initially set to 0), whether or not it is part of the nucleus, and whether or not it is “active” (explained below). Orientations are limited to be in the range $[0, \pi)$, and correspond to the a_2 -axis of orthorhombic rubrene. All pixels that are within a distance `nucleus_radius` of the center of the grid are assigned as being part of the nucleus,

and their orientation is set to point radially. All pixels just outside the nucleus are assigned as “active.” Active pixels can be addressed during a Monte Carlo step.

It would be interesting to allow the nucleus to be non-circular. This has not been done yet, but some arbitrary shape could be implemented and the orientation at the edge of the nucleus be set perpendicular to a tangent line to the nucleus at that point. Such a configuration would result in lots of crystallites growing at odd angles, similar to what is observed in Figure 5.17b, and it would be interesting to see what pattern emerges.

With the simulation initialized, the Monte Carlo steps can be run. Each step consists of several parts, described below. The number of Monte Carlo steps to run is arbitrary, but the general practice has been to almost fill the simulation space. If the crystal growth front advances at a rate g (see section “Advancing the Crystal Growth Front”), then the number of steps is $(\text{side length}/2 - \text{nucleus.radius})/g$.

Selecting a Pixel

With each Monte Carlo step, a pixel is selected at random from those that are active. Active pixels are those that are crystalline or adjacent to a crystalline or nucleus pixel.

Generating a New Configuration

If the selected pixel is not crystalline, its state is set to be crystalline. A crystalline orientation is then chosen. Four strategies have been considered so far.

1. Choose an orientation uniformly at random from $[0, \pi)$. This strategy should work in theory, but the downside is that it drastically slows down the simulation because it has to try many orientations before landing on one that is physically reasonable.

2. Discretize orientation into n equally-spaced values on $[0, \pi)$ and select one randomly. This would be faster than #1, but makes it difficult to distinguish between a spherulite structure and lots of platelets.

3. Calculate the average orientation of some randomly-selected subset of crystalline neighbors. If only one neighbor is selected, the pixel adopts the same orientation as that neighbor. If two pixels are selected, the pixel is the average of their orientations. If three or four pixels are selected, we do a three-way or four-way average. In this scheme there are only a few orientations to sample, but such high-symmetry orientations are likely to mimic the stable states of a real crystal.

The question of averaging orientations is subtle, because 0 and π are identical. Given orientations $\theta_i \in [0, \pi)$, the average is calculated as follows. First, construct a unit vector in the $x - y$ plane with orientation $2\theta_i$. Add all of these vectors together, and the orientation of the result is the average of $2\theta_i$. Divide by two to get the average of θ_i . Inspection shows that this procedure produces reasonable results.

4. The same as #3, but with an added random fluctuation. This fluctuation could be drawn from a uniform or normal distribution, but its magnitude should scale with the temperature of the simulation.

Calculating the Energy

The energy of a configuration is the sum of the energies at each pixel interface. For a given pair of adjacent pixels with orientations θ_1 and θ_2 , the energy is a function of misorientation, $\Delta\theta = \min((\theta_1 - \theta_2) \bmod \pi, (\theta_2 - \theta_1) \bmod \pi)$. Because there are two directions one can “wrap around the circle” to go between θ_1 and θ_2 , this function selects the shorter distance.




Figure 5.37: The pixel interfacial energy as a function of misorientation, $\Delta\theta$, for $\epsilon_2 = 0$ (solid line) and $\epsilon_2 > 0$ (dashed line).

In traditional spherulite theories, the energy would be a simple linear or quadratic function of misorientation angle [214]. Our model is different in that it includes a critical misorientation angle, θ_{crit} , beyond which the interfacial energy is constant. Let $-E_0$ be the energy when the neighboring pixels are perfectly aligned, θ_{crit} be the angle such that $E(\Delta\theta \geq \theta_{crit}) = 0$, and let ϵ_2 be a quadratic parameter. The energy function is

$$E(|\Delta\theta|) = \begin{cases} -E_0 + \left(\frac{E_0}{\theta_{crit}} - \epsilon_2\theta_{crit}\right)|\Delta\theta| + \epsilon_2|\Delta\theta|^2 & |\Delta\theta| < \theta_{crit} \\ 0 & |\Delta\theta| \geq \theta_{crit}. \end{cases} \quad (5.7)$$

This function smoothly interpolates between $E(0) = -E_0$ and $E(\theta_{crit}) = 0$, going from a straight line when $\epsilon_2 = 0$ to a more parabolic curve as ϵ_2 increases. Two examples are shown in Figure 5.37.

Note that each pixel in the simulation does not represent one molecule, but some larger number n . Generally, nucleus_radius has been set to 15, and since nuclei are often about $1 \mu\text{m}$ in diameter, this would correspond to pixels that are 33 nm on a side, or $n = 33 \text{ nm}/0.7 \text{ nm} = 47$ molecules (the size of a rubrene molecule in the substrate plane is about $7 \text{ \AA} \times 7 \text{ \AA}$ [169]). The misorientation $|\Delta\theta|$ between two pixels could actually consist of n^2 misorientations (as there are n^2 molecules per pixel), each of which is $|\Delta\theta|/n$ (assuming a linear change in θ from one pixel to its neighbor). Alternatively there could be a sharp discontinuity, but it would involve n molecular pairs. In either case, E_0 should not be the binding energy between a pair of rubrene molecules, but should be increased by a factor of n .

Some consideration shows that Equation 5.7 could plausibly support the formation of both platelets and spherulites, as well as the transition between the two around $170 \text{ }^\circ\text{C}$.

Consider the simplified scenario of three pixels in a horizontal line. The orientations of the left- and right-hand side pixels are, without loss of generality, $-\theta_{offset}$ and $+\theta_{offset}$, where $\theta_{offset} < \pi/4$ (shown in Figure 5.38a). We ask: “what is the energy of this system, as a function of the orientation of the middle pixel, θ ?”

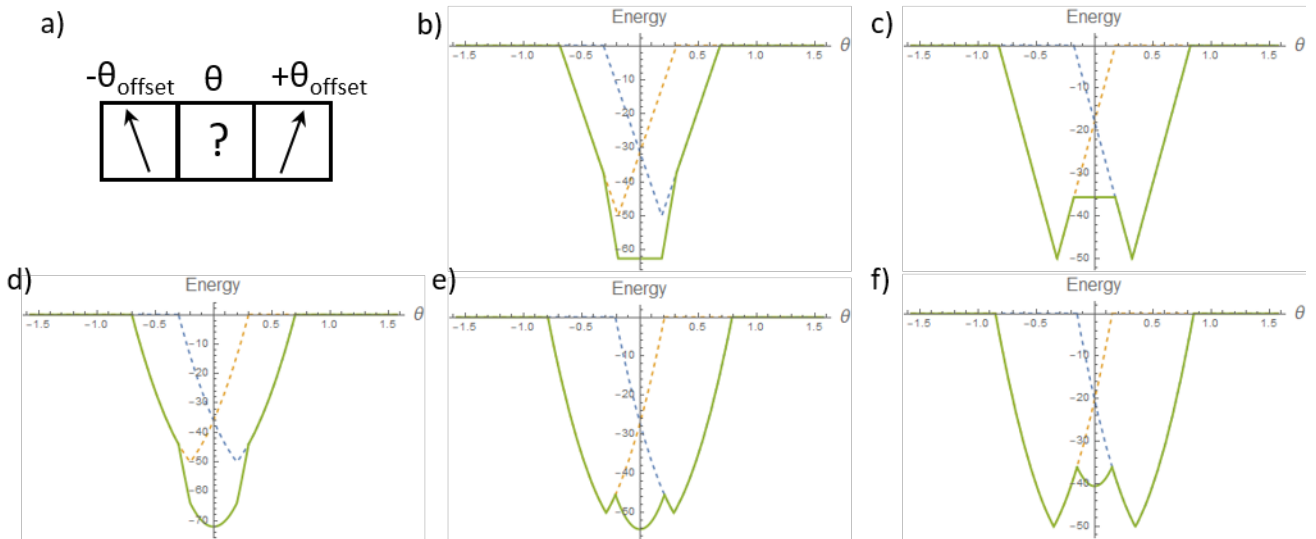


Figure 5.38: Example energy profiles created by (a) two adjacent interfaces, as a function of the orientation of the middle pixel. The orange and blue dashed lines are the energies due to the two interfaces; the solid green line is the total energy. In all plots, $E_0 = -50$ and $\theta_{crit} = 0.5$. Values of ϵ_2 and θ_{offset} are: (b) 0, 0.2, (c) 0, 0.3, (d) 100, 0.2, (e) 100, 0.29, (f) 100, 0.35.

Several examples are shown in Figure 5.38b-f (solid green lines), all for $E_0 = -50$ and $\theta_{crit} = 0.5$. In Figures 5.38b-c, the second order term, ϵ_2 , is set to 0, and θ_{offset} increases from left to right. We see that for smaller values of θ_{offset} (specifically, $\theta_{offset} < \theta_{crit}/2$ as in Figure 5.38b), there is a broad global minimum, indicating that the middle pixel can adopt any orientation in $[-\theta_{offset}, \theta_{offset}]$. On the other hand, for $\theta_{offset} > \theta_{crit}/2$ (Figure 5.38c), there are global minima at $\theta = \pm\theta_{offset}$. That is, the middle pixel aligns with one of its neighbors, but not the others. This is a platelet-like solution. Note that this scenario is only possible because we included the critical angle. In the ordinary case there is no critical angle (set $\theta_{crit} = \pi/2$), so the “platelet” solution arises for $\theta_{offset} > \pi/4$, but it does not make sense for θ_{offset} to be greater than $\pi/4$. Hence, the incorporation of critical angle stabilizes platelets.

Next we include the second order term, setting $\epsilon_2 = 100$ and increasing θ_{offset} in Figures 5.38d-f. For smaller values of θ , the global energy minimum occurs at $\theta = 0$, indicating a compromise between the two neighboring pixels. This is a spherulite-like solution. Increasing θ_{offset} decreases the depth of that central well until local minima form at $\theta = \pm\theta_{offset}$ and eventually become the global minima. Therefore both the “spherulite” and the “platelet”

solution can be stable for the same experimental conditions, depending on the orientation gradient. It is particularly important to note that smaller values of θ_{offset} stabilize spherulites, as this allows both spherulites and platelets to propagate stably. As spherulites grow larger, their misorientation gradient becomes smaller (the spherulite has a larger circumference, so the orientation does not have to change over such a short distance), which is equivalent to θ_{offset} becoming smaller. Therefore, if spherulites are initially the stable configuration then they will remain that way as they grow. On the other hand, because the misorientation boundary is sharp in platelets it does not change as the platelets grow. If platelets are initially the stable configuration, they will continue to be so.

The above analysis indicates that the energy profile described in Equation 5.7 could explain the platelet-to-spherulite transition at 170 °C. Perhaps platelets are the lower energy configuration, but only slightly, and above 170 °C the barrier between them is surmountable. Experimentally we do observe both types of structures, even at higher temperatures, though platelets admittedly become rare (see Section 5.5.4). It is also possible the nucleus matters, as our analysis indicates that nuclei that template gentler misorientations will produce spherulites. Increasing the temperature could change the nucleus formation process in some way that we do not understand, so that the misorientations become less significant (θ_{offset} decreases) and spherulites are more stable. A third possibility is that the spherulite structure is always lower in energy, but platelets are more likely to be templated initially. Forming spherulite would therefore require the system to surmount a potential energy barrier through some transition state, and the height of this barrier could determine the platelet \rightarrow spherulite transition temperature. Such a scenario would not be captured by this simulation.

Accepting/Rejecting the Step

The energies are compared with both the initial (E_1) and new (E_2) configurations. If $E_2 < E_1$, the new configuration is accepted. Otherwise a random number is generated between 0 and 1, and the configuration is accepted with probability $e^{-(E_2-E_1)/k_B T}$, as is standard for Monte Carlo simulations [220].

Advancing the Crystal Growth Front

If the selected pixel has any adjacent and inactive neighbors, they may be activated with some probability r . In the simplest case, r is constant. In a more complex implementation, r may depend on the growth direction because growth along certain crystal axes is faster than others. In this case the fastest growth, r_{max} , would occur when the growth direction is parallel to the crystalline orientation, and the slowest growth, r_{min} , would occur when the growth direction is perpendicular to the crystalline orientation. If the angle between the crystalline orientation and the vector from the pixel to its as-yet-inactive neighbor is ϕ , then the chance of activating this neighbor is:

$$r = r_{min} + (r_{max} - r_{min}) \cos^2 \phi. \quad (5.8)$$

The value of the activation probability should be informed by the ratio of the crystalline growth rate to the molecular reorientation rate. This value, $u_{kin}\tau/a$ has been measured in Ref. [202] for a variety of organic and inorganic molecules. In general it is lower for organics, but given rubrene’s low-fragility behavior it is probably more similar to that of the inorganics. Values for r of about 10^{-2} are therefore reasonable. Values as low as $r = 10^{-3}$ have been tested, but they significantly increase the simulation time and memory, without appearing to affect the results.

5.5.3 Simulation Results

AS of yet only a limited set of simulations have been run, so the results below are highly preliminary. For all of the simulations describe here, the nucleus radius = 15 pixels, the nucleus energy scaling = 0.1, and there is no growth anisotropy ($r_{min} = r_{max} \equiv r$). At first the binding energy was set to 1 eV, which is the correct order of magnitude for binding between rubrene molecules, but as discussed in the previous section (“Calculating the Energy”), about ~ 50 eV is more suitable. The set of calculations shown in Figure 5.41 shown were done with $E_0 = 50$ eV.

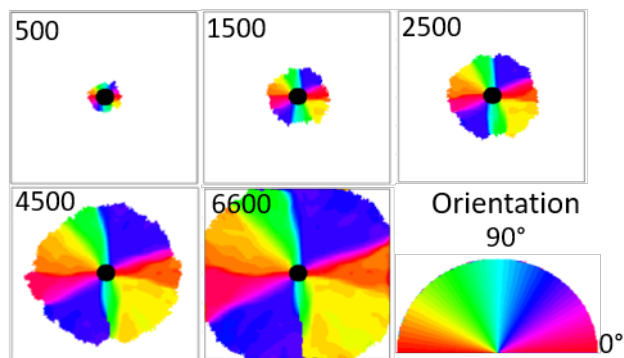


Figure 5.39: Snapshots of simulated rubrene crystallization.

Snapshots of simulated rubrene growth on a 301×301 pixel grid are shown in Figure 5.39 for $t = 500, 1500, 2500, 4500,$ and 6600 “time steps.” In one time step, each active pixel is reoriented on average one time. Color encodes for orientation, as shown in the lower-right panel. The orientation is seen to be smoothly-varying, hence a spherulite-like morphology has formed. The parameters in this simulation were growth rate $r = 0.01$, critical angle $\theta_{crit} = 1.0$, and temperature $T = 30$ meV (75°C). The simulation results depend strongly on growth rate and critical angle, but not on temperature. As shown in Figure 5.38, there is only a narrow range of parameters for which the spherulite and platelet structures are close in energy relative to thermal energy. Outside of this range the system orients into its lowest energy state, regardless of temperature.




Figure 5.40: Results of simulated rubrene crystallization for different values of critical angle.

In order to understand the effect of the critical angle, a series of simulations were performed on a 221×221 pixel grid in which θ_{crit} was varied. Using $r = 0.02$ and $T = 50$ meV (307°C), results for several values of θ_{crit} are shown in Figure 5.40. At $\theta_{crit} = 0.1$ we see uniformly-oriented regions with sharply-defined boundaries between them, reminiscent of platelets. As hypothesized, the inclusion of a critical angle stabilizes sharp interfaces. Increasing θ_{crit} , which is equivalent to making it less significant (in the limit $\theta_{crit} \rightarrow \pi/2$ there is no critical angle), results in more spherulite-like results.




Figure 5.41: Results of simulated rubrene crystallization for different values of temperature.

We have seen that the simulation results can be tuned from platelet to spherulite by changing θ_{crit} , but the ultimate test of this theory is whether or not this transition can be generated by increasing the temperature. Results so far are inconclusive, as only a small region of parameter space has been explored. Preliminary simulations on a 301×301 grid, with $r = 0.01$, $\theta_{crit} = 0.4$, $E_0 = 50$, and $\epsilon_2 = 50$ are shown in Figure 5.41. These values were chosen because a rough calculation showed that for a typical misorientation templated by the nucleus ($\theta_{offset} \approx 2\pi/(2\pi \times \text{nucleus_radius}) = 1/15$, the “spherulite” and “platelet” local

minima would have roughly the same energy (see Section 5.5.2, “Calculating the Energy”). Yet no pattern is seen. I suspect that better results would be achieved if the pixels interacted with their eight nearest neighbors instead of just four (both for calculating energy and for advancing the crystal growth front). At the moment, we observe single-crystalline lines propagating along the four Cardinal directions, which is clearly un-physical. Interacting with more neighbors would mitigate this artifact.

If a region of parameter space is found in which the morphology seems to change from platelet-dominated to spherulite-dominated, then many simulations will be needed to gather statistics, and a way must be found to quantify spherulite and platelet fraction. This can be done by counting the number of neighbors with the same orientation and assigning a pixel as a spherulite if most of its neighbors are misaligned, although care must be taken to ensure that platelet interfaces are not classified as spherulites. Comparing with experiment also requires a detailed knowledge of the experimental results, as discussed in the next section.

5.5.4 Comparing Simulation to Experiment Using Machine Learning

Evaluating the results of the simulation will require a quantitative understanding of how frequently spherulites and platelets form for rubrene annealed at a given temperature. Examining crossed polarizer images from films annealed at several temperatures (Figure 5.22), it appears as if higher temperature leads to relatively more spherulites, but making this impression quantitative would require painstakingly classifying every section of dozens of images.

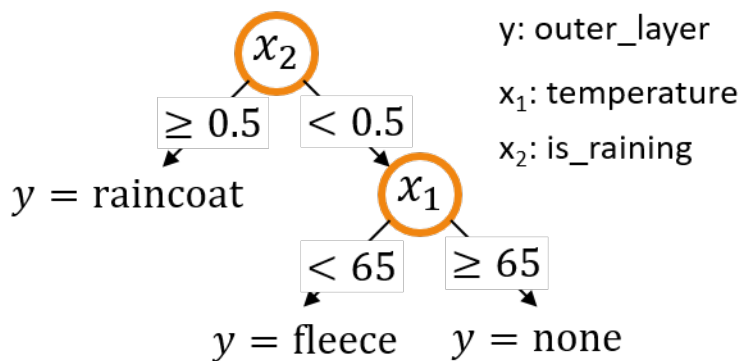


Figure 5.42: Example diagram of a single tree in a random forest classifier.

Instead, this problem was solved with a random forest classifier. Random forest is a machine learning technique that uses a vector of attributes to classify a sample into one of several distinct classes. Classification is done with decision trees, which use a series of logical operations to go from features to output. For example, consider the simple tree diagrammed in Figure 5.42. This tree tries to predict what outer layer someone will wear based on two

features: the temperature x_1 and a variable x_2 that is 1 if it is raining, 0 if not. At the first node the tree compares x_2 to 0.5. If $x_2 \geq 0.5$ then it is raining and it predicts that the person will wear a raincoat. Otherwise it continues to the second node, which compares the temperature to 65. If $x_1 < 65$ it predicts the person will wear a fleece, and if $x_1 \geq 65$ it predicts no outer layer.

Individual decision trees are prone to over-fitting the data, so a random forest classifier builds many trees on randomly chosen subsets of the data and averages their predictions to produce the final output. This procedure reduces the variance in the predictions while maintaining high accuracy [221]. Random forest classifiers have been successfully used in many situations, including classifying land-cover from satellite imagery [222]. The problem faced here is conceptually similar, and random forest was found to work well.

For each film of interest 20 images were collected through crossed polarizers, 10 above glass and 10 above ITO. The images were taken with some pre-determined spacing, so as to represent a random sample of the film. Each image was compressed to 256 by 320 pixels, meaning there are 1.6 million pixels for each sample. The goal was to classify every pixel into one of four categories: amorphous, spherulite, platelet, and miscellaneous (areas where the film was damaged or scratched or otherwise disordered).




Figure 5.43: Training data used for random forest classifier.

Training data was created by manually classifying several regions from 10 different images. The images are shown in Figure 5.43, with colored boxes overlaying the manually classified regions. Red corresponds to amorphous, blue to spherulite, green to platelet, and yellow to miscellaneous. The images have been cleaned up by subtracting the background image, to remove spots, and dividing by the average background signal, to correct for the fact that the lamp brightness varied between films.

Features were chosen in order to bring out the difference between spherulites and platelets. By eye, platelets are uniform in color whereas spherulites are gradually changing. Amorphous regions are also uniform, but they have a predictable brightness level. Miscellaneous regions tend to be highly non-uniform. For each pixel the following quantities are evaluated in its surrounding: the mean, variance, and average gradient. But what defines “surroundings?” One could use all pixels within some fixed distance, but we can be a bit more sophisticated. Consider the image in Figure 5.44a. Clearly the material in the upper left, labeled “1,” belongs to a single feature. There is a band of amorphous material, labeled “2,” and a large structure labeled “3,” which one might see fit to break into smaller regions. How can we work this understanding into the model?




Figure 5.44: Finding regions in an image of crystalline rubrene. (a) Regions one might identify by eye, (b) regions found with agglomerative clustering, (c) those regions overlaid on the image.

An initial attempt to segment the image used edge detection, however many of the edges are too weak to be reliably pulled out of the background noise. Even with Canny edge detection [223], which uses hysteresis to track edges through noisy regions, it was impossible to simultaneously track an edge all the way around a region and also avoid spurious detection.

More successful was agglomerative clustering, a technique in which each pixel starts out as its own “cluster,” and then similar pixels are linked up until a set number of clusters remain [224] (computation was done using the PyAMG module [225]). Example clusters determined with this method and an overlay of the image are shown in Figure 5.44b,c. The algorithm does a good job of finding the borders between areas 1, 2, and 3.

With the image segmented into regions, we can calculate features. The mean and variance are calculated for each region. The local mean, variance, and average gradient are also calculated. “Local” was defined as a square of side length $2d + 1$ pixels for some integer d , centered on the pixel of interest. The variance and average gradient were also normalized by the average signal, to create additional features. The most important features were found to be the variance both within the cluster and within the local region. The relative importance of each feature is tabulated in Table 5.1. The code for this and other aspects of the random forest classifier can be found in Appendix D.

Feature description	importance (full model)	importance (subset models)
Cluster mean	0.12	0.12 ± 0.03
Cluster variance	0.28	0.26 ± 0.04
Cluster variance (normalized)	0.18	0.18 ± 0.04
Local mean	0.07	0.07 ± 0.02
Local variance	0.15	0.15 ± 0.04
Local variance (normalized)	0.08	0.09 ± 0.03
Local gradient magnitude	0.07	0.08 ± 0.03
Local gradient magnitude (normalized)	0.05	0.04 ± 0.01

Table 5.1: Relative importance of features in random forest classifier. The “full model” uses all training data, whereas the “subset models” each use a subset of the training data, as described later in this section. The local area around each pixel was defined with $d = 4$.

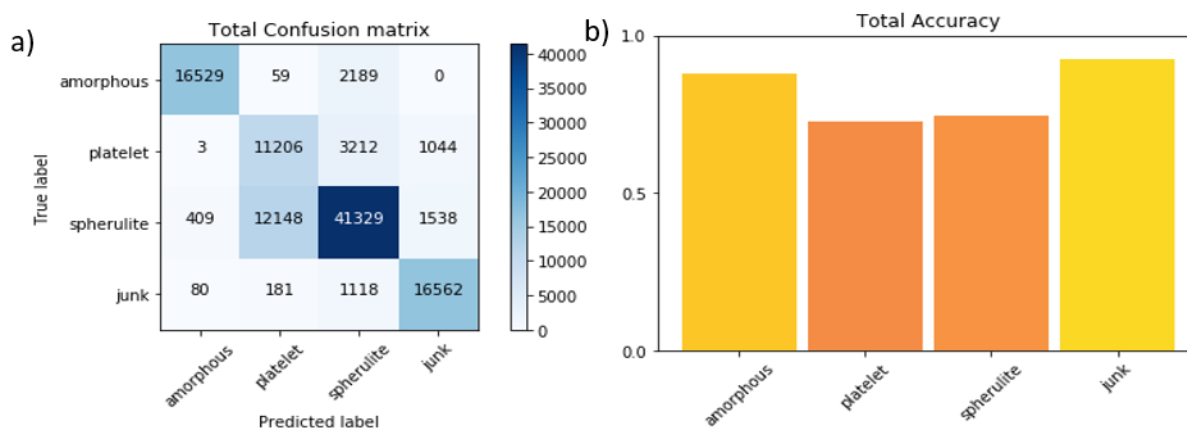


Figure 5.45: Testing the accuracy of a random forest classifier. (a) Confusion matrix and (b) accuracy bar chart for $d = 4$ and $n_{trees} = 80$.

The random forest classifier was created in Jupyter notebooks [218] using the `RandomForestClassifier()` method in Python package `scikit-learn` [226]. The accuracy of the model was tested using k -fold cross validation [227] with $k = 8$. The 79 chunks of training data (rectangles in Figure 5.43) are split into k groups of roughly equal size. In turn each group is used as a “testing” set while the other $k - 1$ are the “training set.” A model is trained with the training data and applied to the testing data to ascertain how accurate it is. Accuracy is visualized with a confusion matrix, shown in Figure 5.45a—it shows how often a pixel of one type is classified in each way. Collapsing each row yields the accuracy bar chart in Figure 5.45b. We see that amorphous and miscellaneous (“junk” in this figure) are often correctly identified, but platelet and spherulite are misclassified about 25% of the time. The confusion matrix shows that they are most often misclassified as each other. This procedure was done for many different parameters, and the best results (shown here) corresponded to a local

area of size $d = 4$ and using $n_{trees} = 80$ trees per forest.




Figure 5.46: Random forest classifier applied to training images.

The results of applying the random forest classifier to every image in the training set are shown in Figure 5.46 (compare to Figure 5.43). The model is seen to do quite a good job classifying the images. For example, consider the image in the lower-right corner. It successfully traces the edges of the amorphous regions in both the upper-right and lower-left corner. It has been told that the material in the upper-left is spherulite and the lower-right is platelet, but has not been told where the boundary is between them. It nonetheless does a reasonable job of drawing that boundary. On the other hand a mistake can be seen in the left image of the top row. The region in the upper-left corner is a spherulite misclassified as a platelet. In general, the model has the most trouble in dark regions where the signal level is low and it can be difficult to distinguish between a smoothly varying and a uniform section.

We can now apply this model to the entire data set in order to determine the relative fraction of amorphous, spherulite, platelet, and miscellaneous for each film. In order to compare to simulation, the most important quantity is the ratio of spherulite fraction to platelet fraction—when a crystalline region nucleates, how much of each type of crystal forms? Calculating this ratio is easy, but we would also like to know its uncertainty. Although there is literature on quantifying uncertainty in random forest estimators [228] (which predict a number as opposed to a class), a good source for random forest classifiers was not found. Inspired by reference [228], the following procedure was used: create 200 models, each of which was trained on a randomly chosen subset of 25/79 chunks of training data. 25 was chosen because it is large enough so that each model has several examples of each class, but small enough so that two arbitrary models were trained on significantly different data sets. This could probably be done more rigorously.

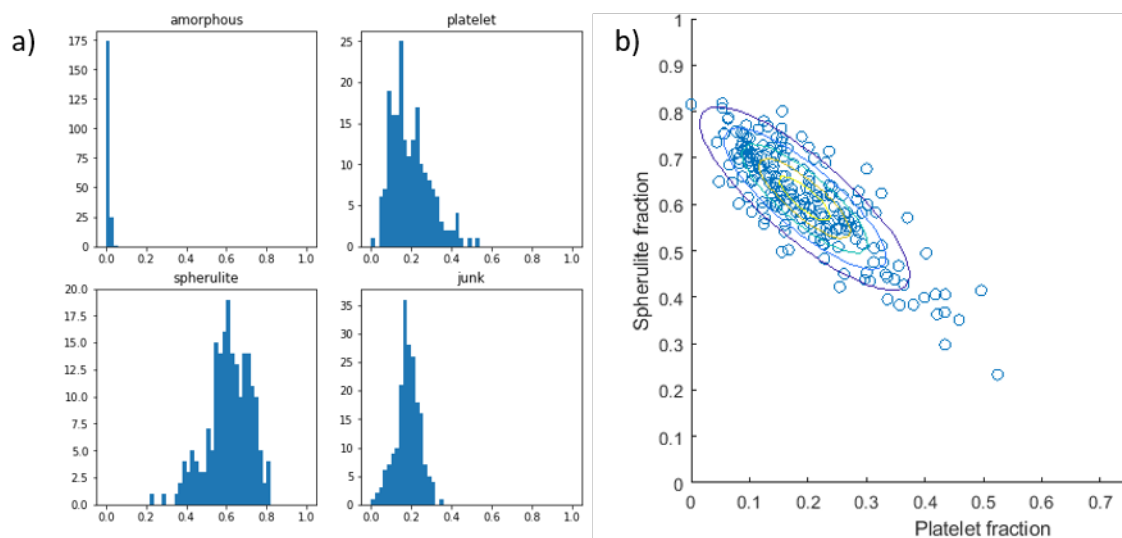


Figure 5.47: Results of 200 random forest classifier models built on subsets of the training data. (a) Histogram of the results and (b) scatter plot of the predicted amount of spherulite vs. platelet.

The results from those 200 models applied to the ITO region of a film annealed at 170 °C in the beamline are shown in Figure 5.47. Figure 5.47a shows a histogram for all four classes, whereas Figure 5.47b is a scatter plot of just the amorphous and platelet regions. They are seen to be strongly anti-correlated, as is to be expected from the confusion matrix (Figure 5.45a). This scatter plot is fit to a Gaussian model and used to calculate the value and uncertainty of the spherulite/platelet ratio. Let the mean spherulite and platelet amounts be s and p , respectively. Let σ_s^2 and σ_p^2 be the variances, while σ_{sp}^2 (which is negative) is the covariance. Then the ratio is $r = s/p$ and the uncertainty is given by Equation 5.9.

$$\delta r = \frac{s}{p} \sqrt{\frac{\sigma_s^2}{s^2} + \frac{\sigma_p^2}{p^2} - 2 \frac{\sigma_{sp}^2}{sp}} \quad (5.9)$$

Results are shown in Figure 5.48 for the ITO regions of the beamline-annealed films. Although the uncertainty is high, there is a clear trend of more spherulites at higher temperatures, which confirms what we see by eye. There is also another data point: we know that at 140 °C there are only platelets. Hence we can compare experiment to simulation and say that the simulation should return only platelets at 140 °C, a moderate amount of spherulites at 160 °C, and almost entirely spherulites at 180 °C. More simulation results are necessary before we can make a direct comparison between theory and experiment, however.




Figure 5.48: Ratio of spherulite to platelet coverage for the ITO region of rubrene films annealed *in situ* at temperatures from 160 to 180 °C.

5.6 Putting It All Together: Steps Toward Multimodal Imaging

The combination of spatially- and temporally-resolved X-ray scattering along with polarized optical inspection and AFM has provided insight into the self-assembly dynamics of rubrene, and in particular why the dominant morphology transitions from platelets to spherulites around 170 °C. Both spherulite and platelet structures include orientational discontinuities where the crystalline orientation changes abruptly. AFM shows that these discontinuities are templated by the nuclei. Microdiffraction shows that shear strain is localized around discontinuities in spherulites, but not in platelets. Because high-angle boundaries (such as the ones between platelets with very different orientations) have been found to act as barriers to charge transport [173], we hypothesize that at poorly-aligned crystalline regions physically separate from each other, entering a regime in which the interfacial energy does not depend on orientational misalignment. Monte Carlo simulations incorporating the above idea have been performed, and although the theory has proven itself capable of stabilizing both spherulites and platelets, it remains to be seen if it can explain the platelet to spherulite transition at higher temperatures. Lastly, we have observed evidence of compressive strain in the GIWAXS patterns, which time-resolved *in-situ* GIWAXS shows becomes more pronounced with time. We hypothesize that this strain is localized at spherulite edges, but higher-resolution measurements are needed to investigate further.

In the future we would like to solve the mystery of the compressive strain, look for

heterogeneity between individual sub-micron crystallites, explain why crystallites have their characteristic size, determine what transient structures exist at the spherulite edges, and use TAM to test if heterogeneity in local structure leads to heterogeneity in dynamics. Answering these questions requires a truly universal sample platform, which is why we have started to fabricate rubrene spherulites on silicon nitride (SiN) windows. Free-standing SiN windows are thin enough to be amenable electron transmission, meaning we can do 4D-STEM on the as-grown sample. The sample is addressable for X-ray scattering and AFM measurements. Finally, as the SiN window is only 10 nm thick it is optically transmissive, and we should be able to perform optical imaging and TAM. Rubrene spherulites present a rich model system for studying molecular self-assembly, and with our suite of complementary techniques it should be possible to make significant progress on this problem.

Chapter 6

TAM (and friends) Investigate Quantum Confinement and Dimensionality Effects in Inorganic Perovskite Nanowires

In this chapter, I discuss studies of quantum-confined CsPbBr₃ nanowires. Lead halide perovskite semiconductors, of which CsPbBr₃ is an example, have been intensely studied for their promise in optoelectronic devices (particularly photovoltaics) and their unique material properties. Here we focus on nanowires, in which nanoscale dimensions along two axes partially confine the excitons, altering their energy levels and dynamics.

We form bundles of aligned nanowires and observe exciton diffusion both along the long- and short-axis of the bundles. We find that transverse diffusion and hence inter-wire coupling is weak, whereas longitudinal diffusion is almost as rapid as in bulk CsPbBr₃. We measure the density of trap states, and find that it is low but still capable of significantly affecting transport.

We also perform polarization-resolved transient absorption microscopy on single bundles. We find that the ultrafast dynamics are essentially the same as in bulk CsPbBr₃, but that the nanowire geometry gently breaks the degeneracy of the band-edge states. Incident light polarized along the longitudinal axis of the nanowires couples to transitions that are about 5 meV lower in energy than those coupled to by transverse-polarized light. Strangely, this shift of 5 meV appears to be the same for nanowires in the weak confinement (10 nm wide) and strong confinement (3 nm wide) regimes. The origin of the shift and its lack of size-dependence remain unclear, although crystalline asymmetry in the nanowires is our best hypothesis.

Finally, we perform polarization-resolved PL studies of nanowires in solution. We find a fluorescence anisotropy that is too large to be due entirely to classical electrodynamics. The nanowires orient non-isotropically in solution, which contributes, but we also conclude that there must be some anisotropy in the strength of the transition dipole moments. Crystalline

asymmetry could also cause these observations, though further measurements are needed to be certain.

6.1 An Introduction to Lead Halide Perovskite Semiconductors

Lead-halide perovskites of the form APbX_3 , where A is a cation and X is a halogen, are hotly studied both for their interesting physical properties and for their promise as the active material in optoelectronic devices. The marquee example is to make photovoltaics out of an organic-inorganic hybrid perovskite, in which the cation is a small organic molecule. Efficiencies above 22% have been achieved [229, 39], which is all the more remarkable because the photovoltaic cells are solution processed and hence highly disordered. Their performance in the face of such uncontrollable messiness seems to be due to a sharp absorption onset [229] and to the presence of the highly-polarizable organic cation, which leads to large polarons that screen carriers from defect recombination [230]. In addition, they are strong absorbers, which means that thin, flexible devices can be fabricated [231, 232], similar to those made out of organic semiconducting thin films.

One drawback of hybrid organic-inorganic perovskites is that the organic molecule is susceptible to damage, a situation that can be remedied by using an inorganic atom instead of an organic molecule [233, 234, 235]. A common choice is cesium, and while its low polarizability means that devices do not reap the beneficial effects of large polarons, the stability (which can be enhanced further by encapsulation [236]), strong absorption, and efficient PL have inspired people to invest significant research effort into improving device performance. Solar cells made of CsPbI_2Br have displayed efficiencies above 10% [237], and substituting a small amount of lead with bismuth can push the efficiency up even higher, towards 15% [233]. Inorganic halide perovskites with cesium have been used to make high-sensitivity X-ray scintillators [238], nanowire lasers with low lasing thresholds and high quality factors [40, 235], polarized photodetectors with high photovoltage [231], and other types of devices [239].

6.1.1 Why Nanostructures Are Interesting

Many of the applications described above involve nanostructures, either nanocubes [234], nanowires [231, 40], or nanosheets [240, 241]. One reason for studying nanostructures is increased tunability. Lead halide perovskites are already considered tunable because by mixing various proportions of different halides (Cl, Br, I) the band gap can be shifted across the visible spectrum [235, 242, 243]. With nanostructures, morphology offers a second knob to tune as pushing dimensions down to the scale of the exciton Böhr radius alters the band-edge structure [241, 244, 245, 246]. This has implications for both electronic structure and dynamics, as will be discussed in Section 6.2.2. In the case of lasers and waveguides, the morphology is crucial to supporting the photon modes [239]. For photodetectors, the

nanowire geometry sets a preferred direction and allows the device to discriminate between different optical polarizations [41, 231].

Nanostructure morphology can also lead to novel physics. For example, in some inorganic semiconductors the transition from nanocube to nanowire changes the ordering of the band-edge states due to coupling between valence bands [247, 248]. This transition has also been shown to change the dominant recombination pathway [249]. In CsPbBr₃, nanocubes were found to display a bright triplet state, which is unusual. This observation was used to infer the presence of a strong Rashba effect, which does not exist in the symmetric crystal structure of bulk CsPbBr₃, but can occur in a nanocube due to symmetry breaking [250]. The large number of surface states also leads to new behaviors. For example, in thin films of Ruddlesden-Popper perovskites, the initially-generated confined excitons dissociate on surface states to form protected free carriers [251]. Finally, both nanowires and nanosheets can self-assemble into unique structures, and their morphology-dependent interactions affect the electronic structure [241, 252].

6.1.2 Previous Studies of CsPbBr₃ Nanostructures

CsPbBr₃, the material discussed in this chapter, has been extensively studied in nanocrystal form, generally about 10 nm on a side. The observation of a bright triplet state [250] was mentioned above. Two studies have used TA and time-resolved PL (TRPL) to measure decay dynamics of nanocubes, both finding strong biexciton interactions [54, 253]. One group did a thorough analysis of the TA dynamics in CsPbBr₃ nanocubes, finding among other things that traps are shallow and few, although substituting some bromine for iodine creates more traps [254]. Furthermore, single-particle studies of CsPbBr₃ indicate that even though trap states are shallow, they exert enough influence to drive the observed blinking dynamics [255]. Finally, THz spectroscopy has been used to argue for extremely high carrier mobilities (4500 cm²/V·s) and diffusion lengths (9 μm) [256], but given that the nanocubes are only 10 nm in size it is unclear what physical relevance these numbers have.

CsPbBr₃ nanowires have been fabricated as well. Most studies so far have focused on the synthesis process, showing that nanowires of a controllable width can be made and that quantum confinement increases the band gap [257, 244, 246], using anion exchange to tune the halide and hence the band gap [258, 242], measuring the crystal structure [246, 259], and performing conductivity measurements [242]. However, very little is known about how nanowire anisotropy impacts electronic behavior. Does confinement along two dimensions affect transport along the third? Do the band-edge states lose their degeneracy, and does the nanowire interact differently with light polarized along its long vs. short axes? How do the nanowires bundle and align with each other? These questions are addressed in the following discussion.

6.2 Properties and Deposition of CsPbBr₃ Nanowire Bundles

Polarization-resolved and spatially-resolved optical techniques can shed light on the electronic structure and dynamics of CsPbBr₃ nanowires, however individual nanowires are just too small. Finding them in an optical setup is very difficult, and the signal to noise ratio is far too low. To get around this problem, we form *bundles*—many nanowires aligned parallel to each other in a micron-scale structure. Bundles are large enough to be found by inspection and they give strong signals, yet the nanowires act independently. Furthermore because they are all aligned, we can measure how optical interactions are affected by light polarization relative to the nanowire. Hence, bundles act as a proxy by which we can investigate the polarization-resolved properties of extremely thin CsPbBr₃ nanowires.

6.2.1 Fabrication, Deposition, and Structure




Figure 6.1: CsPbBr₃ orthorhombic unit cell.

CsPbBr₃ nanowires were fabricated by Jianmei Huang and Mengyu Gau in the Yang lab at UC Berkeley, following procedures from several published works [257, 246, 244]. Wires 10 nm in width and up to several microns long are relatively easy to make, whereas smaller widths (about 3 nm) require more trial and error. We did not measure the crystal structure of our samples, but nanowires made with the same procedures have been investigated using high-resolution transmission electron microscopy [244, 257] and X-ray diffraction [246]. Both measurements indicate that the nanowires have an orthorhombic unit cell [244, 257], which is shown in Figure 6.1, and is not surprising because bulk CsPbBr₃ is orthorhombic [260]. For nanowires less than 4 nm in width, however, the diffraction peaks are too broad to distinguish between the orthorhombic and cubic structures [244, 246]. The growth direction appears to depend on width, with the long axis corresponding to the (001) direction of 10 nm nanowires but the (1 $\bar{1}$ 0) direction of nanowires 5 nm or less in width. A cartoon image

of a nanowire bundle is shown in Figure 6.2, with a few ligands drawn in yellow (in reality, many more are present).




Figure 6.2: Cartoon of CsPbBr_3 nanowire bundle. Yellow lines represent ligands.

Solutions of CsPbBr_3 nanowire bundles were diluted and drop cast onto substrates for further study. It is difficult to quantify the concentration of the solution, so drop casting was done by trial and error—a small amount of the stock solution was taken aside and repeatedly diluted until a good sample was obtained. It is important that the solvent evaporates quickly, because otherwise the bundle will aggregate. When depositing onto TEM grids, a piece of filter paper underneath the grid does the job by wicking away solvent. When depositing onto glass coverslips, they were put in a petri dish on a hot plate heated to 60°C , and the petri dish was regularly tilted to make the solvent slide around and facilitate evaporation.

Figure 6.3 shows TEM images of nanowire bundles. Figure 6.3a shows 10 nm nanowires, which are seen to be straight and monodisperse in width. Lower magnification images (not shown) show that the bundles can be several microns long and remain fairly straight. Figure 6.3b shows 3 nm nanowires. The bundles are thinner, shorter, and more prone to curving.

Figure 6.4 shows images of CsPbBr_3 bundles deposited on glass coverslips and imaged via linear transmission in the TAM setup. These are not representative images—in most regions the bundles are very dense or there are no bundles. One can usually find a transition zone between the two containing isolated bundles. A few bundles are circled. Figure 6.4b shows that straight bundles of 3 nm nanowires are tough to find, but there are at least two specimens in this field of view. One is oriented at about 20° and the other at 90° .

6.2.2 Optical Properties

The properties of a material are affected by quantum confinement if the material has dimensions comparable to or smaller than the exciton Böhr radius. The exciton Böhr radius goes as ϵ/μ , where ϵ is the relative permittivity and μ is the reduced effective mass of an




Figure 6.3: TEM images of (a) 10 nm and (b) 3 nm CsPbBr₃ nanowire bundles. Courtesy of Jianmei Huang.

electron-hole pair [261]. The low-frequency dielectric constant of CsPbBr₃ is calculated to be about 4 [262], and the electron and hole effective masses are 0.128 and 0.134 [250], so the reduced mass is given by $1/\mu = 15.3$. The exciton Böhr radius is therefore about 61 in atomic units, for which the hydrogen atom Böhr radius (0.53 Å) is the unit of length, or 3.2 nm (more precise calculations give 3 nm [263, 243]). Therefore, CsPbBr₃ nanowires 10 nm in width are in the weak confinement regime, and those less than 5 nm in width are in the strong confinement regime.

Absorption and photoluminescence (PL) of CsPbBr₃ nanowire bundles in solution are shown in Figure 6.5. The most prominent indication of quantum confinement is that both the absorption and PL are blue-shifted [264] from their bulk values—the PL peaks shown here are 457 nm for 3 nm nanowires and 523 nm for 10 nm nanowires, whereas bulk CsPbBr₃ has a PL peak at 530 nm [265]. Quantum confinement also increases the electron-hole binding energy and hence the excitonic character of the excited state. This can be seen in the narrowing of the absorption peak [266, 246], particularly for 3 nm nanowires. Increased confinement also leads to more electron-hole overlap, a stronger coupling to the ground state, and hence a shorter radiative lifetime [267]. While the TRPL of bulk CsPbBr₃ can be fit to time constants 23 ns and 233 ns [259], we find faster PL for both 10 nm and 3 nm nanowires (Figure 6.6). The PL of 10 nm nanowires decays with time constants of 2 ns, 16 ns, and 110 ns, while the PL of 3 nm nanowires decays with time constants of about 1 ns, 4 ns, and 15 ns.




Figure 6.4: Optical transmission images of (a) 10 nm and (b) 3 nm CsPbBr₃ nanowire bundles deposited on glass. Several distinct bundles are circled.




Figure 6.5: Absorption (solid lines) and photoluminescence (dashed lines) of 10 nm (orange) and 3 nm (blue) nanowire bundles.

6.2.3 Stability

An early question that arose was how stable the nanowires would be when exposed to oxygen and/or light. To answer this question, several films of 10 nm CsPbBr₃ nanowire bundles were deposited, and the strength of the PL from individual bundles was monitored over the course of several weeks. Ultimately, we found that no encapsulation was necessary, and the nanowires were stable even when they were simply sandwiched between two coverslips




Figure 6.6: Normalized time-resolved photoluminescence of 10 nm and 3 nm nanowire bundles in solution.

with tape (unencapsulated). Results are shown for an unencapsulated sample in Figure 6.7—each point indicates the peak PL spectrum due to a bundle of a certain width, with color indicating how old the sample was when that bundle was measured. Bigger wires are brighter, but there is no trend with age. After 49 days, the bundles are just as bright as they were initially. We did find that intense laser pulses in TAM can damage the nanowires and change their optical properties, but the powers required for our experiments did not cause damage (Section 6.4).

Although the 10 nm nanowires are very stable, thinner nanowires are not. For nanowires less than 4 nm in width, the solution becomes both less bright and more aggregated over the course of a few weeks. They are also prone to damage when exposed to ultrafast laser pulses, even at low energies. This made TAM measurements especially difficult (Section 6.4.3).

6.2.4 Effects of Bundling

Another important question is to what extent bundling and the attendant inter-wire interactions change the properties of an individual nanowire. To test this, we also created solutions of isolated nanowires. The degree of bundling depends on the solvent—toluene promotes bundles, whereas cyclohexane promotes isolated nanowires. Comparing optical measurements done on solutions of both bundled and isolated nanowires indicates that the effects of bundling are minimal—the optical spectra in general red-shift by about 5-10 meV upon bundling. This indicates some interaction, but it is small compared to the effects of quantum confinement. For details, see Section 6.4.2.




Figure 6.7: Stability of CsPbBr_3 10 nm nanowire bundles. The PL is stronger for thicker bundles, but does not diminish over the course of several weeks.




Figure 6.8: Optical image of a thick CsPbBr_3 nanowire.

6.2.5 Control Sample

In order to ensure that our findings are due to quantum confinement, it would be handy to have a sample that is similar in shape to a CsPbBr_3 nanowire bundle, except that it is not in the quantum confined regime. This is achieved with what we refer to as a “thick nanowire” of CsPbBr_3 . Confusingly these are also called “nanowires” in the literature—

they tend to be about 200 nm in width and micron-scale in length, which is similar in size to a nanowire bundle, but large enough so that no quantum confinement occurs. Thick nanowires of CsPbBr₃ have been extensively studied for their lasing properties [40, 268, 235]. An individual thick nanowire, synthesized by Minliang Lai in the Yang lab at UC Berkeley, is shown in Figure 6.8, circled in red. It is about 20 μm long. This and other similar thick nanowires were used as control samples for several of the subsequent experiments.

6.3 Effects of Shape Anisotropy on Diffusion in Nanowires

6.3.1 StrobeSCAT

Stroboscopic interferometric scattering microscopy (strobeSCAT) is a technique, pioneered in our lab by Milan Delor, that takes interferometric scattering microscopy (iSCAT) and extends it to have ultrafast temporal resolution [269]. iSCAT works by shining light onto a sample that is directly pressed against a coverslip, and recording the reflection on a detector. This reflection is a superposition of light reflected off of the coverslip and light scattered off of the sample, and the interference of these two can provide exquisite sensitivity, even down to the single molecule level. Ultimately, scattering is due to the difference in refractive index between the object and the material surrounding it. In stroboSCAT, this idea is extended by using a pulsed pump laser to excite a diffraction-limited spot within the sample. Excitations alter the refractive index of the material, hence when the iSCAT imaging beam arrives some fixed delay time after the pump beam, the excitations change the scattering and hence the image. By comparing the iSCAT images with and without excitation, we create a map of where excited states are within the material. Varying the delay time allows us to visualize the diffusion of excited states [269].

6.3.2 Results of Longitudinal and Transverse Diffusion Measurements

strobeSCAT was performed on several bundles of 10 nm CsPbBr₃ nanowires, eleven of which yielded clear data. An iSCAT image of a bundle is shown in Figure 6.9a, while Figures 6.9b-d show stroboSCAT images for three different time delays. The dark spot in the center of the bundle indicates the presence of excitations, and as time goes on those excitations diffuse along the bundle. Given that the exciton binding energy of CsPbBr₃ was measured to be about 40 meV [243], which is greater than the thermal energy of room temperature, we expect excitons to be the primary excited species.

In order to quantify the diffusion of excitations in CsPbBr₃ nanowires, line cuts were taken along the longitudinal axis of the bundle, as shown by the dashed red line in Figure 6.9b. The three line cuts are shown in Figure 6.10a; clearly the excitons are diffusing outwards. At each delay time the distribution is fit to a Gaussian, and the corresponding σ^2 is plotted




Figure 6.9: strobeSCAT on a CsPbBr₃ 10 nm nanowire bundle. (a) iSCAT image of bundle and (b)-(d) strobeSCAT images for three time delays.

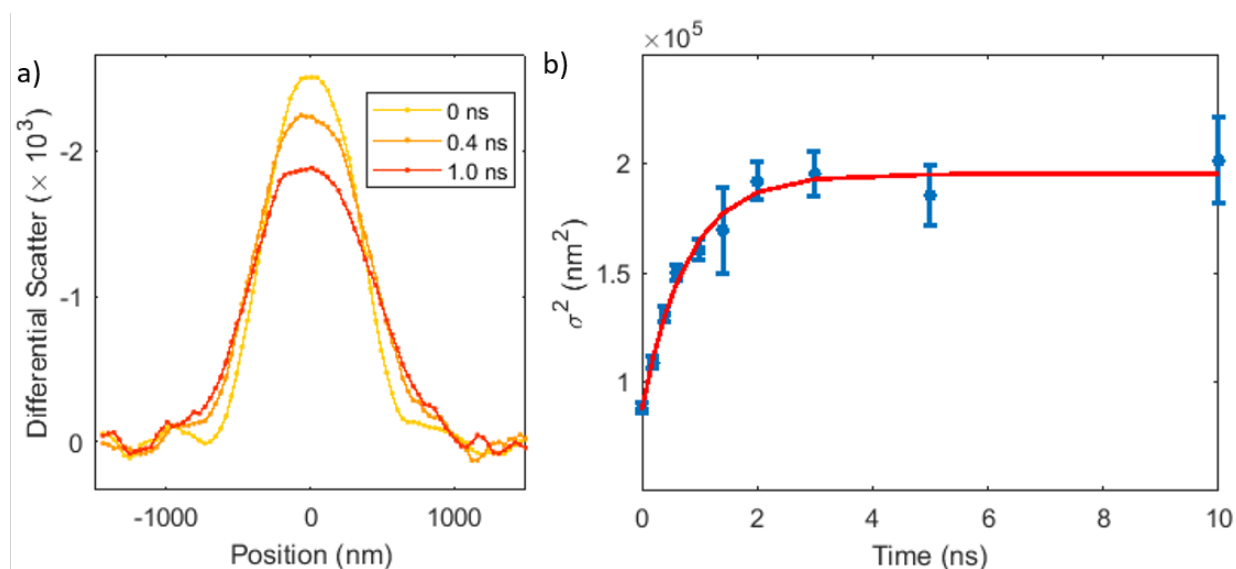


Figure 6.10: Longitudinal diffusion in CsPbBr₃ 10 nm nanowire bundles. (a) Selected line cuts along the long axis of a bundle. (b) Width of the excitation spot, fit to an exponentially decaying diffusivity.

as a function of time in Figure 6.10b. For the first 2 ns the excitations spread out, but then σ^2 plateaus. This behavior is unusual. For ordinary diffusion in one dimension, σ^2 should grow linearly in time:

$$\sigma^2(t) = \sigma^2(0) + 2D_0t, \quad (6.1)$$

where $\sigma^2(0)$ is the variance of the initial distribution and D_0 is the diffusivity [19]. Clearly, this diffusion is not ordinary. The red curve in Figure 6.10 is a fit to an exponentially decaying diffusivity,

$$D(t) = D_0 e^{-t/\tau_d}, \quad (6.2)$$

where τ_d is the “diffusion decay time.” This leads to a distribution that grows as

$$\sigma^2(t) = \sigma^2(0) + 2D_0\tau_d(1 - e^{-t/\tau_d}). \quad (6.3)$$

As seen in Figure 6.10b, Equation 6.3 fits the observed distributions well with parameters $D_0 = 0.69 \pm 0.05 \text{ cm}^2/\text{s}$, $\tau_d = 0.78 \pm 0.08 \text{ ns}$, and $\sigma(0) = 295 \pm 3 \text{ nm}$. But what does it mean, physically? An exponentially decaying diffusivity can be due to hot carrier cooling, but the timescales involved, hundreds of nanoseconds, are much slower than thermalization (see Section 6.4). One common reason for a subdiffusive behavior is a disordered energy landscape. As time goes on the excitations lose energy, and FRET (see Chapter 1) hops become energetically disfavored. This model leads to a power law [19], $D(t) = D_0 t^{\alpha-1}$, an attempted fit to which is shown in Figure 6.11a. Clearly, a power law does not fit our observations.

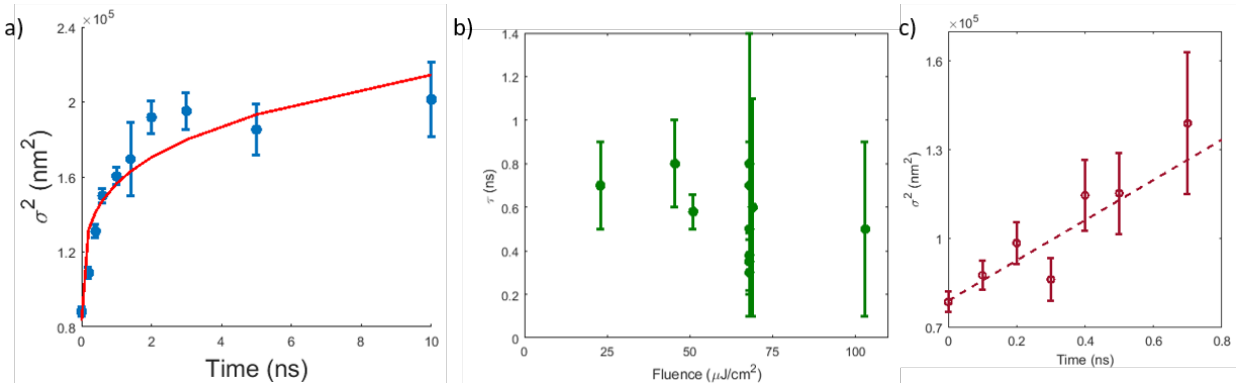


Figure 6.11: Sanity checks related to exponentially decaying diffusivity. (a) Attempted (failed) fit of longitudinal diffusion to a power law. (b) Diffusion decay time for a range of pump fluences, showing no trend. (c) Diffusion in a thick nanowire.

One must also be wary of artifacts. High initial excitation concentration produces non-linear recombination, meaning the center of the distribution decays faster than the edges, hence it *appears* to widen. This widening can be mistaken for initially rapid diffusion. To guard against annihilation, we performed experiments at several different pump fluences. As seen in Figure 6.11b, the diffusion decay time does not show a strong dependence on fluence. Comparing the fluences used here with those used in TA (Figure 6.21), we actually expect that there is some non-linear recombination, but it will be concentrated within the first 100 ps, which is below the stroboSCAT resolution. That could explain why our observed $\sigma(0)$ is bigger than the pump laser spot size (which is about 220 nm), but it cannot explain why the diffusivity decays over approximately 1 nanosecond.

We also studied exciton diffusion in thick nanowires, described in Section 6.2.5. One data set is shown in Figure 6.11c, with a linear fit that yields a diffusivity of $0.39 \pm 0.14 \text{ cm}^2/\text{s}$. The diffusion appears to be ordinary, but our data set is limited. Energy in CsPbBr₃ nanowires of this size is rapidly converted to laser light [40], and the signal is gone within 1 ns. stroboscAT studies of single crystalline CsPbBr₃ have found ordinary diffusion with a constant diffusivity of $1.00 \pm 0.08 \text{ cm}^2/\text{s}$ [269]. After performing all of the checks described here, we are confident that longitudinal diffusivity in CsPbBr₃ nanowires really does decay exponentially.




Figure 6.12: Transverse line cuts showing inter-wire diffusion.

The above discussion involved exciton diffusion along nanowires, but we are interested in inter-nanowire exciton diffusion as well. In one experiment, we excited near the sidewall of a bundle and observed the energy propagate along the transverse direction. Three transverse linecuts are shown in Figure 6.12. They allow us to estimate the transverse diffusivity as $0.22 \pm 0.13 \text{ cm}^2/\text{s}$, but we cannot use any more time delays because after 0.4 ns the excitations reaches the other side of the bundle. Although rough, this value is clearly much lower than the longitudinal diffusivity, and will prove useful for estimating the trap density (see Section 6.3.3).

6.3.3 A Trap-Limited Model of Diffusion

We propose that diffusivity in CsPbBr₃ nanowires appears to decay because of exciton trapping. The model is simple: it includes an initially generated population of mobile excitons, $p_m(x, t)$, with diffusivity D , and a population of trapped excitons, $p_{tr}(x, t)$, that is stationary. Mobile excitons become trapped at a rate k_{tr} , and both populations decay with a rate k_{fl} .

In reality we can expect different decay rates for trapped excitons, but we want to limit the number of free parameters. The equations for the model are therefore:

$$\begin{aligned}\frac{\partial p_m}{\partial t} &= D \frac{\partial^2 p_m}{\partial x^2} - k_{fl} p_m - k_{tr} p_m \\ \frac{\partial p_{tr}}{\partial t} &= k_{tr} p_m - k_{fl} p_{tr},\end{aligned}\tag{6.4}$$

with initial conditions

$$\begin{aligned}p_m(x, t = 0) &= \frac{1}{\sqrt{2\pi\sigma_0^2}} e^{-x^2/2\sigma_0^2} \\ p_{tr}(x, t = 0) &= 0.\end{aligned}\tag{6.5}$$

We can make these equations dimensionless by setting the unit of time to be $1/k_{fl}$ and the unit of length to be σ_0 . Define $k \equiv k_{tr}/k_{fl}$ as the dimensionless trapping rate and let D now be the dimensionless diffusivity. The equations are now

$$\begin{aligned}\frac{\partial p_m}{\partial t} &= D \frac{\partial^2 p_m}{\partial x^2} - p_m - k p_m \\ \frac{\partial p_{tr}}{\partial t} &= k p_m - p_{tr},\end{aligned}\tag{6.6}$$

with initial conditions the same as in Equation 6.5 but with $\sigma = 1$.

Equations 6.6 can be solved exactly. The solution to the first equation, for mobile excitons, is an expression for ordinary diffusion times a decaying exponential with rate $1 + k$:

$$p_m(x, t) = \frac{e^{-x^2/2(1+2Dt)}}{\sqrt{2\pi(1+2Dt)}} e^{-(1+k)t} \theta(t).\tag{6.7}$$

The second equation can be solved by taking its Fourier Transform. We find that the result is a convolution between a decaying exponential of rate 1 and $k p_m(x, t)$, where $p_m(x, t)$ is given by Equation 6.7:

$$p_{tr}(x, t) = \int_{-\infty}^{\infty} k p_m(x, \tau) e^{\tau-t} \theta(t - \tau) d\tau.\tag{6.8}$$

Equation 6.8 can be solved formally [270], but not in closed form:

$$\begin{aligned}p_{tr}(x, t) &= \frac{\sqrt{k}}{2\sqrt{\pi D}} e^{-t+k/2D} \left[e^{-|x|\sqrt{k/D}} \int_{(\sqrt{k/D}-|x|)/\sqrt{2}}^{(\sqrt{k/D}(1+2Dt)-|x|)/\sqrt{2(1+2Dt)}} e^{-\xi^2} d\xi \right. \\ &\quad \left. + e^{|x|\sqrt{k/D}} \int_{(\sqrt{k/D}+|x|)/\sqrt{2}}^{(\sqrt{k/D}(1+2Dt)+|x|)/\sqrt{2(1+2Dt)}} e^{-\xi^2} d\xi \right].\end{aligned}\tag{6.9}$$

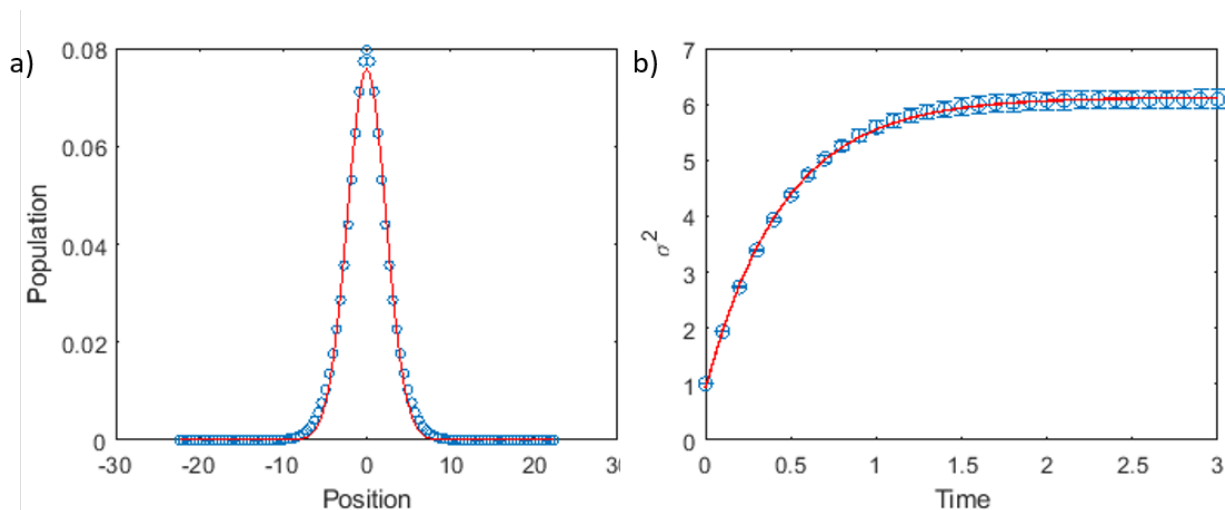


Figure 6.13: Trap-limited decay model in nanowires (Equations 6.7 and 6.9). (a) Distribution at a fixed time fit to a Gaussian and (b) Evolution of σ^2 fit to a decaying exponential.

We cannot derive much insight from this equation, but we can plot its predictions and compare them to our observations. Specifically, we plot the total distribution, $p(x, t) = p_m(x, t) + p_{tr}(x, t)$ and attempt to fit the results to a Gaussian, as is done in Figure 6.10a. Figure 6.13a shows $p(x, t)$ for dimensionless diffusivity $D = 5$ and dimensionless trapping rate $k = 1$ at dimensionless time $t = 1$. It fits well to a Gaussian, and for simulations with lower values of the ratio D/k , the fit is even better. We extract σ from this fit, plot $\sigma^2(t)$ and fit to Equation 6.3 in Figure 6.13. The fit is quite good, as in Figure 6.10b.

From the fit in Figure 6.13b we can extract values of $\sigma(0)$, D_0 , and τ_d and compare them to the input parameters σ_0 , D , and k . We find $\sigma(0) = 0.96$, which is similar to the input value $\sigma_0 = 1$. We find $D_0 = 5.7$, which is slightly different than the input value $D = 5$. And we find $\tau_d = 0.45$, which is significantly different from the input $1/k = 1$.

The above results show that the values of D_0 and τ_d that we extract from fitting do not correspond exactly to the real diffusivity and inverse trapping rate. However, by running simulations for many values of D and k , we can create a map from the fit parameters to the physically meaningful parameters. This is shown in Figure 6.14. The fit vs. actual diffusivity is shown in Figure 6.14a for several values of k : as long as $k > 1$ the relationship $D = 1.145D_0$ approximately holds. The diffusion decay rate, $1/\tau_d$, is shown vs. D and k as a contour plot in Figure 6.14b.

Using these maps, we can convert our fits to actual diffusivities, D , and trapping times, $1/k_{tr}$, for each of the eleven clean data sets. The results are shown as box plots in Figure 6.15. The fluorescence lifetime, $1/k_{fl}$, was set to 2.1 ns based on the results of a time-resolved fluorescence measurement. We find a significant spread in values, with averages $D = 0.78 \text{ cm}^2/\text{s}$ and $1/k_{tr} = 0.79 \text{ ns}$. The longitudinal diffusivity is similar to, but slightly smaller than, the bulk diffusivity of $1.00 \text{ cm}^2/\text{s}$ [269]. Hence, diffusion is not significantly hindered




Figure 6.14: Mapping simulation parameters to those returned by the fit. (a) Fit vs. actual diffusivity, and (b) a contour plot of diffusion decay rate, $1/\tau_d$, vs. trapping rate and diffusivity.

by weak confinement. This extends the work of Tian et. al. into the quantum confinement regime. Tian et. al. studied larger nanowires and nanoplates of methylammonium lead halide perovskites, and found that diffusivity does not depend on shape, but that there is heterogeneity between different samples [271].

The inverse trapping rate, 0.79 ns, is not so easily understood. To convert the trapping rate to a trap density, we perform a simulation. Let λ be the linear trap density on a single nanowire. A square bundle of n by n nanowires, each of length l , is initialized. For each of the n^2 wires, the number of traps on that nanowire is drawn from a Poisson distribution with mean $l\lambda$. Those traps are placed randomly by drawing from a uniform distribution. An excitation starts at an arbitrary position, and at every time step hops either left or right along its current nanowire. Hops to neighboring nanowires are possible as well—these occur with probability D_r , which is the ratio of the transverse to the longitudinal diffusivity. Periodic boundary conditions are used in all dimensions. Steps continue until the excitation encounters a trap. Running many such trajectories produces a trapping time distribution and hence a mean trapping time, $\langle\tau_{tr}\rangle$.

For a given value of λ and D_r , 20,000 trajectories were run. The system size was increased until results converged— $n = 30$ and $l = 20/\lambda$ were found to be large enough. Figure 6.16a shows the results of $\langle\tau_{tr}\rangle$ vs. λ for several values of D_r on a log-log plot. The data fit very well to a line with slope -1, but an intercept that depends on D_r . We can therefore write $\langle\tau_{tr}\rangle \propto A(D_r)/\lambda$, for some constant $A(D_r)$. Converting back from dimensionless to real variables, if D_0 is the measured initial diffusivity along the nanowire bundle, k_{tr} is the measured trapping rate, and $w = 10$ nm is the nanowire width, then the average trap spacing




Figure 6.15: Box plots of (a) diffusivity and (b) inverse trapping rate for eleven CsPbBr₃ 10 nm nanowire bundles.

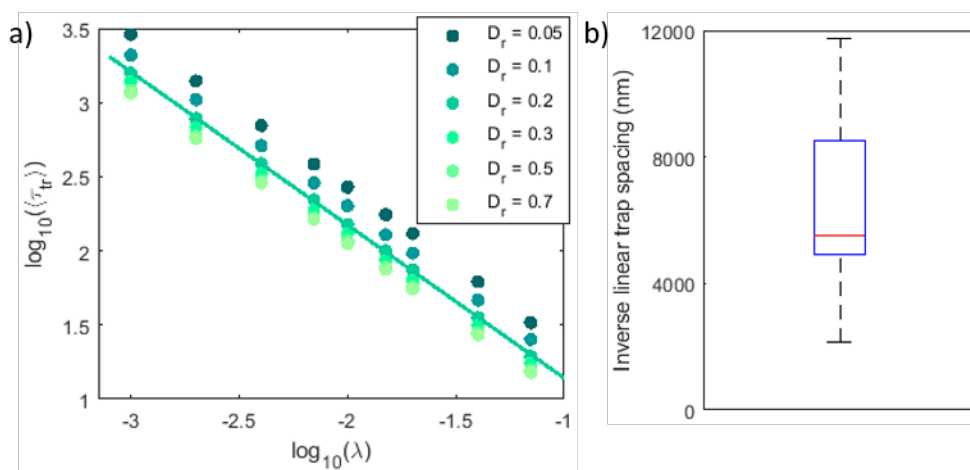


Figure 6.16: Calculating trap density from trapping rate. (a) Simulation results for trapping time vs. trap density, showing the linear fit with slope -1 on a log-log plot. (b) Box plot of inverse linear trap density.

is:

$$\frac{1}{\lambda} = \frac{\frac{2D_0}{1.145}}{A(D_r)wk_{tr}}. \quad (6.10)$$

But what is D_r ? This is where the transverse diffusivity measurement comes in handy—we recorded 0.22 ± 0.13 cm²/s (Figure 6.12), and for that same bundle found a longitudinal diffusivity of 0.76 ± 0.09 cm²/s. The ratio is therefore $D_r = 0.29 \pm 0.17$. Interpolating between

the values of $A(D_r)$ found from simulation, we obtain $A(0.29) = 1.40$, $A(0.29 + 0.17) = 1.22$, and $A(0.29 - 0.17) = 1.90$. There is therefore some uncertainty in A , but it is smaller than the spread in k_{tr} , so we use $A = 1.40$. Using Equation 6.10, a box plot of the average trap spacings for the eleven bundles studied here are shown in Figure 6.16b. The linear trap spacing varies between about $1 \mu\text{m}$ and $10 \mu\text{m}$. As the nanowires are roughly several microns in length, it is possible that the traps are concentrated at the ends of the bundles, although we cannot say for sure. Using a wire width of 10 nm , one trap every 5000 nm corresponds to a volumetric trap density of $2 \times 10^{15} \text{ cm}^{-3}$. For comparison, electrical measurements on nanowires made with the same protocol indicated trap density of $10^{14} - 10^{15} \text{ cm}^{-3}$, so we are roughly in agreement.

6.4 Effects of Anisotropy on Electronic Structure and Ultrafast Dynamics in Nanowires

We set out to discover whether CsPbBr_3 nanowires display anisotropic light-matter interactions. That is, does the nanowire interact differently with light polarized along its longitudinal axis than it does with light polarized along its transverse axis? The following measurements are only possible because of bundling, which incorporates enough nanowires to yield a measurable signal, yet keeps all of those nanowires aligned with the same orientation.

6.4.1 Polarized Local Absorption Shows That Band-Edge Degeneracy is Broken by Nanowire Geometry

One of the simplest optical experiments we can perform with polarized light is absorption. Polarized local absorption spectra are measured in the microscope we use for TAM by passing light from a monochromator through a polarizer and half-wave plate, then through the sample and onto an imaging camera. Varying the output wavelength and the half-wave plate, we measure the amount of absorption at each point in the sample for each optical polarization and wavelength. This technique has previously been used to map the orientation of crystalline domains [71].

Polarized local absorption spectra were measured for a sample of CsPbBr_3 10 nm nanowire bundles, and after using the blank space on either side of each bundle to perform background subtraction, representative results from one bundle are shown in Figure 6.17a. The absorption spectrum shows its first peak around 520 nm , as expected, and it clearly depends on incident light polarization. The spectrum is redder at 0° , which is defined as the longitudinal axis of the bundle, and bluer at 90° . To quantify this shift, the wavelength at which the spectrum reaches half of its maximum is measured for each curve, and the results plotted in Figure 6.17b as red dots. The shift of the absorption onset between longitudinally polarized light and transverse polarized light is roughly 5 meV . The same experiment was performed

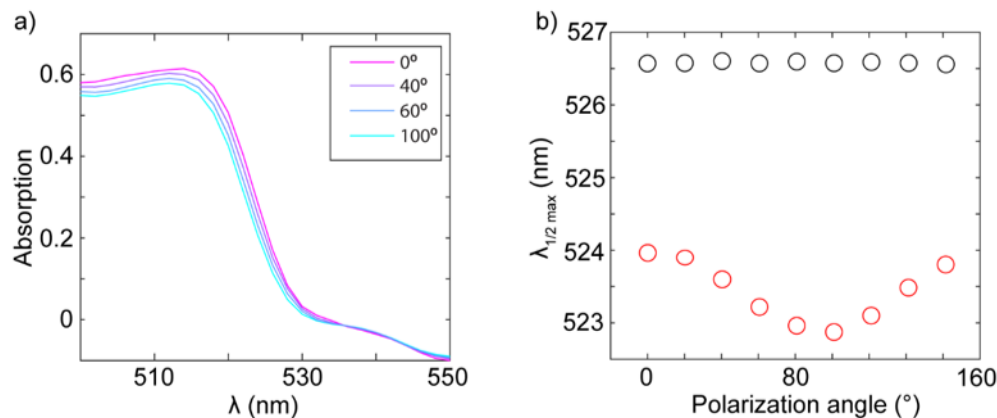


Figure 6.17: (a) Polarized absorption spectra of a CsPbBr₃ nanowire bundle. 0° corresponds to the longitudinal axis. (b) Wavelength at half-max plotted vs. polarization, showing the oscillatory trend for nanowires (red). The same measurement on a thick nanowire (black) shows no polarization-dependence. Figure courtesy of Jenna Tan.

for a thick nanowire (Figure 6.8), and the results, in black circles in Figure 6.17b, show no polarization-dependence. We therefore hypothesize that band-edge degeneracy breaking occurs in 10 nm nanowires, but we cannot say if quantum confinement alters the band structure or if the crystal structure in nanowires is slightly asymmetric. We attempted this measurement on 3 nm nanowire bundles as well, but due to their small size and the ensuing strong diffraction we were unable to measure the local absorption spectrum.

6.4.2 Transient Absorption Spectroscopy of CsPbBr₃ Nanowires in Solution

The next logical step is to perform TAM, to see if the nanowires' electronic transition display polarization-dependent coupling to light. However, I will first take a step back and discuss ordinary transient absorption measurements of nanowires in solution. Understanding the nanowires' ultrafast dynamics using a simpler experiment will make it easier to then interpret the polarization-resolved TAM results.

Transient absorption measurements of a solution of isolated 10 nm CsPbBr₃ nanowires in cyclohexane, excited with a 400 nm pump, are shown as a pseudocolor plot in Figure 6.18. The pump power was 24 μW and the 1/e² diameter, as measured with a razor blade, was about 200 μm, making the peak fluence 60 μJ/cm². Converting this fluence to an initial excitation density requires knowing the absorption coefficient, which has not been measured for CsPbBr₃ nanowires, but it has been measured for both single crystals [272] and nanocrystals [273]. Both works find about 1.0 × 10⁵ cm⁻¹ at 400 nm. Though quantum confinement strongly affects the band-edge absorption spectrum it has a negligible effect at 400 nm [273]. Using this value, we calculate an initial exciton density of 1.2 × 10¹⁹ cm⁻³ =




Figure 6.18: Transient absorption spectroscopy of 10 nm CsPbBr₃ isolated nanowires in cyclohexane.

$12/(10 \text{ nm})^3$. These unusual units are useful because the nanowires are 10 nm in width so that we can imagine a $10 \times 10 \times 10 \text{ nm}^3$ chunk of the nanowire, which would then initially contain about 12 excitons. Given that the exciton Böhr radius is around 3 nm, we can expect non-linear exciton-exciton interactions at early times.

Figure 6.18 shows a strong negative signal around 515 nm, probably due to a combination of ground state bleach (GSB) and stimulated emission (SE). There is a positive, photoinduced absorption (PIA) feature on the high-energy side of the GSB, and a weak negative signal on the low-energy side of the GSB. As this extends far below the band gap, its origin is not immediately clear.

For more insight, a global fit is applied to the TA spectra after subtracting the background and correcting for time-zero. A sequential model is found to fit the data well, with three decaying components and a fourth long-lived constant component (corresponding to nanosecond-scale decay components, including fluorescence). The four evolution associated spectra (EAS, see Section 2.1.3 for explanation) are shown in Figure 6.19a. In order to better see how the spectral shapes evolve in time, the EAS are normalized in Figure 6.19b.

Over the first timescale, $\tau_1 = 0.6 \text{ ps}$, the TA spectrum evolves from the red to the black curve in Figure 6.19. A broad, negative wing on the high-energy side of the GSB peak is lost, as is a positive peak on the low-energy side of the GSB peak. Both of these features can be explained as being due to hot carrier cooling and exciton formation. The pump pulse creates a population of hot electrons that populate the conduction band, bleaching above-gap transitions [274, 275, 276, 254]. If we assume the TA signal is proportional to the state filling, then we can fit the high-energy wing to a Boltzmann distribution [274] and extract a




Figure 6.19: (a) Evolution Associated Spectra (EAS) of CsPbBr₃ 10 nm nanowire TA data fit to a four-component sequential model. (b) Normalized EAS.

temperature $T = 56$ meV (650 K). As the carriers cool, the GSB distribution shifts lower in energy [254]. Comparing the EAS for several pump powers, as in Figure 6.20a, shows that higher fluence leads to a broader wing and hence hotter distribution. A hotter distribution takes longer to cool— τ_1 is as slow as 3 ps for initial excitation densities of 10^{20} cm⁻³. This is likely due to Auger heating—at these densities Auger recombination is dominant, and each recombination event relaxes one carrier at the expense of heating another [274]. In hybrid organic-inorganic lead halide perovskites, a “hot phonon bottleneck” has been identified at modest (10^{18} cm⁻³) excitation densities, in which cooling is inhibited because phonon modes saturate [59]. That does not appear to be the case for CsPbBr₃ nanowires, indicating that there are other cooling pathways involving impurity scattering [274].

Cooling of hot carriers can also explain the briefly-observed photoinduced absorption (PIA) at sub-gap energies. The existence of hot carriers causes band gap renormalization—the conduction and valence bands slightly decrease/increase in energy, respectively, leading to a smaller band gap [274, 277]. Photons with slightly sub-gap energies are able to be absorbed as a result, leading to the emergence of a positive ΔOD . As hot carriers cool, they fill the bottom/top of the renormalized conduction/valence band, and the transition is bleached. This is a manifestation of the Burstein-Moss effect [65].

After the hot carriers cool to form some mixture of excitons and free carriers, the TA spectrum (red line in Figure 6.19), consists of a GSB peak, a PIA at above-gap energies around 450-500 nm, and a mysterious negative ΔOD tail that extends from about 550 to 600 nm. The PIA has been seen in other lead-halide perovskites [276, 254], but has not been commented upon as far as we know. It is most likely due to absorption from the excited state to a higher-lying conduction band state, but its exact shape depends on excitation

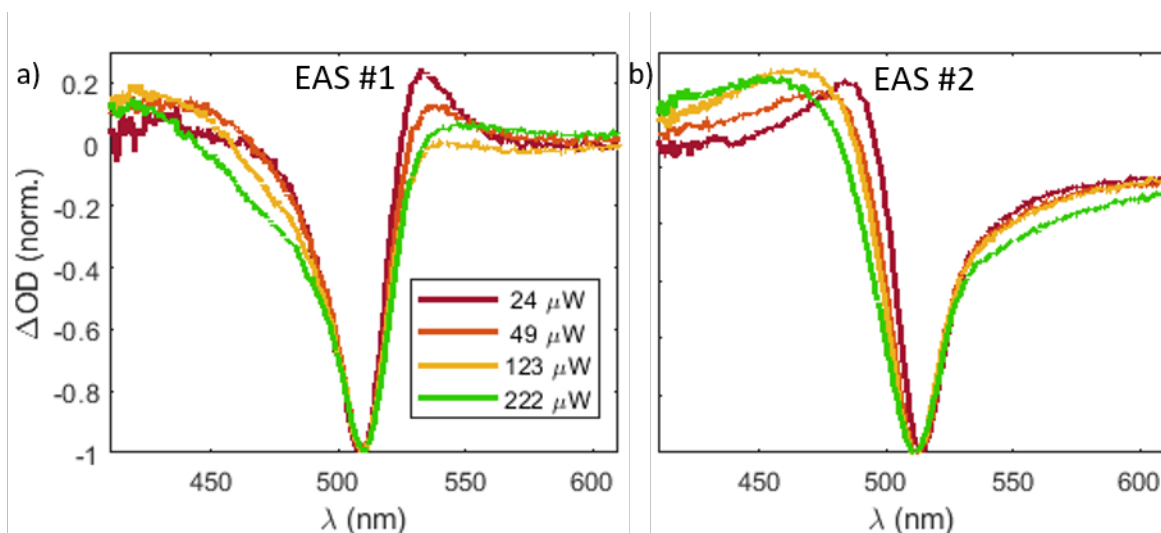


Figure 6.20: Evolution associated spectra (EAS) for CsPbBr₃ 10 nm nanowires in solution, excited with several pump powers.

concentration—using higher pump fluences results in a broader PIA (see Figure 6.20b), and the PIA narrows as time goes on (see Figure 6.19b). Without detailed calculations of the electronic structure of CsPbBr₃, we cannot draw any conclusions about this feature.

The broad negative tail is unexpected. Negative ΔOD signals generally correspond to band-edge transitions, but this signal appears at energies significantly below the band gap. One possibility is two-photon absorption, but the signal’s relative strength does not increase with increasing probe fluence. Another possibility is that the presence of excited states in the nanowires decreases the scattering of probe light, but the fact that stroboSCAT yields a negative signal (Section 6.3.2) implies that the excited nanowires scatter more light, as is expected [269]. This feature has also been seen in nanostructures of the hybrid lead-halide perovskite CH₃NH₃PbBr₃, and is particularly strong in nanowires [277]. It was attributed to electron-phonon coupling allowing for phonon-assisted absorption of low-energy photons. This is in agreement with the broad low-energy tail we see in absorption (Figure 6.5), although the question of why nanowires have such strong electron-phonon coupling requires more investigation.

The decay of the GSB can be fit to two exponentials (25 ps and 130 ps in Figure 6.19), but the time constants decrease with increasing pump fluence. This is an indication that a mixture of linear and non-linear decay processes are occurring. One way to check if the recombination is second order is to plot the inverse of the GSB signal, $1/\Delta OD$. In cases of second-order recombination, $1/\Delta OD$ is linear in time [278]. As seen in Figure 6.21, $1/\Delta OD$ is linear for about 100 ps at 60 $\mu\text{J}/\text{cm}^2$, and about 200 ps at 555 $\mu\text{J}/\text{cm}^2$, indicating that biexciton recombination dominates initially. At longer times this relationship breaks down, as the exciton concentration decreases and nonlinear decay gives way to linear processes




Figure 6.21: $1/\Delta OD$ vs. time for the GSB in CsPbBr₃ 10 nm nanowires. A linear relationship (dashed lines guide the eye) indicate second order dominates initially.

(including fluorescence). Similar kinetics have been used to extract biexciton and trion lifetimes in nanocrystals of CsPbBr₃, for which a countable number of excitons exist on each structure [279, 54, 253]. However, a concept like “biexciton lifetime” is not so meaningful for a nanowire because diffusion is necessary for the excitons to interact, hence the biexciton decay rate depends on exciton concentration.

Having established that 10 nm nanowires of CsPbBr₃ have ultrafast dynamics similar to those of other lead-halide perovskites, we now investigate whether bundling affects those dynamics. The short answer is that it does not. TA spectra for isolated and bundled 10 nm CsPbBr₃ nanowires are shown in Figure 6.22 for $\tau = 1$ ps and $\tau = 10$ ps. Bundling induces a slight red-shift, indicating weak inter-nanowire coupling, but the spectra and dynamics are essentially unchanged.

Lastly, we consider the dynamics of ultrathin nanowires. A pseudocolor plot of the time-resolved TA spectrum for CsPbBr₃ 3 nm nanowire bundles in solution is shown in Figure 6.23a. The features and dynamics are largely similar to those for thicker nanowires, but the GSB signal appears to be longer-lived. To see if this is really the case, the GSB decay is shown in Figure 6.23b, along with the GSB decay for 10 nm nanowires. Similar powers are being compared (101 μW for 3 nm nanowires, 90 μW for 10 nm nanowires), and both samples are bundles of nanowires, so the conditions are similar. The signal in thinner nanowires clearly decay more slowly: after 10 ps of rapid decay they settle into first-order kinetics. A global fit returns one sub-ps component corresponding to hot carrier cooling, one decay component




Figure 6.22: Transient absorption spectra at $\tau = 1$ ps and $\tau = 10$ ps (offset for clarity), comparing solutions of bundled and isolated CsPbBr₃ 10 nm nanowires.

of about 10 ps, and a slow decay component of about 2 ns, which matches up with the first component of the TRPL (Figure 6.6). The paucity of higher-order decay could be due to the nanowires' width—at only 3 nm wide, a given number of excitons are distributed across 16 times as many nanowires as they would be for nanowires 10 nm wide, and so a given pair of nearby excitons are more likely to be on different nanowires, inhibiting their ability to interact and recombine.

6.4.3 Polarized Transient Absorption Microscopy Measurements

Having discussed the TA spectra and dynamics of CsPbBr₃ nanowires, we find individual 10 nm nanowire bundles (as seen in Figure 6.4a) and study them with polarized light in the TA microscope. The TA spectra and dynamics are the same as those measured in solution for comparable fluences. We used both 400 nm and 488 nm pump light, and polarize the probe pulse along both the longitudinal and transverse axis of the bundle. The results were identical, except that the TA spectrum shows a slight dependence on the probe polarization. Normalized TA spectra for specific time slices are shown in Figure 6.24, contrasting probe light parallel (red) and perpendicular (blue) to the longitudinal axis of the nanowire bundle. There is a small, but persistent shift between the two sets of curves. The GSB under parallel probing is about 5 meV lower in energy than under perpendicular probing, just as we observed in linear absorption (Figure 6.17). Several bundles were studied with a variety of orientations in the lab-frame, and in all cases the results were the same.

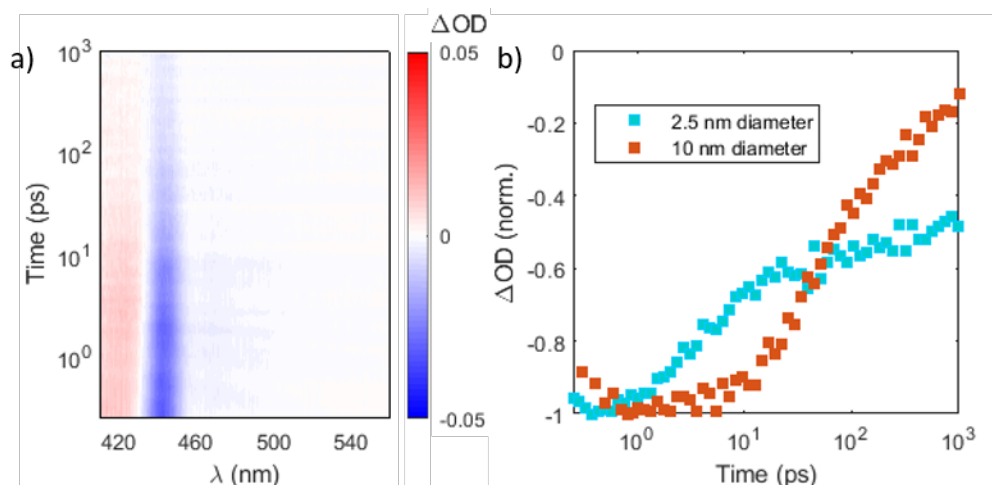


Figure 6.23: (a) TA plot of CsPbBr₃ 3 nm nanowires in solution. (b) Decay of GSB signal, comparing 3 nm and 10 nm nanowires.




Figure 6.24: Selected normalized TA slices for an individual 10 nm CsPbBr₃ nanowire bundled probed with polarized light. Rotating the probe polarization from perpendicular to parallel (relative to the longitudinal axis of the nanowires) shifts the TA spectrum about 5 meV.

The presence of a polarization-dependent GSB suggests again that the degeneracy of the band-edge states has been lifted and that the valence→conduction transitions with TDMS along the longitudinal axis of the nanowires are lower in energy. We cannot say if this shift is due to the width of the nanowire or to an asymmetry of the crystal structure that develops in nanowires.

The TA spectra in Figure 6.24 are normalized, but the un-normalized spectra show

another interesting feature. The overall strength of the TA signal is about 20% lower for transverse-polarized light than it is for longitudinally-polarized light. This is likely due to a well-known effect in nanowires, in which the mismatch in dielectric constant between the nanowires and their environment means that light polarized along the transverse dimension experiences a reduction in electric field and hence a reduction in both absorption and emission cross sections. If the relative dielectric constant between the nanowire and its surroundings is ϵ_r , then the optical intensity is reduced by $\left| \frac{2}{1 + \epsilon_r} \right|^2$ [280]. Given that the optical-frequency dielectric constant of CsPbBr₃ is about 5 [262], a single nanowire in air would be expected to display an absorption of transverse polarized light that is only 10% the absorption of longitudinally polarized light. However, we are studying bundles instead of single nanowires, so the environment for a given nanowire is some combination of other nanowires, ligands, and air. Hence this effect is attenuated.




Figure 6.25: Normalized TA slices for an individual 3 nm CsPbBr₃ nanowire bundle probed with polarized light. Rotating the probe polarization from perpendicular to parallel (relative to the longitudinal axis of the nanowires) shifts the TA spectrum about 5 meV. Gaussian fits to the GSB peaks are shown in an inset.

We also located and performed TAM measurements on bundles of 3 nm nanowires. This was extremely difficult, as even under gentle irradiation (70 nW, or 250 $\mu\text{J}/\text{cm}^2$, which produces a very weak signal), the nanowires degrade within two minutes. Complete TA dynamics are out of the question, but we were able to measure the TA spectra at two

polarizations (parallel and perpendicular) and three time points (-5 ps, 5 ps, and 15 ps). The normalized results for $\tau = 15$ ps are shown in Figure 6.25. There is very little difference between the two spectra, but zooming in on the GSB peak we can see a slight shift from parallel to perpendicular polarization. Fitting each peak to a Gaussian, we find that the shift between the peaks' centers is 5.4 ± 0.3 meV.

Why is the energy splitting seemingly independent of nanowire width? Why should nanowire geometry split the degeneracy in the first place? Let's consider other types of nanowires. In GaAs nanowires, the optical absorption is predicted to depend on light polarization, principally because the strength of a given transition depends on polarization. Different transitions depend on polarization to different degrees, so as the polarization is rotated some transitions are coupled to more strongly than others and the absorption line appears to shift [41]. CdSe nanowires also show a broken degeneracy, and the orientation of the lowest-energy transition changes depending on the ratio of nanowire length to width [248]. Understanding these observations and making similar predictions for CsPbBr₃ would require sophisticated calculations that take into account band coupling [41], electron-hole exchange interactions [281], spin-orbit coupling, and the crystal field term [263]. Doing these calculations in CsPbBr₃ is presently infeasible, largely due to the strong spin-orbit coupling.

The fact that the energy splitting is independent of nanowire width hints that it may not be due to confinement, but instead due to the crystal structure. Deformation would break the symmetry and hence the degeneracy of the band-edge states, as has been shown to occur in CH₃NH₃PbI₃ [282]. The orthorhombic structure of CsPbBr₃ is symmetric, and nanowires appear to have the same structure, [244, 246] so any deformation would have to be small enough so that it is not observed in high-resolution TEM experiments. Single-particle polarized PL studies of CsPbBr₃ nanocubes have found that a mixture of tetragonal and orthorhombic structures exist at low temperatures, and the two structures have distinct polarization-dependent PL spectra [263]. Therefore it is reasonable that something similar occurs for nanowires.

6.5 Investigating Fluorescence Anisotropy in CsPbBr₃ Nanowires

Polarization-resolved PL of nanowires in solution can provide information on the anisotropic orientation of nanowires or their transition dipole moments (TDMs). As a final characterization, samples were therefore diluted in either toluene (bundles) or cyclohexane (isolated nanowires) until the absorption through a 1 cm cuvette was less than 0.3. Samples were then placed in a fluorimeter and excited with a ~ 100 ps pulse from a diode laser. The excitation laser was rotated so that its polarization was alternately vertical (V) and horizontal (H) in the lab frame. Emission was collected in a right-angle geometry, with an emission polarizer rotated both vertically and horizontally. This results in four different measured intensities, I_{VV} , I_{VH} , I_{HV} , and I_{HH} , where the first subscript letter denotes the excitation polarization

and the second letter denotes the emission polarization.




Figure 6.26: Transmission efficiency through the fluorimeter monochromator for horizontally polarized light relative to vertically polarized light. The black dashed line indicates $\eta_H = \eta_V$.

To compare polarized spectra we must correct for the fact that the measured intensity depends on both the excitation conditions (which depends on how the laser is inserted) and the efficiency of optical transmission through the detection monochromator (which is a function of emission wavelength and polarization). The monochromator transmission efficiency ratio was measured by studying a small “control” molecule for which the rotation timescale is much less than 1 nanosecond. The PL of this molecule was measured for both V and H polarized emission light, and since the rapid molecular rotation means that the two should be the same, we can say that any discrepancy between $I_{VV}(\text{control})$ and $I_{VH}(\text{control})$ must be due to the efficiency. Let the transmission efficiencies for V and H emitted light be η_V and η_H . Let primes denote the real emitted (as opposed to measured) intensity. We must have $I'_{VV}(\text{control}) = I'_{VH}(\text{control}) \Rightarrow I_{VV}(\text{control})/\eta_V = I_{VH}(\text{control})/\eta_H \Rightarrow \eta_V/\eta_H = I_{VV}(\text{control})/I_{VH}(\text{control})$ for each wavelength. The measured value of η_V/η_H is shown in Figure 6.26. In the subsequent discussion all intensities have been corrected for this factor.

In Figure 6.28 we also vary the excitation polarization, which involves physically rotating the diode laser and re-inserting it into the fluorimeter. It was later found that this rotation changes the measured intensity, perhaps because the output port is not perfectly centered on the laser and hence rotating the laser changes the beam pointing, the focal spot in the sample, and ultimately the PL intensity. The discrepancy was usually about 5%, but depended on the exact position and the laser used. We therefore decided to re-perform the measurements, and after measuring the nanowires for a given laser insertion immediately study the small control molecule. This will tell us the proper correction factor to apply to that particular

data set. We therefore expect the data in Figure 6.28 to change slightly, but expect the qualitative discussion below to still be valid.

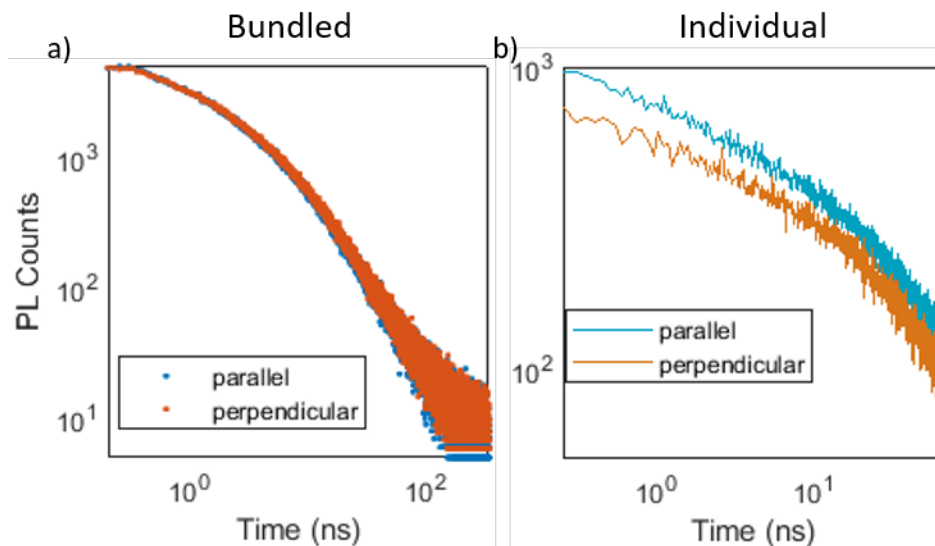


Figure 6.27: Time-resolved polarized PL of 10 nm CsPbBr₃ nanowires: (a) bundled and (b) isolated.

The TRPL of 10 nm nanowires at 520 nm is shown in Figure 6.27, both for bundled nanowires in toluene and individual nanowires in cyclohexane. The nanowires were excited with vertically polarized light. We see that the bundles display isotropic emission, whereas the individual nanowires display anisotropy. Recall from Section 6.4.3 that the mismatch in dielectric constant between the nanowires and their environment leads to diminished absorption and emission for light polarized along dimensions that are much less than the optical wavelength. For isolated nanowires, this results in fluorescence anisotropy. We see that the absorption is time-independent, implying that the nanowires are not rotating significantly during the fluorescence lifetime. Therefore in subsequent anisotropy measurements we record the steady-state PL spectra; TRPL is unnecessary. The lack of anisotropy for bundles we explain as being due to the fact that the bundles are much thicker than the relevant optical wavelengths (see Figure 6.4a).

Polarized fluorescence spectra are shown in Figure 6.28 for 10 nm nanowires excited at both 465 nm and 408 nm, as well as for 3 nm nanowire bundles. Because the 3 nm nanowire bundles are generally thinner than the optical wavelength (as evident in Figure 6.3b and 6.4b), we still see anisotropy when they are bundled. In fact we see anisotropy in all samples, but what exactly does it mean? Can it be explained with classical electrodynamics? The answer turns out to be “no,” and we eventually conclude that the absorption and/or emission TDMs are anisotropic.

The following treatment was inspired by discussions with Pete Sercel, Sasha Efros, John Lyons, Noam Bernstein, and Eran Rabani. Consider light of frequency ω_{ex} incident on a




Figure 6.28: Fluorescence anisotropy of CsPbBr₃ nanowires: (a) 10 nm isolated nanowires in cyclohexane excited at 465 nm, (b) 10 nm isolated nanowires in cyclohexane excited at 408 nm, (c) 3 nm bundles in toluene excited at 408 nm. The first letter of the legend indicates excitation polarization (Vertical or Horizontal), and the second letter indicates emission polarization.

nanowire whose width is much smaller than the wavelength of the light. Let the nanowire have dielectric constant $\epsilon_{nw}(\omega_{ex})$, and the surrounding medium have dielectric constant ϵ_m . As a consequence of Maxwells equations, the electric field for light polarized along the short axis of the nanowire is reduced by $2\epsilon_m/(\epsilon_m + \epsilon_{nw}(\omega_{ex}))$ [280, 283]. The same is true for light emitted at frequency ω_{em} , so the nanowire excitation and emission can both be cast in the following forms:

$$\begin{aligned} P(\omega_{ex}) &= 1 + \kappa(\omega_{ex})(\hat{c} \cdot \hat{e}_{ex}) \\ P(\omega_{em}) &= 1 + \kappa(\omega_{em})(\hat{c} \cdot \hat{e}_{em}), \end{aligned} \quad (6.11)$$

where \hat{c} is the unit vector along the nanowire axis and \hat{e} is the polarization of the excited or emitted light. κ has the form

$$\kappa(\omega) = \frac{(\epsilon_m + \epsilon_{nw}(\omega))^2}{4\epsilon_m^2} - 1. \quad (6.12)$$

In order to calculate the intensity for a given configuration, such as I_{HV} , we fix the directions of \hat{e}_{ex} and \hat{e}_{em} and take the product $P(\omega_{ex})P(\omega_{em})$. Define the excitation propagation direction as \hat{x} , the emission propagation direction as \hat{y} , and the vertical direction as \hat{z} . Letting the nanowire have orientation given by polar angle θ and azimuthal angle ϕ , the four intensities are:

$$\begin{aligned}
 I_{VV} &= (1 + \kappa_{ex} \cos^2 \theta)(1 + \kappa_{em} \cos^2 \theta) \\
 I_{VH} &= (1 + \kappa_{ex} \cos^2 \theta)(1 + \kappa_{em} \sin^2 \theta \cos^2 \phi) \\
 I_{HV} &= (1 + \kappa_{ex} \sin^2 \theta \sin^2 \phi)(1 + \kappa_{em} \cos^2 \theta) \\
 I_{HH} &= (1 + \kappa_{ex} \sin^2 \theta \sin^2 \phi)(1 + \kappa_{em} \sin^2 \theta \cos^2 \phi).
 \end{aligned} \tag{6.13}$$

For an isotropic distribution of nanowires we take the average over θ and ϕ and easily verify the following:

$$\begin{aligned}
 I_{VV}(\text{isotropic}) &= 1 + \kappa_{ex}/3 + \kappa_{em}/3 + \kappa_{ex}\kappa_{em}/5 \\
 I_{VH}(\text{isotropic}) &= 1 + \kappa_{ex}/3 + \kappa_{em}/3 + \kappa_{ex}\kappa_{em}/15 \\
 I_{HV}(\text{isotropic}) &= 1 + \kappa_{ex}/3 + \kappa_{em}/3 + \kappa_{ex}\kappa_{em}/15 \\
 I_{HH}(\text{isotropic}) &= 1 + \kappa_{ex}/3 + \kappa_{em}/3 + \kappa_{ex}\kappa_{em}/15.
 \end{aligned} \tag{6.14}$$

We see that $I_{VV} > I_{VH} = I_{HV} = I_{HH}$ if the nanowires are uniformly oriented and the only anisotropy arises due to classical electrodynamics. We observe $I_{HV} \neq I_{VH}$, so the situation is more complicated. Furthermore, we can consider the polarization anisotropy between VV and VH :

$$A_V \equiv \frac{I_{VV} - I_{VH}}{I_{VV} + I_{VH}} = \frac{\kappa_{ex}\kappa_{em}}{15 + 5(\kappa_{ex} + \kappa_{em}) + 2\kappa_{ex}\kappa_{em}}. \tag{6.15}$$

The optical-frequency dielectric constant of cyclohexane is 2.02 [284]. The optical-frequency dielectric constant of CsPbBr₃ is not well known, but computations indicate that it is about 5 [262]. Using these values we find $\kappa_{ex} = \kappa_{em} = 2.02$, and $A_v = 0.09$. The measured value based on the PL peak in Figure 6.28a is $A_V = 0.21$. Clearly, there are other effects at play.

We therefore consider that the absorption and emission TDMs could be anisotropic. For a transition at a given frequency, let the TDM strength be d_{\parallel} for light polarized along the long axis of the nanowire, and d_{\perp} for light polarized along the short axis of the nanowire. We then have:

$$\kappa(\omega) = \frac{d_{\parallel}^2(\epsilon_m + \epsilon_{nw}(\omega))^2}{4d_{\perp}^2\epsilon_m^2} - 1. \tag{6.16}$$

If the occupation of band-edge states is non-uniform that would also modify κ , but we do not consider this possibility because the TAM measurements in Section 6.4.3 indicate that the energy shift between states of different polarizations is only about 5 meV, which is small compared to the thermal energy at room temperature.

We also consider orientational anisotropy of the nanowires. Let the orientational distribution be isotropic in ϕ , but have some polar probability distribution function $f(\theta)$, where the normalization condition is

$$\int_0^{\pi/2} f(\theta) \sin \theta d\theta = 1. \quad (6.17)$$

Define the following moments of the distribution:

$$\begin{aligned} m_2 &\equiv \int_0^{\pi/2} f(\theta) \cos^2 \theta \sin \theta d\theta \\ m_4 &\equiv \int_0^{\pi/2} f(\theta) \cos^4 \theta \sin \theta d\theta. \end{aligned} \quad (6.18)$$

For a uniform distribution $m_2 = 1/3$ and $m_4 = 1/5$.

Combining the two effects described above, we have the following intensities:

$$\begin{aligned} I_{VV} &= 1 + \kappa_{ex}m_2 + \kappa_{em}m_2\kappa_{ex}\kappa_{em}m_4 \\ I_{VH} &= 1 + \kappa_{ex}m_2 + \kappa_{em}\frac{1-m_2}{2}\kappa_{ex}\kappa_{em}\frac{m_2-m_4}{2} \\ I_{HV} &= 1 + \kappa_{ex}\frac{1-m_2}{2} + \kappa_{em}m_2\kappa_{ex}\kappa_{em}\frac{m_2-m_4}{2} \\ I_{HH} &= 1 + \kappa_{ex}\frac{1-m_2}{2} + \kappa_{em}\frac{1-m_2}{2} + \kappa_{ex}\kappa_{em}\frac{1-2m_2+m_4}{8}. \end{aligned} \quad (6.19)$$

We cannot make much progress without knowing something about the orientational distribution, but it is enlightening to consider the anisotropy between I_{VH} and I_{HV} :

$$A_{cross} \equiv \frac{I_{HV} - I_{VH}}{I_{HV} + I_{VH}} = \frac{2(\kappa_{em} - \kappa_{ex})(m_2 - 1/3)}{4 + (\kappa_{em} + \kappa_{ex})(1 + m_2) + 2\kappa_{ex}\kappa_{em}(m_2 - m_4)}. \quad (6.20)$$

The anisotropy is proportional to both $(\kappa_{em} - \kappa_{ex})$ and also $(m_2 - 1/3)$. Therefore fluorescence anisotropy between I_{HV} and I_{VH} (which we observe in all panels of Figure 6.28) implies there must be orientational anisotropy **and** there must be a difference between κ for excitation and emission. We do not have enough information to determine the magnitude of this difference, but we can consider possible causes. It could be because ϵ_{nw} is a strong function of frequency, in which case the fact that excitation and emission are at different frequencies would produce different values of κ . However the amount of anisotropy is uniform in emission frequency, and as can be seen in Figure 6.28a-b, using a different excitation frequency does not significantly change the results. We therefore conclude that anisotropic TDMs are the cause: $d_{||} \neq d_{\perp}$.

Although we conclude that there is some anisotropy in the strength of the TDMs, uncertainty about the nanowire orientation in solution prevents us from saying how strong this anisotropy is and whether it occurs for absorption, emission, or both. Studies of CdSe nanoplates have found that band-edge emission is anisotropic while absorption high above the band edge is isotropic [285] (“high above” in this case was 1.1 eV). We might expect something similar to occur in CsPbBr₃, in which case the anisotropy would be purely in

emission, but the fact that the results vary slightly when going from 465 nm to 408 nm excitation (compare Figure 6.28 panels a and b) means we cannot be certain we are in the isotropic absorption regime. Note that 465 nm is only 280 meV above the band gap, while 408 nm is 650 meV above the band gap.

Recall from Section 6.4.3 that we found the band-edge polarization-dependent **shift** in energy was apparently size-independent, suggesting that an asymmetric unit cell was the cause. An asymmetric unit cell could also lead to anisotropic emission TDM strength, in which case we might expect to measure the same anisotropy in thinner nanowires, whereas if the anisotropy is due to a quantum confinement effect we might expect the anisotropy to increase. Looking at the fluorescence anisotropy for 3 nm nanowires (Figure 6.28c), we see that it is still present ($I_{HV} \neq I_{VH}$), but that it is *lessened*. Unfortunately, interpretation is confounded by the orientational anisotropy. Because 3 nm nanowires are smaller they may be better at orienting randomly in solution, and a decrease in orientational anisotropy will diminish the observed fluorescence anisotropy. These effects can be disentangled by measuring the polarized PL of individual bundles (see Section 6.6.2).

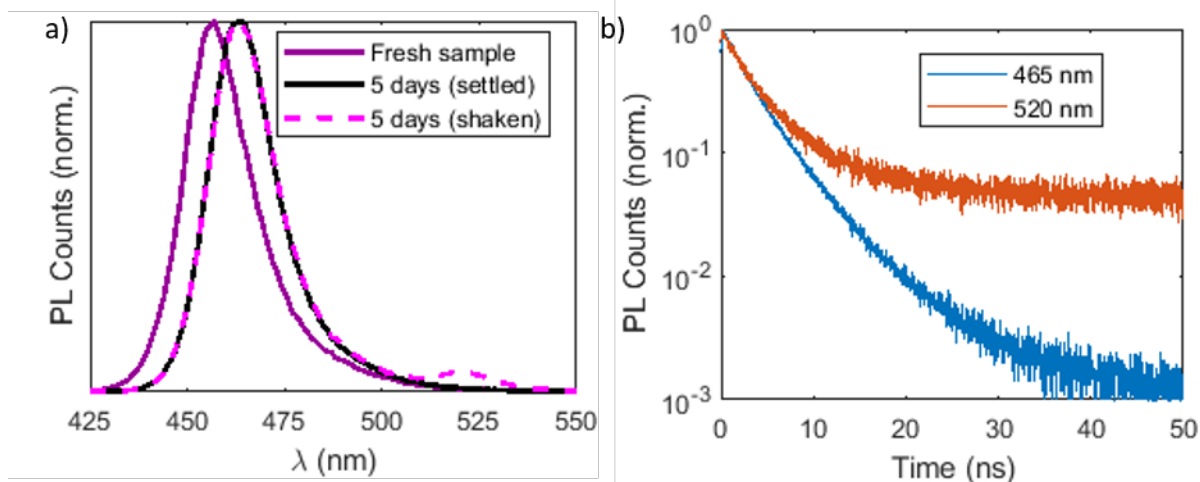


Figure 6.29: Defect formation and aggregation in 3 nm CsPbBr_3 nanowires. (a) PL spectrum of a fresh sample, that same sample 5 days later (red-shifted), and then after shaking (a bump appears at 520 nm). (b) TRPL on the 5 day-old sample at both 465 nm (blue) and 520 nm (orange), showing long-lived species that emit at 520 nm.

I will end on a word of caution, and note that one must be careful when working with the ultrathin nanowires as they rapidly aggregate and then disintegrate. The normalized PL spectrum (vertical excitation and emission) of 3 nm nanowires is shown in Figure 6.29a (purple curve), collected less than one hour after synthesis and surface treatment. After five days of sitting in the fluorimeter chamber, the spectrum had red-shifted by about 35 meV, indicative of aggregation [252]. Larger particles were also seen to collect at the bottom of the cuvette. Shaking the cuvette and re-measuring the PL produces the spectrum given by the

pink dashed line—in addition to the excitonic high-energy emission there is a peak at 520 nm, which corresponds to the intrinsic band edge. Measuring the TRPL at 520 nm (Figure 6.29b, orange curve), we see it is extremely long-lived. Data were collected out to 250 ns (not shown), which we use to estimate a lifetime of 700 ns. It seems that the ligands are detaching, leading to trap states on the bromine atoms [273]. But why is the fluorescence at the intrinsic band edge? It could be that the nanowires have so thoroughly aggregated that the electrons no longer experience a confined environment. A more intriguing possibility is that defects promote exciton dissociation into free carriers that occupy states associated with the intrinsic electronic structure, as was observed for Ruddlesden-Popper perovskites [251]. In that case the resulting carriers are protected, explaining their extremely long lifetime.

6.6 Future Directions

6.6.1 Applications

Our results indicate that CsPbBr₃ nanowires act as mostly independent entities. Even when bundled, they only couple weakly to each other—they retain their excitonic behavior due to quantum confinement, and rapid transport occurs primarily along one dimension. There is much interest in combining nanostructures with different morphologies in order to make unique hybrid devices [286], and this ability to retain their distinct behavior and transport properties means that CsPbBr₃ nanowires are good candidates for hybrid devices. Another benefit of weak coupling, as we saw in Figure 6.23b, is that high exciton concentrations can be sustained without nonlinear relaxation occurring. This makes ultrathin nanowires of CsPbBr₃ candidates for concentrated solar cells.

6.6.2 Further Studies of Band Edge Exciton Fine Structure

In this chapter, I discussed several measurements of the band-edge electronic structure of CsPbBr₃ nanowires. Using local absorption and TAM we learn that the nanowires display anisotropic optical coupling—the band-edge transition is slightly lower for light polarized along the long nanowire axis than it is for light polarized along the short nanowire axis. Furthermore, polarized PL indicates that the strength of the TDM is polarization-dependent as well. An asymmetric unit cell could be the cause [263, 282], but the orientational anisotropy in solution and the broad transitions makes it difficult to be certain.

A good next step would be to conduct polarization-resolved PL measurements of single bundles at low temperature. Studying single bundles with a known orientation (as we did for TAM and local absorption) will allow emitted light of different polarizations to be measured independently [282], and making the measurements at low temperature produces sharper emission peaks [287] that are easier to interpret. These measurements should be supplemented with electronic structure calculations [229] for various crystal structures and nanowire growth directions. As asymmetry in the unit cell may play a central role, the

*CHAPTER 6. TAM (AND FRIENDS) INVESTIGATE QUANTUM CONFINEMENT
AND DIMENSIONALITY EFFECTS IN INORGANIC PEROVSKITE NANOWIRES 172*

crystal field term will need to be included [288, 289]. In this way we can hope to determine the origin of band-edge degeneracy breaking in CsPbBr₃ nanowires.

Chapter 7

Concluding Remarks

This dissertation has presented several applications of ultrafast microscopy and related techniques to better understand the behavior of semiconductors with micron-scale structure. Chapters 1 and 2 provided the necessary background, covering the electronic structure and dynamics of organic semiconductors, the applications of these materials to optoelectronic devices, and the basics of transient absorption microscopy (TAM). TAM combines the exquisite temporal resolution of ultrafast laser pulses with the spatial resolution of microscope objectives, allowing us to study microscopic samples or observe micron-scale heterogeneity within a single sample. Because micron-scale samples tend to display some structural anisotropy, TAM has the added benefit of allowing us to selectively couple to individual transition dipole moments by rotating the optical polarization. These capabilities were crucial to deriving the results presented in the subsequent chapters.

Chapters 3 and 4 focused on small-molecule organic semiconducting thin films, in which solution processing resulted in distinct micron-scale crystalline domains. By measuring the polarization-resolved TAM signal within individual domains, we derived new insight into the electronic dynamics. In one material, diF-*TES*-ADT, we observed heterogeneous behavior between domains. We fit the data to a model and discovered that there are two distinct types of dynamics. In another material, TIPS-Pn, we focused on the process of singlet fission, in which a singlet exciton splits into two initially-correlated triplet excitons. We resolved a disagreement related to the triplet pair formation timescale, quantified the triplet-triplet binding energy, and measured how the charge transfer character of the correlated triplets perturbs their electronic structure and hence absorption spectrum.

In Chapter 5, we studied the self-assembly of rubrene spherulites. Spherulites are kinetically-trapped polycrystalline structures that contain numerous small- and large-angle grain boundaries. There are many open questions about fundamental forces driving self-assembly, the differences in structure and electronic properties between different regions of a spherulite, and how these properties can be tuned to optimize performance in a device. We found that many large-angle grain boundaries are templated by the spherulite nucleus, but that only in spherulites do these interfaces carry significant strain. This observation was incorporated into a model in an attempt to explain why single crystalline platelets form at

lower annealing temperatures and spherulites form at higher temperatures. In this chapter we have done significant ground work to understand rubrene spherulite self-assembly and behavior, but many open questions remain. Future work will involve revising and extending this model to extract parameters relevant to rubrene, fabricating spherulites on TEM grids so that we can use transmission electron microscopy to measure local strain with nanoscale spatial resolution, and applying TAM to determine how local spherulite morphology impacts electronic dynamics.

Lastly, in Chapter 6 we studied aligned bundles of inorganic perovskite nanowires. Time-resolved interferometric scattering imaged exciton diffusion, and the diffusivity was found to be similar to that of the bulk material. TAM and other optical measurements revealed that the nanowire-optical coupling depends on the polarization of light with respect to the long nanowire axis. This anisotropy did not appear to depend strongly on the width of the nanowires, indicating that the effect is due not to quantum confinement, but instead to a crystalline asymmetry that develops in all nanowires that are only a few unit cells in width. Unfortunately, there were several ambiguities in our measurement. In order to unambiguously determine the optical anisotropy of nanowires, future experiments could measure the polarization-resolved fluorescence measurements of individual bundles. This would provide significant insight into the electronic structure of perovskite nanowires, and how it is affected by both crystal structure and geometry.

The self-assembly of disordered semiconductors is an important problem both for elucidating fundamental principles of statistical physics and for its relevance to next-generation optoelectronic devices. We have combined several state of the art microscopy techniques to better understand the structure and dynamics of single crystalline organic semiconductors, polycrystalline organic semiconductors, and inorganic nanostructures. This thesis provides an example of how structural and dynamic probes can be combined to address a variety of different problems in semiconducting materials, and hopefully serves as an inspiration for many more such investigations.

Bibliography

- [1] Charles P. Enz. *No Time to be Brief. A Scientific Biography of Wolfgang Pauli*. Oxford University Press, 2002, p. 157.
- [2] Peter Y. Yu and Manuel Cardona. *Fundamentals of Semiconductors*. 4th ed. Springer, 2010.
- [3] Charles Kittel. *Introduction to Solid State Physics*. 8th ed. John Wiley and Sons, Inc, 2005.
- [4] John S. Townsend. *A Modern Approach to Quantum Mechanics*. McGraw-Hill Inc., 1992.
- [5] Chetvorno. *Solid state electronic band structure*. Wikimedia Commons, distributed under Creative Commons License 1.0. 2017. (Visited on 08/11/2018).
- [6] J. R. Schrieffer. “What is a Quasi-Particle?” In: *Journal of Research of the National Bureau of Standards - A. Physics and Chemistry* 74A.4 (1970), pp. 537–541.
- [7] Richard D. Mattuck. *A Guide to Feynman Diagrams in the Many-Body Problem*. 2nd ed. McGraw-Hill, 1976.
- [8] Marvin L. Cohen and Steven G. Louie. *Fundamentals of Condensed Matter Physics*. Cambridge University Press, 2016.
- [9] Sigma-Aldrich. URL: <https://www.sigmaaldrich.com/> (visited on 08/19/2018).
- [10] John E. Anthony. “The Larger Acenes: Versatile Organic Semiconductors”. In: *Angewandte Chemie - International Edition* 47 (2008), pp. 452–483. DOI: 10.1002/anie.200604045.
- [11] Anna M. Hiszpanski and Yueh-Lin Loo. “Directing the film structure of organic semiconductors via post-deposition processing for transistor and solar cell applications”. In: *Energy Environ. Sci.* 7.2 (2014), pp. 592–608. DOI: 10.1039/C3EE42615G.
- [12] J. E. Northrup R. A. Street and A. Salleo. “Organic single-crystal field-effect transistors”. In: *Materials Today* 71 (16 2005), p. 165202.
- [13] Colin Reese and Zhenan Bao. “Transport in polycrystalline polymer thin-film transistors”. In: *Physical Review B* 10.3 (2007), pp. 20–27.

- [14] Jonathan Rivnay et al. “Quantitative Determination of Organic Semiconductor Microstructure from the Molecular to Device Scale”. In: *Chemical Reviews* 112 (2012), pp. 5488–5519.
- [15] F. London. “The general theory of molecular forces”. In: *Transactions of the Faraday Society* 33 (1937).
- [16] Sean M. Ryno, Chad Risko, and Jean Luc Brédas. “Impact of molecular packing on electronic polarization in organic crystals: The case of pentacene vs TIPS-pentacene”. In: *Journal of the American Chemical Society* 136.17 (2014), pp. 6421–6427. DOI: 10.1021/ja501725s.
- [17] Steffen Illig et al. “Reducing dynamic disorder in small-molecule organic semiconductors by suppressing large-amplitude thermal motions”. In: *Nature Communications* 7 (2016), p. 10736. ISSN: 2041-1723. DOI: 10.1038/ncomms10736.
- [18] Alexander S. Eggeman et al. “Measurement of molecular motion in organic semiconductors by thermal diffuse electron scattering”. In: *Nature Materials* 12.11 (2013), pp. 1045–1049. DOI: 10.1038/nmat3710. (Visited on 07/02/2015).
- [19] Samuel B. Penwell. “Spatially Resolving Dynamics and Nanoscale Migration of Excitons in Organic Semiconductors Using Transient Absorption Imaging and STED Microscopy”. PhD thesis. University of California, Berkeley, 2016.
- [20] David P. McMahon and Alessandro Troisi. “Organic Semiconductors: Impact of Disorder at Different Timescales”. In: *ChemPhysChem* 11 (2010), pp. 2067–2074.
- [21] Yuanyuan Hu et al. “Direct Observation of the Dipole-Induced Energetic Disorder in Rubrene Single-Crystal Transistors by Scanning Kelvin Probe Microscopy”. In: *Journal of Physical Chemistry Letters* 9 (2018), pp. 2869–2873.
- [22] Roald Hoffmann. “How Chemistry and Physics Meet in the Solid State”. In: *Angewandte Chemie International Edition in English* 26 (1987), pp. 846–878.
- [23] Mingzhen Liu, Michael B. Johnston, and Henry J. Snaith. “Efficient planar heterojunction perovskite solar cells by vapour deposition”. In: *Nature* 501 (2013), p. 306.
- [24] Matthew T. Lloyd, John E. Anthony, and George G. Malliaras. “Photovoltaics from soluble small molecules”. In: *Materials Today* 10 (2007), pp. 34–41.
- [25] William Shockley and Hans J. Queisser. “Detailed Balance Limit of Efficiency of p-n Junction Solar Cells”. In: *Journal of Applied Physics* 32 (1961), p. 510.
- [26] Department of Energy. *Crystalline Silicon Photovoltaics Research*. URL: <https://www.energy.gov/eere/solar/crystalline-silicon-photovoltaics-research> (visited on 08/13/2018).
- [27] Thomas N. D. Tibbits et al. “New Efficiency Frontiers with Wafer-Bonded Multi-Junction Solar Cells”. In: *Proceedings of the 29th European PV Solar Energy Conference and Exhibition* (2014).

- [28] Xiaozhou Che et al. “High fabrication yield organic tandem photovoltaics combining vacuum- and solution-processed subcells with 15% efficiency”. In: *Nature Energy* 3 (2018), pp. 422–427.
- [29] Woon Seok Yang et al. “Iodide management in formamidinium-lead-halide-based perovskite layers for efficient solar cells”. In: *Science* 356 (6345 2017), pp. 1376–1379.
- [30] Rajib Mondal et al. “Molecular design for improved photovoltaic efficiency: band gap and absorption coefficient engineering”. In: *Journal of Materials Chemistry* 19 (2009), pp. 7195–7197.
- [31] Benjamin Lawrence Cotts. “Time-resolved optical spectroscopy of organic electronics as a function of local environment”. PhD thesis. University of California, Berkeley, 2016.
- [32] Robert A. Street. “Thin-Film Transistors”. In: *Advanced Materials* 21 (2009), pp. 2007–2022.
- [33] Sybille Allard et al. “Organic Semiconductors for Solution-Processable Field-Effect Transistors (OFETs)”. In: *Angewandte Chemie International Edition* 47 (2008), pp. 4070–4098.
- [34] Yongbo Yuan et al. “Ultra-high mobility transparent organic thin film transistors grown by an off-centre spin-coating method”. In: *Nature Communications* (2014). DOI: 10.1038/ncomms4005.
- [35] Bogyu Lim et al. “High Performance Solution Processed Organic Field Effect Transistors with Novel Diketopyrrolopyrrole-Containing Small Molecules”. In: *Scientific Reports* 7 (2017).
- [36] *Si - Silicon Electrical properties*. URL: <http://www.ioffe.ru/SVA/NSM/Semicond/Si/electric.html> (visited on 10/30/2018).
- [37] Ana Claudia Arias et al. “Materials and Applications for Large Area Electronics: Solution-Based Approaches”. In: *Chemical Reviews* 110 (2010), pp. 3–24.
- [38] I. D. W. Samuel and G. A. Turnbull. “Organic Semiconductor Lasers”. In: *Chemical Reviews* 107 (2007), pp. 1272–1295.
- [39] Jiachen Sun et al. “Organic/Inorganic Metal Halide Perovskite Optoelectronic Devices beyond Solar Cells”. In: *Advanced Science* (2018). DOI: 10.1002/adv.201700780.
- [40] Samuel W. Eaton et al. “Lasing in robust cesium lead halide perovskite nanowires”. In: *Proceedings of the National Academy of Sciences* 113 (2016), pp. 1993–1998.
- [41] Peter C. Sercel and Kerry J. Vahala. “Polarization dependence of optical absorption and emission in quantum wires”. In: *Physical Review B* 44 (1991).
- [42] Kidong Park et al. “Light-Matter Interactions in Cesium Lead Halide Perovskite Nanowire Lasers”. In: *The Journal of Physical Chemistry Letters* 7 (2016), pp. 3703–3710.

- [43] Sergii Yakunin et al. “Low-threshold amplified spontaneous emission and lasing from colloidal nanocrystals of caesium lead halide perovskites”. In: *Nature Communications* (2015). DOI: 10.1038/ncomms9056.
- [44] Brantly A. West et al. “Ultrafast dynamics of Frenkel Excitons in Tetracene and Rubrene Single Crystals”. In: *Journal of Chemical Physics* 140.11 (2014), pp. 10580–10591.
- [45] Andrew D. Platt et al. “Optical, Fluorescent, and (Photo)conductive Properties of High-Performance Functionalized Pentacene and Anthradithiophene Derivatives”. In: *The Journal of Physical Chemistry C* 113.31 (Aug. 2009), pp. 14006–14014.
- [46] Aleksandr Ryasnyanskiy and Ivan Biaggio. “Triplet exciton dynamics in rubrene single crystals”. In: *Physical Review B* 84 (2011), p. 193203.
- [47] Na Sai et al. “Optical spectra and exchange-correlation effects in molecular crystals”. In: *Physical Review B* 77 (2008), p. 161306.
- [48] M. Muccini et al. “Optical properties and the photoluminescence quantum yield of organic molecular materials”. In: *Journal of Optics A* 2 (2000), pp. 577–583.
- [49] Hiroki Uoyama et al. “Highly efficient organic light-emitting diodes from delayed fluorescence”. In: *Nature* 492 (2012), p. 234.
- [50] Brendan D. Folie et al. “Long-Lived Correlated Triplet Pairs in a π -Stacked Crystalline Pentacene Derivative”. In: *Journal of the American Chemical Society* 140 (2018), pp. 2326–2335.
- [51] Sam L. Bayliss et al. “Geminate and nongeminate recombination of triplet excitons formed by singlet fission”. In: *Physical Review Letters* 112.23 (2014), pp. 1–5. DOI: 10.1103/PhysRevLett.112.238701.
- [52] Kuo-Chun Tang, Kuan Lin Liu, and I-Chia Chen. “Rapid intersystem crossing in highly phosphorescent iridium complexes”. In: *Chemical Physics Letters* 386 (2004), pp. 437–441.
- [53] J. P. Zheng and H. S. Kwok. “Exciton and biexciton recombination in semiconductor nanocrystals”. In: *Applied Physics Letters* 65 (1994), pp. 1151–1153.
- [54] Naoki Yarita et al. “Dynamics of Charged Excitons and Biexcitons in CsPbBr₃ Perovskite Nanocrystals Revealed by Femtosecond Transient-Absorption and Single-Dot Luminescence Spectroscopy”. In: *The Journal of Physical Chemistry Letters* 8 (2017), pp. 1413–1418.
- [55] Emmanouil Kioupakis et al. “First-principles calculations of indirect Auger recombination in nitride semiconductors”. In: *Physical Review B* 92 (2015), p. 035207.
- [56] D. L. Dexter. “A Theory of Sensitized Luminescence in Solids”. In: *The Journal of Chemical Physics* 21 (1953), p. 836.
- [57] Andrei Tokmakoff. *Time-Dependent Quantum Mechanics and Spectroscopy*. 2014. URL: <https://tdqms.uchicago.edu/> (visited on 08/28/2018).

- [58] Rudi Berera, Riend van Grondelle, and John T. M. Kennis. “Ultrafast transient absorption spectroscopy: principles and applications to photosynthetic systems”. In: *Photosynthesis Research* 101 (2009), pp. 105–118.
- [59] Ye Yang et al. “Observation of a hot-phonon bottleneck in lead-iodide perovskites”. In: *Nature Photonics* 10 (2015), pp. 53–59.
- [60] Mark W B Wilson et al. “Ultrafast Dynamics of Exciton Fission in Polycrystalline Pentacene”. In: *Journal of the American Chemical Society* 133 (2011), pp. 11830–11833.
- [61] R. Joseph Kline et al. “Controlling the Microstructure of Solution-Processable Small Molecules in Thin-Film Transistors through Substrate Chemistry”. In: *Chemistry of Materials* 23.5 (2011), pp. 1194–1203. DOI: 10.1021/cm102834m.
- [62] Simon Gelinas et al. “Ultrafast Long-Range Charge Separation in Organic Semiconductor Photovoltaic Diodes”. In: *Science* 343 (2014), pp. 512–517.
- [63] Christopher Grieco et al. “Direct Observation of Correlated Triplet Pair Dynamics during Singlet Fission Using Ultrafast Mid-IR Spectroscopy”. In: *The Journal of Physical Chemistry C* 122 (2018), pp. 2012–2022.
- [64] S. A. Kovalenko et al. “Femtosecond spectroscopy of condensed phases with chirped supercontinuum probing”. In: *Physical Review A* 59 (1999), pp. 2369–2384.
- [65] Joseph S. Manser and Prashant V. Kamat. “Band filling with free charge carriers in organometal halide perovskites”. In: *Nature Photonics* 8 (2014), pp. 737–743.
- [66] Jeffrey A. Cina et al. “Ultrafast transient absorption revisited: Phase-flips, spectral fingers, and other dynamical features”. In: *The Journal of Chemical Physics* 144 (2016), p. 175102.
- [67] Joris J. Snellenburg et al. “Glotaran: A Java-Based Graphical User Interface for the R Package TIMP”. In: *Journal of Statistical Software* (2012).
- [68] Gregory V. Hartland. “Ultrafast studies of single semiconductor and metal nanostructures through transient absorption microscopy”. In: *Chemical Science* (2010).
- [69] Erik M Grumstrup et al. “Reversible Strain-Induced Electron-Hole Recombination in Silicon Nanowires Observed with Femtosecond Pump-Probe Microscopy”. In: *Nano Letters* 14 (2014), pp. 6287–6296.
- [70] Hongyan Shi et al. “Exciton dynamics in suspended monolayer and few-layer MoS₂ 2D crystals”. In: *ACS Nano* 7.2 (2013), pp. 1072–1080. DOI: 10.1021/nn303973r.
- [71] Cathy Y. Wong et al. “Revealing exciton dynamics in a small-molecule organic semiconducting film with subdomain transient absorption microscopy”. In: *Journal of Physical Chemistry C* 117.42 (2013), pp. 22111–22122. DOI: 10.1021/jp407645k.
- [72] Cathy Y. Wong et al. “Exciton dynamics reveal aggregates with intermolecular order at hidden interfaces in solution-cast organic semiconducting films”. In: *Nature Communications* 6 (2015), p. 5946. DOI: 10.1038/ncomms6946.

- [73] Cathy Y Wong et al. “Discerning Variable Extents of Interdomain Orientational and Structural Heterogeneity in Solution-Cast Polycrystalline Organic Semiconducting Thin Films”. In: *The Journal of Physical Chemistry Letters* (2015), pp. 3155–3162. DOI: 10.1021/acs.jpcllett.5b01416.
- [74] Zhi Guo et al. “Long-range hot-carrier transport in hybrid perovskites visualized by ultrafast microscopy”. In: *Science* 356 (2017), pp. 59–62.
- [75] Christoph Schnedermann et al. “Sub-10 fs Time-Resolved Vibronic Optical Microscopy”. In: *The Journal of Physical Chemistry Letters* (2016).
- [76] Dar’ya Davydova et al. “Transient absorption microscopy: advances in chemical imaging of photoinduced dynamics”. In: *Laser and Photonics Reviews* (2016), pp. 62–81.
- [77] Claudia Ruppert et al. “The Role of Electronic and Phononic Excitation in the Optical Response of Monolayer WS₂ after Ultrafast Excitation”. In: *Nano Letters* (2017), pp. 644–651.
- [78] Giulia Grancini et al. “Transient Absorption Imaging of P3HT : PCBM Photovoltaic Blend”. In: *The Journal of Physical Chemistry Letters* 2.9 (2011), pp. 1099–1105. DOI: 10.1021/jz200389b.
- [79] Emma E. M. Cating et al. “Imaging Spatial Variations in the Dissipation and Transport of Thermal Energy within Individual Silicon Nanowires Using Ultrafast Microscopy”. In: *Nano Letters* 16 (2015), pp. 434–439.
- [80] Tong Zhu et al. “Two Birds with One Stone: Tailoring Singlet Fission for Both Triplet Yield and Exciton Diffusion Length”. In: *Advanced Materials* 28 (2016), pp. 7539–7547. DOI: 10.1002/ADMA.201600968.
- [81] Yan Wan et al. “Cooperative singlet and triplet exciton transport in tetracene crystals visualized by ultrafast microscopy”. In: *Nature Chemistry* 7 (2015), pp. 785–792. DOI: 10.1038/nchem.2348.
- [82] Yan Wan et al. “Transport of Spin-Entangled Triplet Excitons Generated by Singlet Fission”. In: *Journal of Physical Chemistry Letters* 9 (2018), pp. 6731–6738.
- [83] Sankar Subramanian et al. “Chromophore Fluorination Enhances Crystallization and Stability of Soluble Anthradithiophene Semiconductors”. In: *Journal of the American Chemical Society* 130.9 (2008), pp. 2706–2707. DOI: 10.1021/ja073235k.
- [84] Oana D. Jurchescu et al. “Organic Single-Crystal Field-Effect Transistors of a Soluble Anthradithiophene”. In: *Chemistry of Materials* 20.21 (2008), pp. 6733–6737. DOI: 10.1021/cm8021165.
- [85] Stephanie S. Lee et al. “Controlling Nucleation and Crystallization in Solution-Processed Organic Semiconductors for Thin-Film Transistors”. In: *Advanced Materials* 21.35 (2009), pp. 3605–3609. DOI: 10.1002/adma.200900705.

- [86] Oana D. Jurchescu et al. “Effects of polymorphism on charge transport in organic semiconductors”. In: *Physical Review B* 80.8 (2009), p. 085201. DOI: 10.1103/PhysRevB.80.085201.
- [87] Ruipeng Li et al. “Direct Structural Mapping of Organic Field-Effect Transistors Reveals Bottlenecks to Carrier Transport”. In: *Advanced Materials* 24.41 (2012), pp. 5553–5558. DOI: 10.1002/adma.201201856.
- [88] Chaw Keong Yong et al. “The entangled triplet pair state in acene and heteroacene materials”. In: *Nature Communications* 8 (2017), p. 15953. DOI: 10.1038/ncomms15953.
- [89] Liyang Yu et al. “Programmable and coherent crystallization of semiconductors”. In: *Science Advances* 3.3 (2017), pp. 1–10. DOI: 10.1126/sciadv.1602462.
- [90] David T. James et al. “Controlling microstructure of pentacene derivatives by solution processing: Impact of structural anisotropy on optoelectronic properties”. In: *ACS Nano* 7.9 (2013), pp. 7983–7991. DOI: 10.1021/nn403073d.
- [91] Hannah L. Stern et al. “Vibronically coherent ultrafast triplet-pair formation and subsequent thermally activated dissociation control efficient endothermic singlet fission”. In: *Nature Chemistry* 9 (2017).
- [92] Stephanie S. Lee et al. “Orientation-Independent Charge Transport in Single Spherulites from Solution-Processed Organic Semiconductors”. In: *Journal of the American Chemical Society* 134.12 (2012), pp. 5436–5439. DOI: 10.1021/ja2116316.
- [93] Gregory D. Scholes. “Correlated Pair States Formed by Singlet Fission and Exciton-Exciton Annihilation”. In: *Journal of Physical Chemistry A* 119.51 (2015), pp. 12699–12705. DOI: 10.1021/acs.jpca.5b09725.
- [94] Millicent B Smith and Josef Michl. “Recent advances in singlet fission.” In: *Annual review of physical chemistry* 64 (2013), pp. 361–86. DOI: 10.1146/annurev-physchem-040412-110130.
- [95] Geoffrey B Piland et al. “Singlet Fission : From Coherences to Kinetics”. In: *Journal of Physical Chemistry Letters* 5 (2014), pp. 2312–2319.
- [96] Bruno Ehrler et al. “Singlet Exciton Fission-Sensitized Infrared Quantum Dot Solar Cells - Nano Letters (ACS Publications)”. In: *Nano Letters* 12 (2012), pp. 1053–1057. DOI: Doi10.1021/N1204297u.
- [97] Luis. M. Pazos-Outon et al. “A Silicon-Singlet Fission Parallel Tandem Solar Cell Exceeding 100% External Quantum Efficiency”. In: *ArXiv* 1512.07466 (2016).
- [98] Shane R Yost et al. “A transferable model for singlet-fission kinetics.” In: *Nature chemistry* 6.6 (2014), pp. 492–7. DOI: 10.1038/nchem.1945.
- [99] Timothy C. Berkelbach, Mark S. Hybertsen, and David R. Reichman. “Microscopic theory of singlet exciton fission. III. Crystalline pentacene”. In: *Journal of Chemical Physics* 141.7 (2014).

- [100] Paul E. Teichen and Joel D. Eaves. “Collective aspects of singlet fission in molecular crystals”. In: *Journal of Chemical Physics* 143.4 (2015). DOI: 10.1063/1.4922644.
- [101] Xintian Feng, Anatoliy V. Luzanov, and Anna I. Krylov. “Fission of entangled Spins: An electronic structure perspective”. In: *Journal of Physical Chemistry Letters* 4.22 (2013), pp. 3845–3852. DOI: 10.1021/jz402122m.
- [102] Paul M. Zimmerman et al. “Mechanism for singlet fission in pentacene and tetracene: From single exciton to two triplets”. In: *Journal of the American Chemical Society* 133.49 (2011), pp. 19944–19952. DOI: 10.1021/ja208431r.
- [103] Karan Aryanpour, Alok Shukla, and Sumit Mazumdar. “Theory of singlet fission in polyenes, acene crystals, and covalently linked acene dimers”. In: *Journal of Physical Chemistry C* 119.13 (2015), pp. 6966–6979. DOI: 10.1021/jp5124019.
- [104] S. Refaely-Abramson et al. “Origins of singlet fission in solid pentacene from an ab initio Green’s-function approach”. In: *Physical Review Letters* 119 (2017), p. 267401.
- [105] Erik Busby et al. “A design strategy for intramolecular singlet fission mediated by charge-transfer states in donor-acceptor organic materials.” In: *Nature materials* 14.4 (2015), pp. 426–33. DOI: 10.1038/nmat4175.
- [106] Eric A. Margulies et al. “Enabling singlet fission by controlling intramolecular charge transfer in π -stacked covalent terrylenediimide dimers”. In: *Nature Chemistry* 8 (2016), pp. 1–6. DOI: 10.1038/nchem.2589.
- [107] Nicholas R Monahan et al. “Dynamics of the triplet pair state reveals the likely coexistence of coherent and incoherent singlet fission in crystalline hexacene”. In: *Nature chemistry* 8 (2016), pp. 1–18. DOI: 10.1038/nchem.2665.
- [108] Artem A. Bakulin et al. “Real-time observation of multiexcitonic states in ultrafast singlet fission using coherent 2D electronic spectroscopy”. In: *Nature Chemistry* 8.1 (2016), pp. 16–23. DOI: 10.1038/nchem.2371.
- [109] Roel Tempelaar and David R Reichman. “Vibronic exciton theory of singlet fission. I. Linear absorption and the anatomy of the correlated triplet pair state”. In: *arXiv* 1703.01173 (2017).
- [110] Dylan H. Arias et al. “Polymorphism influences singlet fission rates in tetracene thin films”. In: *Chem. Sci.* 7.2 (2016), pp. 1185–1191. DOI: 10.1039/C5SC03535J.
- [111] Xintian Feng, Anatoly B. Kolomeisky, and Anna I. Krylov. “Dissecting the effect of morphology on the rates of singlet fission: Insights from theory”. In: *Journal of Physical Chemistry C* 118.34 (2014), pp. 19608–19617. DOI: 10.1021/jp505942k.
- [112] Robert J. Dillon, Geoffrey B. Piland, and Christopher J. Bardeen. “Different rates of singlet fission in monoclinic versus orthorhombic crystal forms of diphenylhexatriene”. In: *Journal of the American Chemical Society* 135.46 (2013), pp. 17278–17281. DOI: 10.1021/ja409266s.

- [113] Joseph L Ryerson et al. “Two Thin Film Polymorphs of the Singlet Fission Compound 1,3-Diphenylisobenzofuran”. In: *Journal of Physical Chemistry C* 118 (2014), pp. 12121–12132.
- [114] Kalishankar Bhattacharyya and Ayan Datta. “Polymorphism Controlled Singlet Fission in TIPS-Anthracene: Role of Stacking Orientation”. In: *The Journal of Physical Chemistry C* 121 (2017), pp. 1412–1420. DOI: 10.1021/acs.jpcc.6b10075.
- [115] S T Roberts et al. “Efficient singlet fission found in a disordered acene film”. In: *Journal of the American Chemical Society* 134 (2012), pp. 6388–6400.
- [116] Jessica Wade et al. “Charge mobility anisotropy of functionalized pentacenes in organic field effect transistors fabricated by solution processing”. In: *J. Mater. Chem. C* 2.47 (2014), pp. 10110–10115. DOI: 10.1039/C4TC01353K.
- [117] Leonidas C. Palilis et al. “Organic photovoltaic cells with high open circuit voltages based on pentacene derivatives”. In: *Organic Electronics: physics, materials, applications* 9.5 (2008), pp. 747–752. DOI: 10.1016/j.orgel.2008.05.015.
- [118] John E. Anthony et al. “Synthesis, properties, and device applications of functionalized acenes”. In: *Proceedings of SPIE* 5217 (2003), pp. 124–132. DOI: 10.1117/12.506101.
- [119] Juan Arago and Alessandro Troisi. “Regimes of Exciton Transport in Molecular Crystals in the Presence of Dynamic Disorder”. In: *Advanced Functional Materials* 26 (2016), pp. 2316–2325. DOI: 10.1002/adfm.201503888.
- [120] Nicolas Renaud, Paul A. Sherratt, and Mark A. Ratner. “Mapping the relation between stacking geometries and singlet fission yield in a class of organic crystals”. In: *Journal of Physical Chemistry Letters* 4.7 (2013), pp. 1065–1069. DOI: 10.1021/jz400176m.
- [121] Hang Zang et al. “Effects of Charge Transfer State and Exciton Migration on Singlet Fission Dynamics in Organic Aggregates”. In: *Journal of Physical Chemistry C* 120.25 (2016), pp. 13351–13359. DOI: 10.1021/acs.jpcc.6b02943.
- [122] Catherine M Mauck et al. “Singlet Fission via an Excimer-Like Intermediate in 3,6-Bis(thiophen-2-yl)diketopyrrolopyrrole Derivatives”. In: *Journal of the American Chemical Society* 138 (2016), pp. 11749–11761. DOI: 10.1021/jacs.6b05627.
- [123] Wai-Lun Chan et al. “Observing the Multiexciton State in Singlet Fission and Ensuing Ultrafast Multielectron Transfer”. In: *Science* 334.December (2011), pp. 1541–1545. DOI: 10.1126/science.1213986.
- [124] Rui Wang et al. “Magnetic dipolar interaction between correlated triplets created by singlet fission in tetracene crystals.” In: *Nature communications* 6.May (2015), p. 8602. DOI: 10.1038/ncomms9602.

- [125] Wai-Lun Chan, Manuel Ligges, and X-Y. Zhu. “The energy barrier in singlet fission can be overcome through coherent coupling and entropic gain”. In: *Nature Chemistry* 4.10 (2012), pp. 840–845. DOI: 10.1038/nchem.1436.
- [126] Leah R. Weiss et al. “Strongly exchange-coupled triplet pairs in an organic semiconductor”. In: *Nature Physics* 13 (2016), pp. 176–181. DOI: 10.1038/nphys3908.
- [127] Murad J. Y. Tayebjee et al. “Quintet multiexciton dynamics in singlet fission”. In: *Nature Physics* 1.October (2016). DOI: 10.1038/nphys3909.
- [128] Mark W B Wilson et al. “Temperature-independent singlet exciton fission in tetracene”. In: *Journal of the American Chemical Society* 135.44 (2013), pp. 16680–16688. DOI: 10.1021/ja408854u.
- [129] Hannah L. Stern et al. “Identification of a triplet pair intermediate in singlet exciton fission in solution”. In: *Proceedings of the National Academy of Sciences* 112.25 (2015), pp. 7656–7661. DOI: 10.1073/pnas.1503471112.
- [130] Ryan D Pensack et al. “Observation of Two Triplet Pair Intermediates in Singlet Exciton Fission”. In: *The Journal of Physical Chemistry Letters* 7 (2016), pp. 2370–2375. DOI: 10.1021/acs.jpcllett.6b00947.
- [131] Ashok Maliakal et al. “Photochemical Stability of Pentacene and a Substituted Pentacene in Solution and in Thin Films”. In: *Chemistry of Materials* 16.24 (2004), pp. 4980–4986. DOI: 10.1021/cm049060k.
- [132] Alessandro Troisi, Giorgio Orlandi, and John E. Anthony. “Electronic interactions and thermal disorder in molecular crystals containing cofacial pentacene units”. In: *Chemistry of Materials* 17.17 (2005), pp. 5024–5031. DOI: 10.1021/cm051150h.
- [133] Sahar Sharifzadeh et al. “Relating the Physical Structure and Optoelectronic Function of Crystalline TIPS-Pentacene”. In: *Advanced Functional Materials* 25.13 (2015), pp. 2038–2046. DOI: 10.1002/adfm.201403005.
- [134] Jihua Chena et al. “Controlled solution deposition and systematic study of charge-transport anisotropy in single crystal and single-crystal textured TIPS pentacene thin film”. In: *Organic Electronics* 10 (2009), pp. 696–703.
- [135] Le Yang et al. “Solution-processable singlet fission photovoltaic devices”. In: *Nano Letters* 15.1 (2015), pp. 354–358. DOI: 10.1021/nl503650a.
- [136] Brian J Walker et al. “Singlet exciton fission in solution.” In: *Nature chemistry* 5.12 (2013), pp. 1019–24. DOI: 10.1038/nchem.1801.
- [137] Andrew J. Musser et al. “Evidence for conical intersection dynamics mediating ultrafast singlet exciton fission”. In: *Nature Physics* 11 (2015), pp. 352–357. DOI: 10.1038/nphys3241.

- [138] Charusheela Ramanan et al. “Competition between Singlet Fission and Charge Separation in Solution-Processed Blend Films of 6,13-Bis(triisopropylsilylethynyl)-pentacene with Sterically-Encumbered Perylene-3,4:9,10-bis(dicarboximides)”. In: *Journal of the American Chemical Society* 134 (2012), pp. 386–397.
- [139] Ryan D Pensack et al. “Exciton Delocalization Drives Rapid Singlet Fission in Nanoparticles of Acene Derivatives.” In: *Journal of the American Chemical Society* 137.21 (2015), pp. 6790–6803. DOI: 10.1021/ja512668r.
- [140] Murad J Y Tayebjee et al. “Morphological Evolution and Singlet Fission in Aqueous Suspensions of TIPS-Pentacene Nanoparticles”. In: *Journal of Physical Chemistry C* 120.1 (2016), pp. 157–165. DOI: 10.1021/acs.jpcc.5b11353.
- [141] Ryan D Pensack et al. “Solution-processable , crystalline material for quantitative singlet fission”. In: *Materials Horizons* 4 (2017), pp. 915–923. DOI: 10.1039/C7MH00303J.
- [142] Stefan C B Mannsfeld, Ming Lee Tang, and Zhenan Bao. “Thin film structure of triisopropylsilylethynyl-functionalized pentacene and tetraceno[2,3-b]thiophene from grazing incidence x-ray diffraction”. In: *Advanced Materials* 23.1 (2011), pp. 127–131. DOI: 10.1002/adma.201003135.
- [143] Hiroyuki Tamura et al. “First-Principles Quantum Dynamics of Singlet Fission: Coherent versus Thermally Activated Mechanisms Governed by Molecular π Stacking”. In: *Physical Review Letters* 115.10 (2015), pp. 1–5. DOI: 10.1103/PhysRevLett.115.107401.
- [144] Julia Herz et al. “Acceleration of singlet fission in an aza-derivative of TIPS-pentacene”. In: *Journal of Physical Chemistry Letters* 5.14 (2014), pp. 2425–2430. DOI: 10.1021/jz501102r.
- [145] Thomas James McDonough et al. “Triplet Exciton Dissociation and Electron Extraction in Graphene-Templated Pentacene Observed with Ultrafast Spectroscopy”. In: *Phys. Chem. Chem. Phys.* 19 (2017), pp. 4809–4820. DOI: 10.1039/C6CP06454J.
- [146] Murad J Y Tayebjee, Raphaël G C R Clady, and Timothy W Schmidt. “The exciton dynamics in tetracene thin films”. In: *Physical Chemistry Chemical Physics* 15.35 (2013), p. 14797. DOI: 10.1039/c3cp52609g.
- [147] Jack Deslippe et al. “BerkeleyGW: A massively parallel computer package for the calculation of the quasiparticle and optical properties of materials and nanostructures”. In: *Computer Physics Communications* 183.6 (2012), pp. 1269–1289. DOI: 10.1016/j.cpc.2011.12.006.
- [148] Mark S. Hybertsen and Steven G. Louie. “Electron correlation in semiconductors and insulators: Band gaps and quasiparticle energies”. In: *Physical Review B* 34.8 (1986), pp. 5390–5413. DOI: 10.1103/PhysRevB.34.5390.
- [149] Justin C. Johnson, Arthur J. Nozik, and Josef Michl. “The Role of Chromophore Coupling in Singlet Fission”. In: *Acc. Chem. Res.* 46.6 (2013), pp. 1290–1299. DOI: 10.1021/cr1002613.

- [150] Brian Johnson, Mark J. Kendrick, and Oksana Ostroverkhova. “Charge carrier dynamics in organic semiconductors and their donor-acceptor composites: Numerical modeling of time-resolved photocurrent”. In: *Journal of Applied Physics* 114.9 (2013), p. 094508. DOI: 10.1063/1.4820259.
- [151] Kolja Kolata et al. “Molecular Packing Determines Singlet Exciton Fission in Organic Semiconductors”. In: *ACS Nano* 8.7 (2014), pp. 7377–7383.
- [152] Andre Rinn et al. “Molecular Packing Motifs Determine Charge- Transfer and Carrier Dynamics in Molecular Heterosystems : the Case of Pentacene - Perfluoropentacene”. In: *Arxiv* 1612.03478 (2016).
- [153] Sahar Sharifzadeh et al. “Low-energy charge-transfer excitons in organic solids from first-principles: The case of pentacene”. In: *Journal of Physical Chemistry Letters* 4.13 (2013), pp. 2197–2201. DOI: 10.1021/jz401069f.
- [154] Sam L. Bayliss et al. “Localization length scales of triplet excitons in singlet fission materials”. In: *Physical Review B - Condensed Matter and Materials Physics* 92.11 (2015), pp. 1–7. DOI: 10.1103/PhysRevB.92.115432.
- [155] Nadezhda V. Korovina et al. “Singlet Fission in a Covalently Linked Cofacial Alkynyl-tetracene Dimer”. In: *Journal of the American Chemical Society* 138.2 (2016), pp. 617–627. DOI: 10.1021/jacs.5b10550.
- [156] V. Podzorov, V. M. Pudalov, and M. E. Gershenson. “Field-effect transistors on rubrene single crystals with parylene gate insulator”. In: *Applied Physics Letters* 82.11 (2003), pp. 1739–1741. DOI: 10.1063/1.1560869.
- [157] Frederik C. Krebs. “Fabrication and processing of polymer solar cells A review of printing and coating techniques”. In: *Solar Energy Materials and Solar Cells* 93 (2009), pp. 394–412.
- [158] Y. Diao et al. “Morphology control strategies for solution- processed organic semiconductor thin film”. In: *Energy and Environmental Sciences* 7 (2014), pp. 2145–2159.
- [159] Tommie W. Kelley and C. Daniel Frisbie. “Gate Voltage Dependent Resistance of a Single Organic Semiconductor Grain Boundary”. In: *Journal of Physical Chemistry B* 105 (2001), pp. 4538–4540.
- [160] Xiantong Tang et al. “Intersystem Crossing and Triplet Fusion in Singlet-Fission-Dominated Rubrene-Based OLEDs Under High Bias Current”. In: *ACS Applied Materials & Interfaces* (2018). DOI: 10.1021/acsami.7b17695.
- [161] Lin Ma et al. “Singlet fission in rubrene single crystal: direct observation by femtosecond pump-probe spectroscopy”. In: *Physical Chemistry Chemical Physics* 14.23 (2012), p. 8307. DOI: 10.1039/c2cp40449d.

- [162] Ilana Breen et al. “Triplet separation drives singlet fission after femtosecond correlated triplet pair production in rubrene”. In: *Journal of the American Chemical Society* 139 (2017). DOI: 10.1021/jacs.7b02621.
- [163] J. Takeya et al. “Very high-mobility organic single-crystal transistors with in-crystal conduction channels”. In: *Applied Physics Letters* 90.10 (2007). DOI: 10.1063/1.2711393.
- [164] Tatsuo Hasegawa and Jun Takeya. “Organic field-effect transistors using single crystals”. In: *Science and Technology of Advanced Materials* 10 (2009).
- [165] Huanli Dong et al. “25th Anniversary Article: Key Points for High-Mobility Organic Field-Effect Transistors”. In: *Advanced Materials* 25 (2013).
- [166] Xiaopeng Wang et al. “Effect of crystal packing on the excitonic properties of rubrene polymorphs”. In: *CrystEngComm* 18.38 (2016), pp. 7353–7362. DOI: 10.1039/C6CE00873A.
- [167] Stefano Bergantin et al. “Pressure-induced conformational change in organic semiconductors: Triggering a reversible phase transition in rubrene”. In: *Journal of Physical Chemistry C* 118.25 (2014), pp. 13476–13483. DOI: 10.1021/jp503271h.
- [168] Liwei Huang et al. “Rubrene micro-crystals from solution routes: their crystallography, morphology and optical properties”. In: *Journal of Materials Chemistry* 20.1 (2010), p. 159. DOI: 10.1039/b914334c.
- [169] Pil Sung Jo et al. “Control of rubrene polymorphs via polymer binders: Applications in organic field-effect transistors”. In: *Chemistry of Materials* 27.11 (2015), pp. 3979–3987. DOI: 10.1021/acs.chemmater.5b00884.
- [170] V. Podzorov et al. “Hall Effect in the Accumulation Layers on the Surface of Organic Semiconductors”. In: *Physical Review Letters* 95 (2005).
- [171] László Gránásy et al. “Growth and form of spherulites”. In: *Physical Review E* 72.1 (2005). DOI: 10.1103/PhysRevE.72.011605.
- [172] Pavlo Fesenko et al. “Determination of crystal orientation in organic thin films using optical microscopy”. In: *Organic Electronics: physics, materials, applications* 37 (2016), pp. 100–107. DOI: 10.1016/j.orgel.2016.06.011.
- [173] Stephanie S. Lee et al. “Quantifying resistances across nanoscale low- and high-angle interspherulite boundaries in solution-processed organic semiconductor thin films”. In: *ACS Nano* 6.11 (2012), pp. 9879–9886. DOI: 10.1021/nn303446h.
- [174] Anna K Hailey et al. “Quantifying the Energy Barriers and Elucidating the Charge Transport Mechanisms across Interspherulite Boundaries in Solution-Processed Organic Semiconductor Thin Films”. In: *Advanced Functional Materials* 25 (2015).
- [175] G. Ryschenkow and G. Faivre. “Bulk crystallization of liquid selenium Primary nucleation, growth kinetics and modes of crystallization”. In: *Journal of Crystal Growth* 87.2-3 (1988), pp. 221–235. DOI: 10.1016/0022-0248(88)90169-8.

- [176] Yajun Su et al. “Polymer assisted solution-processing of rubrene spherulites via solvent vapor annealing”. In: *RSC Advances* 2.13 (2012), p. 5779. DOI: 10.1039/c2ra20417g.
- [177] Anna K Hailey. “Engineering Materials Properties and Process Technologies for Electronic and Energy Applications”. PhD thesis. 2017.
- [178] Akira Otomo et al. “Photochemical stability of encapsulated laser dyes in dendritic nanoboxes against singlet oxygen.” In: *Optics letters* 27.11 (2002), pp. 891–893. DOI: 10.1364/OL.27.000891.
- [179] C. Kloc et al. “Purity of rubrene single crystals”. In: *Applied Physics A: Materials Science and Processing* 95.1 (2009), pp. 219–224. DOI: 10.1007/s00339-008-5014-0.
- [180] Jack Ly et al. “Reversible and Irreversible Oxidation of Rubrene, and Implications for Device Stability”. In: *Journal of Materials Chemistry C* (2018). DOI: 10.1039/C7TC05775J.
- [181] Bregt Verreet et al. “Microcrystalline organic thin-film solar cells”. In: *Advanced Materials* 25.38 (2013), pp. 5504–5507. DOI: 10.1002/adma.201301643.
- [182] Michael A. Fusella et al. “Use of an Underlayer for Large Area Crystallization of Rubrene Thin Films”. In: *Chemistry of Materials* 29.16 (2017), pp. 6666–6673. DOI: 10.1021/acs.chemmater.7b01143.
- [183] Pavel Irkhin et al. “Absorption and photoluminescence spectroscopy of rubrene single crystals”. In: *Physical Review B - Condensed Matter and Materials Physics* 86.8 (2012), pp. 1–13. DOI: 10.1103/PhysRevB.86.085143.
- [184] Luisa Raimondo et al. “Oxidation of Crystalline Rubrene Films: Evidence of an Epitaxial Native Oxide Layer”. In: *Advanced Materials Interfaces* 1700670 (2017), p. 1700670. DOI: 10.1002/admi.201700670.
- [185] Thomas R. Fielitz and Russell J. Holmes. “Crystal Morphology and Growth in Annealed Rubrene Thin Films”. In: *Crystal Growth and Design* 16.8 (2016), pp. 4720–4726. DOI: 10.1021/acs.cgd.6b00783.
- [186] Alexander Hexemer et al. “A SAXS/WAXS/GIWAXS Beamline with Multilayer Monochromator”. In: *Journal of Physics: Conference Series* 247 (2010).
- [187] Zhang Jiang. “GIXSGUI: a MATLAB toolbox for grazing-incidence X-ray scattering data visualization and reduction, and indexing of buried three-dimensional periodic nanostructured films”. In: *Journal of Applied Crystallography* 48 (2015), pp. 917–926.
- [188] Oana D. Jurchescu, Auke Meetsma, and T. T M Palstra. “Low-temperature structure of rubrene single crystals grown by vapor transport”. In: *Electrochemical and Solid-State Letters* 9.5 (2006), pp. 330–334. DOI: 10.1107/S0108768106003053.
- [189] Mira El Helou, Olaf Medenbach, and Gregor Witte. “Rubrene microcrystals: A route to investigate surface morphology and bulk anisotropies of organic semiconductors”. In: *Crystal Growth and Design* 10.8 (2010), pp. 3496–3501. DOI: 10.1021/cg1003758.

- [190] Takeshi Matsukawa et al. “Solution growth of rubrene single crystals using various organic solvents”. In: *Japanese Journal of Applied Physics* 47.12 (2008), pp. 8950–8954. DOI: 10.1143/JJAP.47.8950.
- [191] Ku-Yen Lin et al. “Role of molecular conformations in rubrene polycrystalline films growth from vacuum deposition at various substrate temperatures”. In: *Scientific Reports* 7 (2017), p. 40824. DOI: 10.1038/srep40824.
- [192] Mark C. Weidman, Detlef-M. Smilgies, and William A. Tisdale. “Kinetics of the self-assembly of nanocrystal superlattices measured by real-time in situ X-ray scattering”. In: *Nature Materials* 15 (2016), pp. 775–781.
- [193] Jaco J. Geuchies et al. “In situ study of the formation mechanism of two-dimensional superlattices from PbSe nanocrystals”. In: *Nature Materials* 15 (2016), pp. 1248–1254.
- [194] Kang Wei Chou et al. “Late stage crystallization and healing during spin-coating enhance carrier transport in small-molecule organic semiconductors”. In: *Journal of Materials Chemistry C* 2 (2014), pp. 5681–5689.
- [195] Neil D. Treat et al. “In situ measurement of power conversion efficiency and molecular ordering during thermal annealing in P3HT:PCBM bulk heterojunction solar cells”. In: *Journal of Materials Chemistry* 21 (2011), pp. 15224–15231.
- [196] Stephan Proller et al. “Following the Morphology Formation In Situ in Printed Active Layers for Organic Solar Cells”. In: *Advanced Energy Materials* 6 (2016).
- [197] Martin Kunz et al. “A dedicated superbend x-ray microdiffraction beamline for materials, geo-, and environmental sciences at the advanced light source”. In: *Review of Scientific Instruments* 80 (2009).
- [198] Kiyoshi Miyata et al. “Coherent singlet fission activated by symmetry breaking”. In: *Nature Chemistry* 9.10 (2017), pp. 983–989. DOI: 10.1038/nchem.2784.
- [199] Eric A. Muller et al. “Infrared vibrational nano-crystallography and nano-imaging”. In: *Science Advances* 2 (2016).
- [200] Ouliana Panova et al. “Orientation mapping of semicrystalline polymers using scanning electron nanobeam diffraction”. In: *Micron* 88 (2016).
- [201] Wikipedia. *Rubrene*. URL: <https://en.wikipedia.org/wiki/Rubrene> (visited on 11/02/2018).
- [202] M. D. Ediger, Peter Harrowell, and Lian Yu. “Crystal growth kinetics exhibit a fragility-dependent decoupling from viscosity”. In: *Journal of Chemical Physics* 128.3 (2008). DOI: 10.1063/1.2815325.
- [203] Gang Sun, Jenny Xu, and Peter Harrowell. “The mechanism of the ultrafast crystal growth of pure metals from their melts”. In: *Nature Materials* (2018). DOI: 10.1038/s41563-018-0174-6.

- [204] K. L. Ngai, J. H. Magill, and D. J. Plazek. “Flow, diffusion and crystallization of supercooled liquids: Revisited”. In: *Journal of Chemical Physics* 112.4 (2000), pp. 1887–1892. DOI: 10.1063/1.480752.
- [205] Laura A. Deschenes and David A. Vanden Bout. “Comparison of ensemble and single molecule approaches to probing polymer relaxation dynamics near T_g ”. In: *Journal of Chemical Physics* 116 (2002), pp. 5850–5856.
- [206] L. S. Garcia-Colin, L. F. Del Castillo, and Patricia Goldstein. “Theoretical basis for the Vogel-Fulcher-Tammann equation”. In: *Physical Review B* 40.10 (1989), pp. 7040–7044. DOI: 10.1103/PhysRevB.40.7040.
- [207] F. R. Blackburn, Chia-Ying Wang, and M. D. Ediger. “Translational and Rotational Motion of Probes in Supercooled 1,3,5-Tris(naphthyl)benzene”. In: *Journal of Physical Chemistry* 100 (1996), pp. 18249–18257.
- [208] Juan P. Garrahan and David Chandler. “Coarse-grained microscopic model of glass formers”. In: *Proceedings of the National Academy of Sciences* 100 (2003), pp. 9710–9714.
- [209] M. D. Ediger. “Spatially Heterogeneous Dynamics in Supercooled Liquids”. In: *Annual Reviews of Physical Chemistry* 51 (2000), pp. 99–128.
- [210] Alexander Agapov. “Decoupling Phenomena in Dynamics of Soft Matter”. PhD thesis. 2011.
- [211] M. D. Ediger, C. A. Angell, and Sidney R. Nagel. “Supercooled Liquids and Glasses”. In: *Journal of Physical Chemistry* 100 (1996), pp. 13200–13212.
- [212] Se Woong Park et al. “Amorphous-to-crystalline phase transformation of thin film rubrene”. In: *Journal of Physical Chemistry B* 114.17 (2010), pp. 5661–5665. DOI: 10.1021/jp910459p.
- [213] Y. Luo et al. “Growth of Rubrene thin film, spherulites and nanowires on SiO₂”. In: *Physica Status Solidi (A) Applications and Materials Science* 204.6 (2007), pp. 1851–1855. DOI: 10.1002/pssa.200675340.
- [214] László Gránásy et al. “Phase-field modeling of polycrystalline solidification: From needle crystals to spherulites - A review”. In: *Metallurgical and Materials Transactions A: Physical Metallurgy and Materials Science* 45.4 (2014), pp. 1694–1719. DOI: 10.1007/s11661-013-1988-0.
- [215] L. Gránásy et al. “Phase field theory of crystal nucleation and polycrystalline growth: A review”. In: *Journal of Materials Research* 21.2 (2006), pp. 309–319. DOI: 10.1557/jmr.2006.0011.
- [216] Ryo Kobayashi, James A. Warren, and W. Craig Carter. “A continuum model of grain boundaries”. In: *Physica D* 140 (2000), pp. 141–150.

- [217] Keewook Paeng et al. “Fast Crystal Growth Induces Mobility and Tension in Supercooled o-Terphenyl”. In: *Journal of Physical Chemistry Letters* 3 (2012), pp. 2562–2567.
- [218] Thomas Kluyver et al. “Jupyter Notebooks – a publishing format for reproducible computational workflows”. In: *Positioning and Power in Academic Publishing: Players, Agents and Agendas*. Ed. by F. Loizides and B. Schmidt. IOS Press. 2016, pp. 87–90.
- [219] S. Behnel et al. “Cython: The Best of Both Worlds”. In: *Computing in Science Engineering* 13.2 (2011), pp. 31–39. DOI: 10.1109/MCSE.2010.118.
- [220] Luis A. Zepeda-Ruiz and George H. Gilmer. “Monte Carlo Simulations of Crystal Growth”. In: *Handbook of Crystal Growth*. Ed. by Tatau Nishinaga. 2nd ed. Elsevier, 2015. Chap. 10, pp. 445–475.
- [221] Tin Kam Ho. “The Random Subspace Method for Constructing Decision Forests”. In: *IEEE Transactions of Pattern Analysis and Machine Intelligence* 20 (1998), pp. 832–844.
- [222] V. F. Rodriguez-Galiano et al. “An assessment of the effectiveness of a random forest classifier for land-cover classification”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* 67 (2012), pp. 93–104.
- [223] John Canny. *A Computational Approach to Edge Detection*. 1986.
- [224] P. P. Mohanta, D. P. Mukherjee, and S. T. Acton. “Agglomerative clustering for image segmentation”. In: Proceedings of the 16th International Conference on Pattern Recognition. 2002, pp. 664–667. DOI: 10.1109/ICPR.2002.1044838.
- [225] L. N. Olson and J. B. Schroder. *PyAMG: Algebraic Multigrid Solvers in Python v4.0*. Release 4.0. 2018. URL: <https://github.com/pyamg/pyamg>.
- [226] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [227] Ron Kohavi. “A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection”. In: International Joint Conference on Artificial Intelligence. 1995.
- [228] Lucas Mentch and Giles Hooker. “Quantifying Uncertainty in Random Forests via Confidence Intervals and Hypothesis Test”. In: *Journal of Machine Learning Research* 17 (2016).
- [229] Robert F. Berger. “Design Principles for the Atomic and Electronic Structure of Halide Perovskite Photovoltaic Materials: Insights from Computation”. In: *Chemistry - A European Journal* 24.35 (2018), pp. 8708–8716. DOI: 10.1002/chem.201706126.
- [230] Kiyoshi Miyata, Timothy L. Atallah, and X.-Y. Zhu. “Lead halide perovskites: Crystal-liquid duality, phonon glass electron crystals, and large polaron formation”. In: *Science Advances* 3 (2017).

- [231] Ying Zhou et al. “Flexible Linearly Polarized Photodetectors Based on All-Inorganic Perovskite CsPbI₃ Nanowires”. In: *Advanced Optical Materials* 1800679 (2018), p. 1800679. DOI: 10.1002/adom.201800679.
- [232] Yanqiang Hu et al. “Highly Efficient Flexible Solar Cells Based on Room-Temperature Processed Inorganic Perovskite”. In: *Journal of Materials Chemistry A* (2018). DOI: 10.1039/C8TA06719H.
- [233] Haining Chen et al. “Inorganic Perovskite Solar Cells: A Rapidly Growing Field”. In: *Solar RRL* 6 (2018), p. 1700188. DOI: 10.1002/solr.201700188.
- [234] Yue Wang et al. “All-Inorganic Colloidal Perovskite Quantum Dots: A New Class of Lasing Materials with Favorable Characteristics”. In: *Advanced Materials* 27.44 (2015), pp. 7101–7108. DOI: 10.1002/adma.201503573.
- [235] Yongping Fu et al. “Broad Wavelength Tunable Robust Lasing from Single-Crystal Nanowires of Cesium Lead Halide Perovskites (CsPbX₃, X = Cl, Br, I)”. In: *ACS Nano* 10.8 (2016), pp. 7963–7972. DOI: 10.1021/acsnano.6b03916.
- [236] Shilpa N. Raja et al. “Encapsulation of Perovskite Nanocrystals into Macroscale Polymer Matrices: Enhanced Stability and Polarization”. In: *ACS Applied Materials & Interfaces* 8 (2016), pp. 35523–35533. DOI: 10.1021/acsmi.6b09443.
- [237] Jingru Zhang et al. “3D-2D-0D Interface Profiling for Record Efficiency All-Inorganic CsPbBrI₂ Perovskite Solar Cells with Superior Stability”. In: *Advanced Energy Materials* (2018), p. 1703246. DOI: 10.1002/aenm.201703246.
- [238] Qiushui Chen et al. “All-inorganic perovskite nanocrystal scintillators”. In: *Nature* 561 (2018), pp. 88–93.
- [239] Yupeng Zhang et al. “Synthesis, properties, and optical applications of low-dimensional perovskites”. In: *Chemical Communications* 52.94 (2016), pp. 13637–13655. DOI: 10.1039/c6cc06425f.
- [240] Meng-Ya Zhu et al. “Ppt-level benzene detection and gas sensing mechanism using (C₄H₉NH₃)₂PbI₂Br₂ organic-inorganic layered perovskite”. In: *Inorganic Chemistry Frontiers* (2018). DOI: 10.1039/C8QI00803E.
- [241] Yehonadav Bekenstein et al. “Highly Luminescent Colloidal Nanoplates of Perovskite Cesium Lead Halide and Their Oriented Assemblies”. In: *Journal of the American Chemical Society* 137.51 (2015), pp. 16008–16011.
- [242] Dandan Zhang et al. “Synthesis of Composition Tunable and Highly Luminescent Cesium Lead Halide Nanowires through Anion-Exchange Reactions”. In: *Journal of the American Chemical Society* 138.23 (2016), pp. 7236–7239. DOI: 10.1021/jacs.6b03134.
- [243] Loredana Protesescu et al. “Nanocrystals of Cesium Lead Halide Perovskites (CsPbX₃, X = Cl, Br, and I): Novel Optoelectronic Materials Showing Bright Emission with Wide Color Gamut”. In: *Nano Letters* 15 (2015).

- [244] Muhammad Imran et al. “Colloidal Synthesis of Strongly Fluorescent CsPbBr₃Nanowires with Width Tunable down to the Quantum Confinement Regime”. In: *Chemistry of Materials* 28.18 (2016), pp. 6450–6454. DOI: 10.1021/acs.chemmater.6b03081.
- [245] Vikash Kumar Ravi et al. “Excellent green but less impressive blue luminescence from CsPbBr₃ perovskite nanocubes and nanoplatelets”. In: *Nanotechnology* 27 (2016), p. 325708.
- [246] Dandan Zhang et al. “Ultrathin Colloidal Cesium Lead Halide Perovskite Nanowires”. In: *Journal of the American Chemical Society* (2016), jacs.6b08373. DOI: 10.1021/jacs.6b08373.
- [247] Q. Zhu et al. “Transition from two-dimensional to three-dimensional quantum confinement in semiconductor quantum wires/quantum dots”. In: *Nano Letters* 7.8 (2007), pp. 2227–2233. DOI: 10.1021/nl0706650.
- [248] Jiangtao Hu et al. “Linearly Polarized Emission from Colloidal Semiconductor Quantum Rods”. In: *Science* 292 (2001), pp. 2060–2064.
- [249] H. Htoon et al. “Effect of zero- to one-dimensional transformation on multiparticle auger recombination in semiconductor quantum rods”. In: *Physical Review Letters* 91.22 (2003), pp. 1–4. DOI: 10.1103/PhysRevLett.91.227401.
- [250] Michael A. Becker et al. “Bright triplet excitons in caesium lead halide perovskites”. In: *Nature* 553.7687 (2018), pp. 189–193.
- [251] J. C. Blancon et al. “Extremely efficient internal exciton dissociation through edge states in layered 2D perovskites”. In: *Science* 355.6331 (2017), pp. 1288–1292. DOI: 10.1126/science.aal4211.
- [252] Francesco Di Stasio et al. “Reversible Concentration-Dependent Photoluminescence Quenching and Change of Emission Color in CsPbBr₃Nanowires and Nanoplatelets”. In: *Journal of Physical Chemistry Letters* 8.12 (2017), pp. 2725–2729. DOI: 10.1021/acs.jpcclett.7b01305.
- [253] Juan A Castaneda et al. “Efficient Biexciton Interaction in Perovskite Quantum Dots Under Weak and Strong Confinement”. In: *ACS Nano* 10 (2016), pp. 8603–8609. DOI: 10.1021/acsnano.6b03908.
- [254] Navendu Mondal and Anunay Samanta. “Complete ultrafast charge carrier dynamics in photo-excited all-inorganic perovskite nanocrystals (CsPbX₃)”. In: *Nanoscale* 9 (2017), pp. 1878–1885. DOI: 10.1039/c6nr09422h.
- [255] Natalie A. Gibson et al. “Excitation Intensity Dependence of Photoluminescence Blinking in CsPbBr₃ Perovskite Nanocrystals”. In: *Journal of Physical Chemistry C* 122 (2018), pp. 12106–12113. DOI: 10.1021/acs.jpcc.8b03206.
- [256] Gurivi Reddy Yettapu et al. “Terahertz Conductivity within Colloidal CsPbBr₃ Perovskite Nanocrystals: Remarkably High Carrier Mobilities and Large Diffusion Lengths”. In: *Nano Letters* 16 (2016), pp. 4838–4848. DOI: 10.1021/acs.nanolett.6b01168.

- [257] Dandan Zhang et al. “Solution-Phase Synthesis of Cesium Lead Halide Perovskite Nanowires”. In: *Journal of the American Chemical Society* 137.29 (2015), pp. 9230–9233. DOI: 10.1021/jacs.5b05404.
- [258] Yu Tong et al. “From Precursor Powders to CsPbX₃ Perovskite Nanowires: One-Pot Synthesis, Growth Mechanism, and Oriented Self-Assembly”. In: *Angewandte Chemie - International Edition* 56.44 (2017), pp. 13887–13892. DOI: 10.1002/anie.201707224.
- [259] Makhsud I. Saidaminov et al. “Inorganic Lead Halide Perovskite Single Crystals: Phase-Selective Low-Temperature Growth, Carrier Transport Properties, and Self-Powered Photodetection”. In: *Advanced Optical Materials* 5 (2017).
- [260] Ji-Hyun Cha et al. “Photoresponse of CsPbBr₃ and Cs₄PbBr₆ Perovskite Single Crystals”. In: *Journal of Physical Chemistry Letters* 8 (2017), pp. 565–570.
- [261] Roman Krahné et al. “Physical properties of elongated inorganic nanoparticles”. In: *Physics Reports* 501 (2011), pp. 75–221. DOI: 10.1016/j.physrep.2011.01.001.
- [262] Murad Ahmad et al. “Structural, electronic and optical properties of CsPbX₃(X=Cl, Br, I) for energy storage and hybrid solar cell applications”. In: *Journal of Alloys and Compounds* 705 (2017), pp. 828–839.
- [263] Julien Ramade et al. “Fine structure of excitons and electron-hole exchange energy in polymorphic CsPbBr₃ single nanocrystals”. In: *Nanoscale* 10 (2018), pp. 6393–6401.
- [264] Michael C. Brennan et al. “Origin of the Size-Dependent Stokes Shift in CsPbBr₃ Perovskite Nanocrystals”. In: *Journal of the American Chemical Society* 139.35 (2017), pp. 12201–12208. DOI: 10.1021/jacs.7b05683.
- [265] Haiming Zhu et al. “Screening in crystalline liquids protects energetic carriers in hybrid perovskites”. In: *Science* 353 (2016), pp. 1409–1413.
- [266] Abhishek Swarnkar et al. “Colloidal CsPbBr₃ Perovskite Nanocrystals: Luminescence beyond Traditional Quantum Dots”. In: *Angewandte Chemie* 127.51 (2015), pp. 15644–15648. DOI: 10.1002/ange.201508276.
- [267] Jiaming Li et al. “Temperature-dependent photoluminescence of inorganic perovskite nanocrystal films”. In: *RSC Advances* 6.82 (2016), pp. 78311–78316. DOI: 10.1039/c6ra17008k.
- [268] Haiming Zhu et al. “Lead halide perovskite nanowire lasers with low lasing thresholds and high quality factors.” In: *Nature materials* 14.6 (2015), pp. 636–642. DOI: 10.1038/nmat4271.
- [269] Milan Delor et al. “Universal imaging of material functionality through nanoscale tracking of energy flow”. In: *ArXiv* 1805.09982 (2018). DOI: 1805/1805.09982.
- [270] Murray Geller and Edward W. Ng. “A Table of Integrals of the Error Function. II. Additions and Corrections.” In: *Journal of Research of the National Bureau of Standards B*. 75B (1971).

- [271] Wenming Tian et al. “Visualizing Carrier Diffusion in Individual Single-Crystal Organolead Halide Perovskite Nanowires and Nanoplates”. In: *Journal of the American Chemical Society* 137.39 (2015), pp. 12458–12461. DOI: 10.1021/jacs.5b08045.
- [272] Jizhong Song et al. “Ultralarge All-Inorganic Perovskite Bulk Single Crystal for High-Performance Visible-Infrared Dual-Modal Photodetectors”. In: *Advanced Optical Materials* 5.12 (2017), pp. 1–8. DOI: 10.1002/adom.201700157.
- [273] Jonathan De Roo et al. “Highly Dynamic Ligand Binding and Light Absorption Coefficient of Cesium Lead Bromide Perovskite Nanocrystals”. In: *ACS Nano* 10.2 (2016), pp. 2071–2081. DOI: 10.1021/acsnano.5b06295.
- [274] Michael B. Price et al. “Hot-carrier cooling and photoinduced refractive index changes in organic-inorganic lead halide perovskites”. In: *Nature Communications* 6.May (2015), pp. 1–8. DOI: 10.1038/ncomms9420.
- [275] Jianhui Fu et al. “Hot carrier cooling mechanisms in halide perovskites”. In: *Nature Communications* 8.1 (2017). DOI: 10.1038/s41467-017-01360-3.
- [276] Mingjie Li et al. “Slow cooling and highly efficient extraction of hot carriers in colloidal perovskite nanocrystals”. In: *Nature Communications* 8 (2017). DOI: 10.1038/ncomms14350.
- [277] Hamzeh Telfah et al. “Ultrafast Exciton Dynamics in Shape-Controlled Methylammonium Lead Bromide Perovskite Nanostructures: Effect of Quantum Confinement on Charge Carrier Recombination”. In: *Journal of Physical Chemistry C* 121.51 (2017), pp. 28556–28565. DOI: 10.1021/acs.jpcc.7b10377.
- [278] Christopher Grieco et al. “Dynamic Exchange During Triplet Transport in Nanocrystalline TIPS-Pentacene Films”. In: *Journal of the American Chemical Society* 138.49 (2016), pp. 16069–16080. DOI: 10.1021/jacs.6b10010.
- [279] Elinore M. L. D. de Jong et al. “Multiexciton Lifetime in All-Inorganic CsPbBr₃ Perovskite Nanocrystals”. In: *The Journal of Physical Chemistry C* 121 (2017), pp. 1941–1947.
- [280] H. E. Ruda and A. Shik. “Polarization-sensitive optical phenomena in semiconducting and metallic nanowires”. In: *Physical Review B* 72.11 (2005), p. 115308. DOI: 10.1103/PhysRevB.72.115308.
- [281] N. Le Thomas et al. “Exciton Fine Structure in Single CdSe Nanorods”. In: *Physical Review Letters* 94 (2005), p. 016803.
- [282] Daniela Täuber et al. “Exploring the Electronic Band Structure of Organometal Halide Perovskite via Photoluminescence Anisotropy of Individual Nanocrystals”. In: *Nano Letters* 16.8 (2016), pp. 5087–5094.
- [283] A. L. Efros et al. “Nonlinear Optical Effects in Porous Silicon: Photoluminescence Saturation and Optically Induced Polarization Anisotropy”. In: *Physical Review B* 56.7 (1997), pp. 3875–3884.

- [284] University of Washington Dalton Research Group. *Dielectric Constants of Common Solvents*. URL: https://depts.washington.edu/eo optic/linkfiles/dielectric_chart%5B1%5D.pdf (visited on 11/15/2018).
- [285] X. Ma et al. “Anisotropic Photoluminescence from Isotropic Optical Transition Dipoles in Semiconductor Nanoplatelets”. In: *Nano Letters* 18 (2018), pp. 4647–4652.
- [286] Jia-Shiang Chen et al. “0D-2D and 1D-2D Semiconductor Hybrids Composed of All Inorganic Perovskite Nanocrystals and Single-Layer Graphene with Improved Light Harvesting”. In: *Particle & Particle Systems Characterization* 1700310 (2017), p. 1700310. DOI: 10.1002/ppsc.201700310.
- [287] B.T. Diroll, H. Zhou, and R.D. Schaller. “Low-Temperature Absorption, Photoluminescence, and Lifetime of CsPbX₃ (X = Cl, Br, I) Nanocrystals”. In: *Advanced Functional Materials* 3 (2018), pp. 1–7. DOI: 10.1002/adfm.201800945.
- [288] Ming Fu et al. “Neutral and Charged Exciton Fine Structure in Single Lead Halide Perovskite Nanocrystals Revealed by Magneto-optical Spectroscopy”. In: *Nano Letters* 17.5 (2017), pp. 2895–2901. DOI: 10.1021/acs.nanolett.7b00064.
- [289] M. O. Nestoklon et al. “Optical orientation and alignment of excitons in ensembles of inorganic perovskite nanocrystals”. In: *Physical Review B* 97.23 (2018), pp. 1–10. DOI: 10.1103/PhysRevB.97.235304.

Appendix A

Alignment Procedure for Transient Absorption Microscope

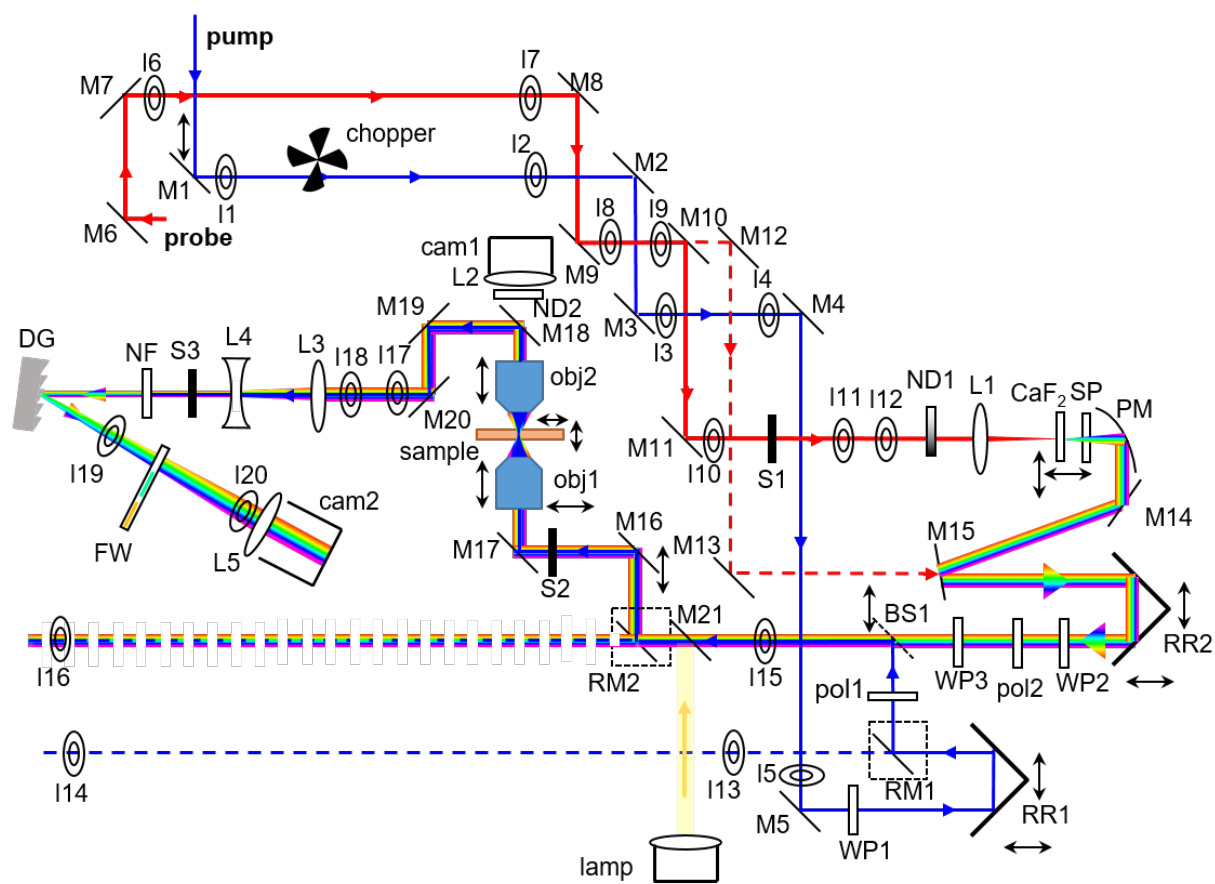


Figure A.1: Diagram of TAM experiment

A diagram of the TAM setup is shown in Figure A.1. A legend can be found in Table A.2, and selected part numbers in Table A.1. This diagram omits upstream parts of the setup in which the pump is generated and shaped. The pump pulse may be generated in the optical parametric amplifier (OPA) or by focusing the 800 nm beam into a BBO crystal to generate 400 nm light. It then may go through a prism compressor, a telescope, and a spatial filter before entering the diagram below.

The probe begins with 800 nm light that comes out of the Regenerative Amplifier. It then goes through several passes to set its path length equal to that of the pump before it enters the diagram below.

In the following discussion, the z -axis refers to the beam's propagation direction. Of the remaining two dimensions, x is parallel and y is perpendicular to the table. A glossary of abbreviations is at the end of this appendix.

Align Pump:

1. Use mirror M1 to align the pump through irises I1 and I2. M1 is on a translation stage (indicated by arrows). Translate to align the beam on I1, use the knobs on M1 to align to I2, and iterate.
2. Use mirrors M2 and M3 to align to irises I3 and I4.
3. Use mirror M4 to align to iris I5.
4. Take out removable mirror RM1. Use retroreflector RR1 and mirror M5 to align to irises I13 and I14. RR1 can be translated along both the x and y directions with micrometers in order to align to I13. RR1 is on a motorized delay stage that can be translated 150 mm along the z direction. Translate RR1 from one end to the other and check that the alignment through I14 remains stable (for an explanation of how to control this and other motorized elements, see Appendix B). Put RM1 back in place.

Align Probe:

5. Use mirrors M6 and M7 to align to irises I6 and I7.
6. Use mirrors M8 and M9 to align to irises I8 and I9.
7. If a degenerate or two-color geometry is being used, M10 can be flipped down and the probe follows the dashed line. Skip the next few steps, flip down M15 and use M13 in its place to align through I15 and I16.
8. Open shutter S1, but make sure to place some other beam block before the CaF₂ crystal. Use M10 and M11 to align through I10 and I11.

9. Use I12 and neutral density filter wheel ND1 to generate white light in calcium fluoride crystal CaF_2 . Important parameters are the focusing conditions (adjusted with I12), the power (adjusted with I12 and also ND1), and the position of the focus inside the crystal (adjusted by translating CaF_2 along the z -axis with a micrometer). The crystal must be continuously translated to avoid burning. This is done with a translational stage connected to a motor that is powered by a 12 Volt power supply and controlled with an Arduino.

If the white light generation needs to be completely re-aligned, remove ND1 and CaF_2 , and insert an iris just after CaF_2 but before PM. Use M10 and M11 to align to I10 and this new iris. Insert lens L1 to approximately the correct height. Rotate it around the y -axis until the back-reflection propagates straight back. Adjust its x - and y -positions until the focal spot is aligned through the new iris. Insert ND1 and align its tilt until the back-reflection propagates almost straight back (you don't want it going all the way back into the Regen). Rotate ND1 so that the neutral density is as high as possible. Close I12 about half-way down. Re-insert CaF_2 and slowly decrease the neutral density of ND1 until white light appears. Translate CaF_2 along the z -axis to find the position that maximizes the white light strength, continuously changing the neutral density of ND1 so the light does not become too bright. Finally play with ND1 and I12 to make a stable white light spot with no filaments that does not burn the crystal.

10. Shortpass filter SP filters out redder wavelengths that carry most of the white light's intensity, but may not be relevant for the experiment. Make sure to insert an appropriate filter.
11. Parabolic mirror PM should collimate the white light. Small adjustments along the y - and z -directions can be made with the mount, as can small rotations about the x -axis. Large re-alignment requires detaching the base from the optical table and moving PM by hand. Flip down M15 and use M14 and another temporary mirror to direct the beam out of the enclosed area. The longer a path length you have to examine the beam spot, the better for collimation. Unfortunately, in this setup there's no easy way to align PM in a controllable manner. Just keep at it and eventually you'll hit upon something that's good enough. Don't be tempted by perfection—you don't want to find yourself an hour later, frustrated and cursing because you've made it worse and you can't get back to your good-enough configuration.
12. Remove RM2 and Use RR2 and M15 to align through beam splitter BS1 to I15 and I16. RR2 can be translated along the x - or y -directions with micrometers to align to I15. M21 should be flipped down. If the initial alignment is poor you can remove half-wave plates WP2 and WP3 and polarizer pol2 at first. They are on magnetic mounts, and can easily be popped back in to position. Make sure pol2 is set to 0° (vertically polarized, which in the microscope will become horizontally polarized). You can adjusted WP2 to control the brightness of the beam.

Image sample:

13. Block both laser beams, flip up M21 and turn on the lamp.
14. Insert the sample. Set sample $X = Y = Z = 100 \mu\text{m}$. Set the bottom objective, obj1, about 3 mm below the sample so that light will illuminate a broad spot. $Z_b \approx -3.0 \text{ mm}$.
15. Flip M18 out of the way so that light is directed into camera cam1.
16. Set the neutral density in front of the camera, ND2. About 2.0 is good for imaging samples.
17. Open S2 to view the sample. The imaging lens L2 should be fixed. Adjust the z -position of top objective obj2 to bring the sample in to focus. If necessary adjust Z_b so that dust spots are not being imaged. Note the Z_t position. Move around in the $x - y$ plane with the large micrometers to find a good spot on the sample to study. Save an image of this spot.
18. Turn off the lamp and flip down M21.

Image Laser Beams:

19. Rotate polarizer poll1 to the desired polarization of the pump beam. Rotate half-wave plate WP1 to maximize the pump brightness.
20. Use RM1 and BS1 to align the pump beam through I15 and I16. About 30% of the pump is lost because it transmits through BS1, and about 70% of the probe is lost because it reflects off of BS1. A beam stop must be in place to catch them.
21. Adjusting BS1 will have slightly altered the path of the probe. Check the alignment of the probe through I15 and I16, adjusting if necessary with RR2 and M15.
22. Set the neutral density ND2. 4.0 works well for imaging beams.
23. Use WP1 and WP2 to set both beams to be fairly dim (lower in power than you would use in an experiment, as we do not want to damage the sample).
24. Block the pump. Insert RM2 and open S2.
25. Move obj1 towards the sample in steps of 0.1 mm until the probe appears on cam1. Use adjustments as small as 0.002 mm to make the probe spot as small and circular as possible.
26. Mark the probe position, block the probe, and image the pump. It should be near the probe and slightly larger. Slightly tweak BS1 to overlap pump and probe. You can also slightly adjust the z -position of obj1 to balance the two beam images.




Figure A.2: Images of focused and overlapped (a) pump and (b) probe.

27. Adjust the sample position with the PI stage, making sure that the beam images do not change. If they do it indicates that the sample is heterogeneous and is scattering the beams. Make sure you're imaging the beams when they pass through a homogeneous region.
28. Use WP1 and WP2 to make the beams as bright as possible on the camera without saturating. Record their individual images. For the pump, record an image when the delay stage is at -75 mm, 0 mm, and +75 mm. For the probe, record at WP3 = 0° (horizontally polarized) and 45° (vertically polarized). The beams should be stable by eye under these transformations, but having the images saved allows you to measure how stable they really are. It also allows you to measure overlap and beam sizes. Example images of an overlapped pump and probe are shown in Figure A.2.
29. If the beams appear in an odd position or are smeared out, this indicates that they are not being directed into the center of obj1. First flip down M16 and insert an iris behind M16. If the beam does not go through this temporary iris, adjust the vertical tilt of RM2 to align it. Flip up M16. Use the tilt of M16 as well as its translation to set the beam images.

Align Into Detector:

30. Flip M18 back into position.
31. Use M19 and M20 to align to I17 and I18. This is difficult for several reasons. One is that it's tough to see I17 and I18. Another is that the probe is really large, and lots of the beam gets clipped during travel, making it tough to pin down the center of the beam. Currently obj1 has a higher numerical aperture than obj2, so they act as a beam expander. Replacing obj2 to match obj1 would solve this problem. It would also allow for higher magnification imaging of the sample, at the expense of not being able to see as large an area on the camera.
32. Lenses L3 and L4 act to shrink the probe. If the probe is not collimated out of L4, adjust the z -position of L4 by sliding its mount along the cage rails.

33. Open S3 and flip up notch filter NF to block the pump beam.
34. Adjust diffraction grating DG to send the probe through I19 and I20. If a larger re-alignment is needed, the entire breadboard containing the camera and irises can be moved. This alignment procedure could use some improvement, perhaps by adding a mirror before DG so that there are two additional degrees of freedom. L5 should be set to the correct distance from cam2, and should not need to be adjusted.
35. See if cam2 picks up the probe spectrum. The signal strength is particularly sensitive to the vertical tilt of DG, so if there's no signal try adjusting that. Make the signal as strong as possible, then turn down the probe power with WP2 so that the camera does not saturate. Rotate the probe polarization with WP3 and make sure there's no saturation for any polarization you will use in your measurement (transmission efficiency through the microscope is a function of polarization).
36. Adjust the horizontal tilt of DG so that the entire probe spectrum is imaged on cam2 without clipping. Rotate filter wheel FW to bring each bandpass filter into the beam path, and note where it is imaged by cam2. Use this information to perform the wavelength calibration procedure (see Appendix B).

Align Sample:

37. From imaging the sample, you should have an idea of approximately where the region of interest on the sample is. Block the pump and scan around in X and Y using the PI stage, measuring the probe transmission. The result may look something like Figure A.3a, which was taken for a horizontally-oriented nanowire. The nanowire is seen to be in the region $80 \leq X \leq 100 \mu\text{m}$ and $107 \leq Y \leq 112 \mu\text{m}$. *A word of caution:* the image in the LabView program is rotated by 180° relative to what you see on the camera.
Once you have a rough idea of where the sample is, do finer scans to find the best scan position. Figure A.3b shows a transmission scan for $X = 88 \mu\text{m}$; the best position is $Y = 110 \mu\text{m}$.
38. Let the pump through and try to observe a ΔOD signal. You should have some idea of where time-0 is; make sure $T > 0$. Scan in Z to find the position that gives you the strongest ΔOD signal.
39. Scan the delay stage and do a rough time-0 scan. An example is shown in Figure A.4 for a specific wavelength—time-0 is around the drop at 3.0 ps. The exact position of time-0 will depend on λ .
40. You are now ready to perform TAM scans. If you need a full wavelength-dependent instrument response function and time-0, move to a blank spot on the sample and finely scan around the time ≈ 0 region. You will have to adjust Z to find the plane of the sample. Otherwise you can go on your merry way and perform whatever scan protocol your heart desires.

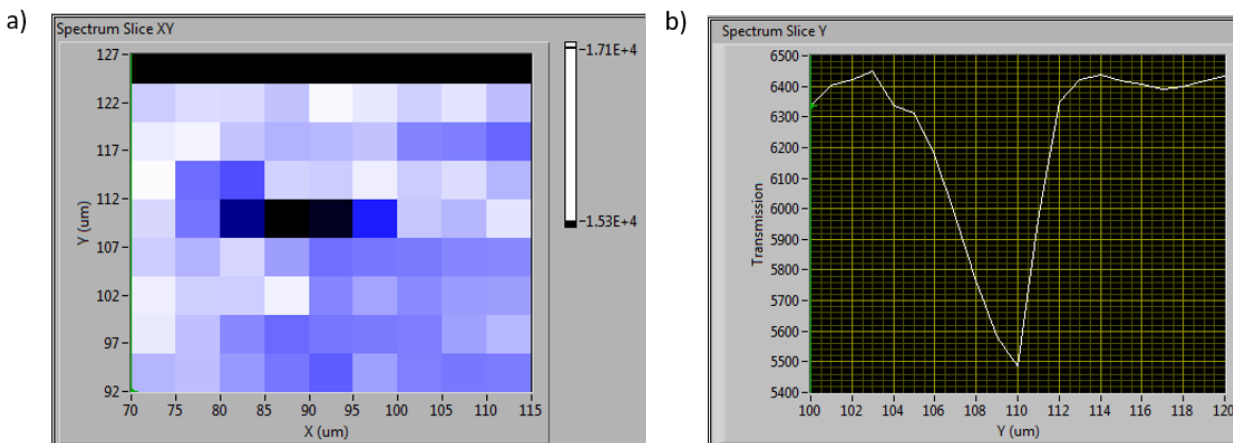


Figure A.3: Transmission scans to find the sample. (a) Rough scan in X and Y , (b) fine scan in Y .




Figure A.4: Time-0 scan at $\lambda = 517$ nm.

Part Information:
Abbreviations:

label	vendor	part number
CaF ₂	Newlight Photonics	CAF1020-C
BS1	ThorLabs	BSS10R
cam1	ThorLabs	DCC1545M
cam2	LightWise	Allegro Orion 2K Line Scan Camera
DG	ThorLabs	GR25-0305
obj1	Leica	HC PL APO 25x/0.70 CS
obj2	Leica	HC PL APO 10x/0.40 CS
sample stage	PI (Physik Instrumente)	PInano XYZ
rotation mounts	Newport	PR50CC
retroreflector delay stage	Newport	ILS150CC
objective translation stages	Newport	VP-25XL

Table A.1: Part numbers of important components of TA microscope

Abbreviation	Description
BS	beam splitter
cam	camera
DG	diffraction grating
FW	filter wheel
I	iris
L	lens
ND	neutral density filter
obj	objective lens
PW	parabolic mirror
pol	polarizer
RM	removable mirror
RR	retroreflector
S	shutter
SP	shortpass filter
WP	half-wave plate

Table A.2: Abbreviations used in Figure A.1

Appendix B

Laview Program for Controlling Transient Absorption Microscope

The transient absorption microscope is controlled with a LabView program described here. The front panel consists of three panes, “Setup,” “Microscope Control,” and “Wavelength Control.” Each is shown below and the functionality of each area is described. The wire diagram is not discussed here.

Setup:

- A1 Connection to PI stage for fine control of sample position. When program starts, the “PI On” button will become selectable. If there is an error when connecting to the PI stage, it will appear in the “PI error” box.
- A2 Connection to XPS controller for objectives, delay stage, and motorized optics. The dropdown menu selects the IP address of the XPS. The option “XPS-C8” is hard-coded to 192.168.0.1. If there is an error when connecting to the XPS, it will appear in the “XPS 1” or “XPS 2” boxes on the right.
- A3 Connection to the camera for data acquisition. The dropdown menu auto-populates with the names of all cameras connected to the computer. In general there should be two—one for imaging the sample and one for data acquisition. Make sure to select the correct option before pressing the “Connect to Camera” button.
- A4 Camera controls and testing. After selecting new options, make sure to press “Set Attributes.” The most commonly changed option is “Height,” which controls the number of lines to acquire at a time. I generally use 500 lines, which at 5 kHz takes 100 ms. “CVC Gain” should always be set to “Low,” unless the signal is very weak. Changing the exposure time has no effect, since the pulse width is much shorter than the exposure time. The pixels can be binned horizontally, but I have found that it yields no improvement in either signal/noise ratio or acquisition time, and introduces

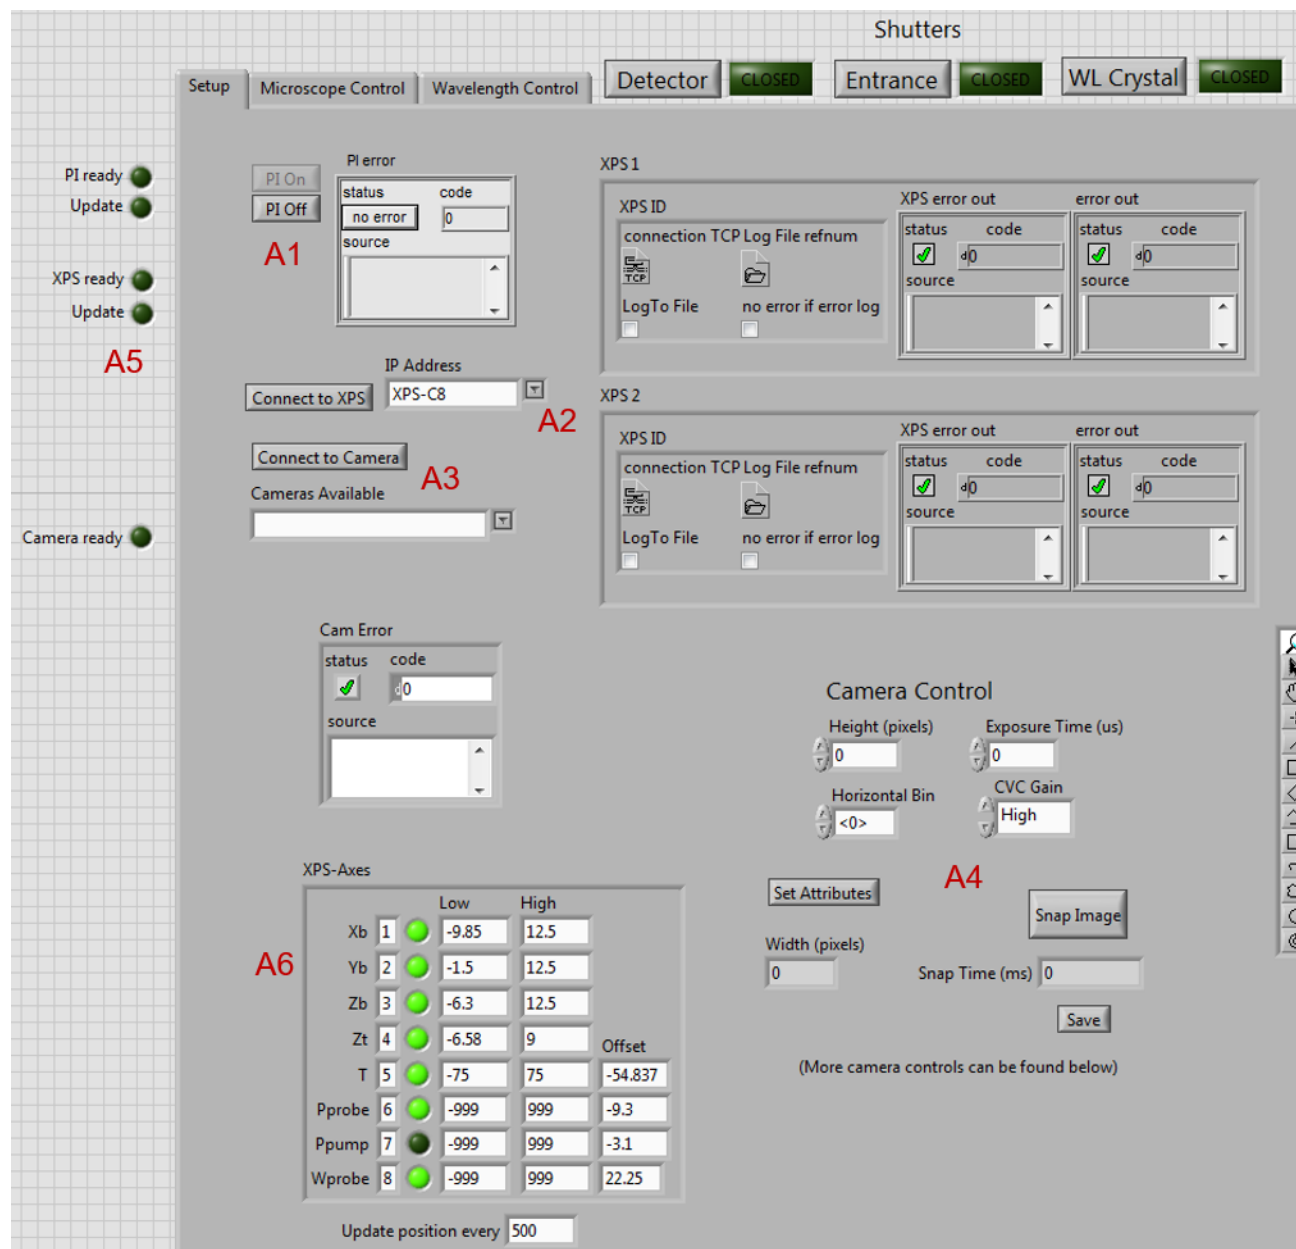


Figure B.1: Left side of “Setup” screen

some complications, so I use no binning which leads to “Width (pixels)” equal to 2064. “Snap Image” takes a series of lines and stacks them vertically to form an image. By comparing the “Snap Time (ms)” with the expected time (number of lines divided by laser rep rate), you can get a measure of how much computational overhead there is. The first image always has more overhead than subsequent images.




Figure B.2: Right side of “Setup” screen

A5 Indicator lights that turn green if the program is connected to the PI stage, XPS, and camera. “Update” buttons should be turned on if you want the stage positions to be updated in real-time during a scan.

A6 Limits and offsets for the XPS stages. The stages control: X position of bottom objective, Y position of bottom objective, Z position of bottom objective, Z position

of top objective, linear position of delay stage, angle of probe polarizer, angle of pump polarizer (not in use), angle of probe half-wave plate. “Low” and “High” are the lowest and highest positions the stages can move to. This is important for the objectives because otherwise they can crash into things. “Offset” is the offset of the stage, which for the delay stage is related to time-0, and for the rotation stages is when their fast axes are oriented vertically. For example, if you tell the probe polarizer to go to 0° , the software will actually move to -9.3° , but the displayed position will be 0° and the optic will be oriented with its fast axis vertical.

- A7 Camera image snapped as described in A4. Can be saved. Useful if you want to record individual laser pulses.
- A8 File paths that are hard coded into the software and should not be changed. The first path stores data files as they are acquired during scans, before being saved. The second stores the the file that logs every time-0 position.

Microscope Control:

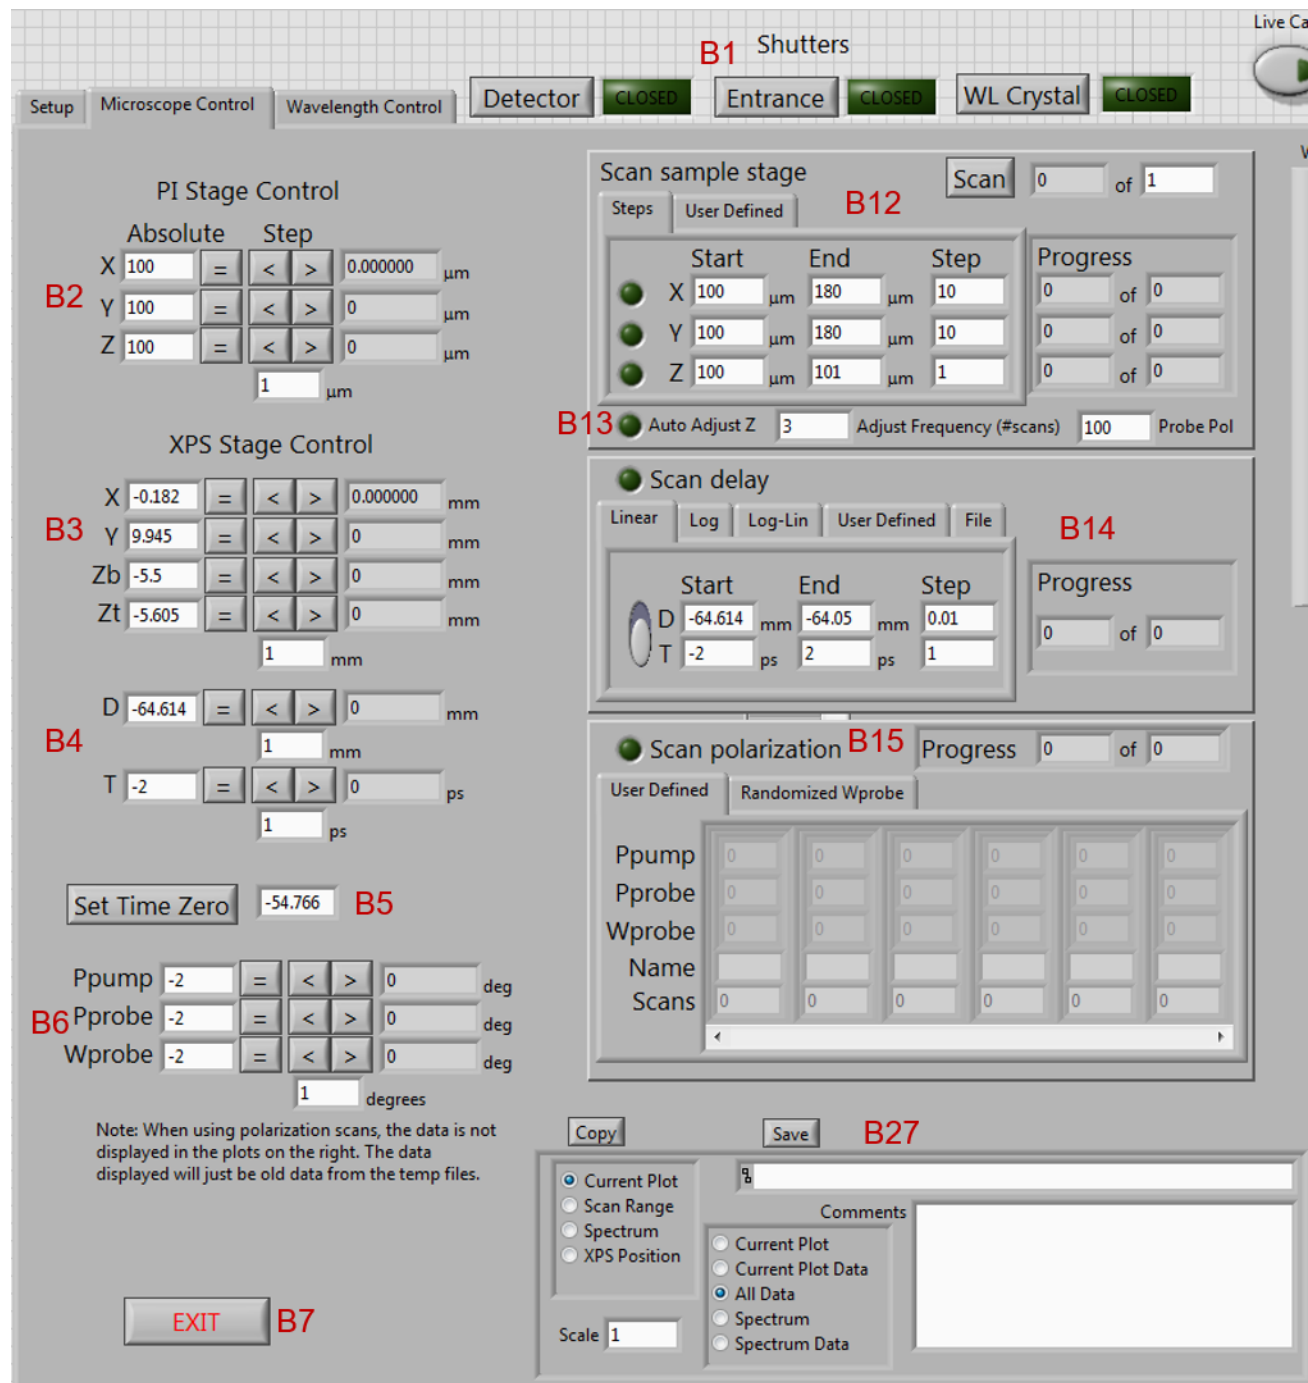


Figure B.3: Left side of “Microscope Control” screen

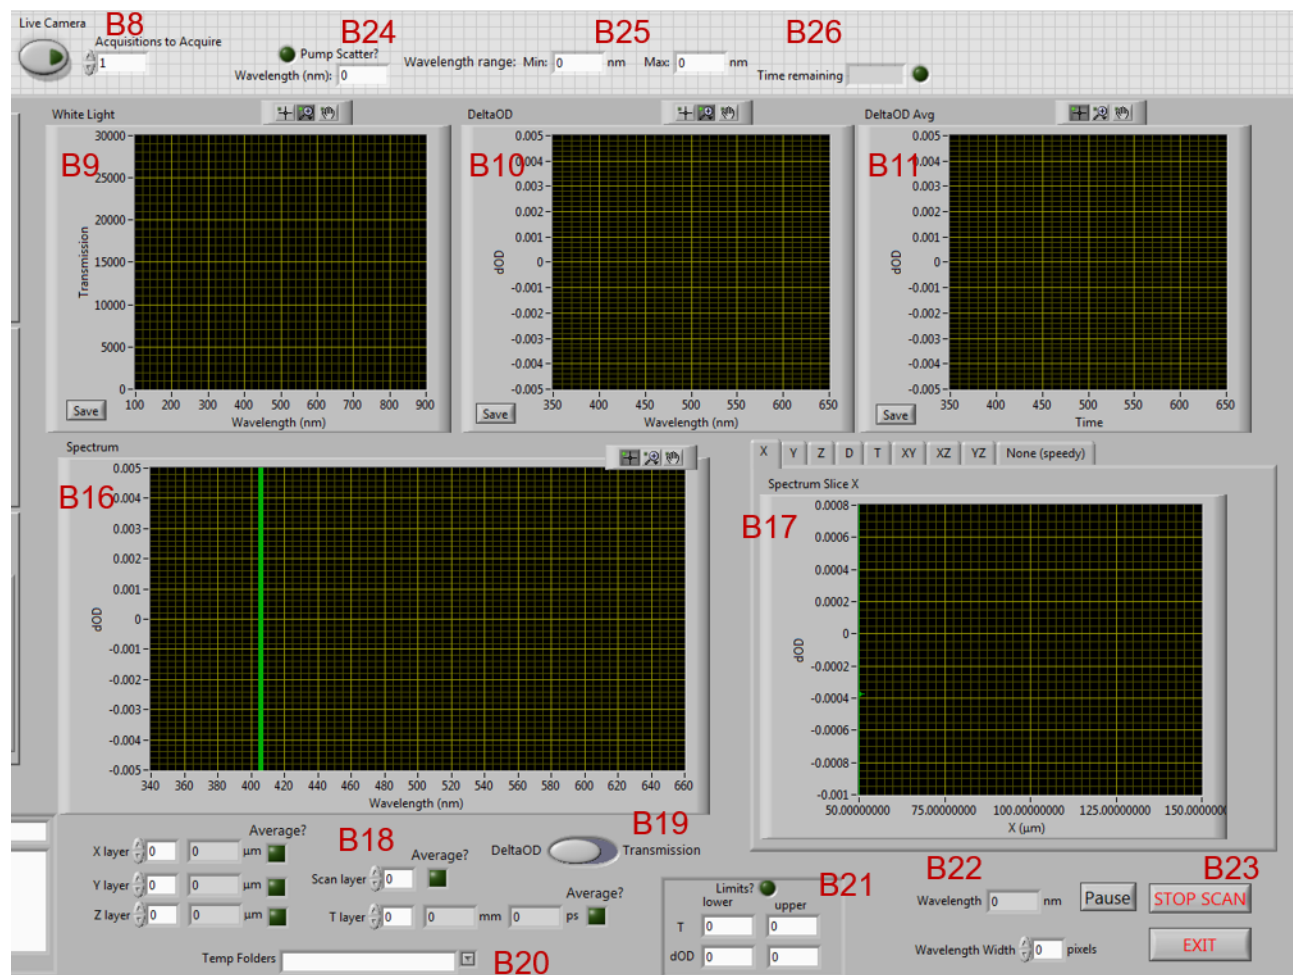


Figure B.4: Right side of “Microscope Control” screen

- B1 Buttons to open/close the shutters before the detector, at the microscope entrance, and before the white light generating crystal (currently removed because it was unreliable). These buttons only work when in the “Microscope Control” pane. If they are not working, check the shutter controller boxes and make sure they are not in manual mode.
- B2 PI stage control. Set the X , Y , and Z positions either absolutely (using the equals sign) or relatively (using the j and i buttons to step by an amount in the lower box). When connected to the PI stage, the current positions are displayed. All positions are in μm .
- B3 XPS stage control for objectives. Set the X , Y , and Z positions of the bottom objective or Z position of the top objective. This can be done absolutely (using the equals sign)

or relatively (using the j and k buttons to step by an amount in the lower box). When connected to the XPS, the current positions are displayed. All positions are in mm.

- B4 XPS stage control for delay stage position. Position can be set either in units of distance, D (mm), or time, T (ps). Distance ranges from -75 to +75 mm. The corresponding time depends on the current position of time-0, d_0 . The conversion is 6.671 ps/mm, so $T = 6.671 \times (D - d_0)$. As above, movement can be done either absolutely or relatively.
- B5 Pressing the “Set Time Zero” button takes the number in the box and sets that position equal to time-0. When the delay stage is set to that position the pump and probe are expected to arrive simultaneously, the the time T will be equal to 0 ps. Every time the button is pressed, the new time-0 is logged (see A8).
- B6 XPS stage control for rotation optics. Set the orientation of the pump polarizer, probe polarizer, or probe half-wave plate. This can be done absolutely (using the equals sign) or relatively (using the j and k buttons to step by an amount in the lower box). When connected to the XPS, the current positions are displayed. All positions are in degrees.
- B7 Exits the program.
- B8 Press “Live Camera” to begin acquiring from the camera. “Acquisitions to Acquire” tells the program how many acquisition to average together. The number of lines in a single acquisition is set in the “Setup” pane (see A4).
- B9 Displays the average transmission signal after each acquisition.
- B10 Displays the average ΔOD signal after each acquisition.
- B11 Displays ΔOD averaged over several acquisition. The number of acquisition to average is set in B8.
- B12 Set scan parameters for the PI stage and begin scan. For each of X , Y , Z it will only scan if the corresponding green button is pressed and the light is on. It scans from the lower value to the higher value in steps given by “Step.” If End - Start is not a multiple of Step, it just goes as far as it can, so if Start = 100, End = 129, and Step = 10 it will visit positions 100, 110, and 120. If Start = 129, End = 100, and Step = 10 it will swap the values (it always wants Start to be lower than End) and once again visit 100, 110, and 120. Be warned that it will always begin the scan by going to the lower value for each parameter, so be careful. If you have found a good spot at $X = 50 \mu\text{m}$ and want to stay there while scanning other parameters, make sure you set “X Start” equal to 50, and make sure that “X End” is bigger then 50. There is also a “User Defined” pane in which you can set specific (X, Y, Z) positions to go to. You fill out a $3 \times n$ array where the first row contains the X positions, the second row contains the Y positions, the third row contains the Z positions, and n is the total

number of spatial positions. The “Progress” values update to tell you where the stage is in its scan. Press “Scan” to begin the scan, but before doing so check the other settings (B14 and B15). Use the box in the upper-right to set how many total scans to perform.

- B13 This is used to auto-adjust the sample Z position in order to compensate for drift. It can be useful when doing experiments with a crossed-beam geometry. Press the “Auto Adjust Z ” button to make adjustments happen. The “Adjust Frequency” box tells it how many scans pass before doing an adjustment, and “Probe Pol” tells it what to set the probe polarizer to (something with a strong signal). It will set $t = 1$ ps and adjust sample stage Z to maximize the signal strength. I’ve never used this feature with a white light probe, and don’t know if it will work in that case.
- B14 Scan settings for delay stage. There are 5 options. “Linear:” set start and end times/positions and a step size, and it will linearly go between start and end. “Log:” set start and end times and a number of steps, and it will construct step logarithmically (linear steps from $\log(\text{start})$ to $\log(\text{end})$). “Log-Lin:” this is the most complicated but also the most useful. It has three regions, and relies on a good knowledge of time-0. First you tell it how many time delays to measure before time-0, and where to end ($t_{pre-end}$). It steps in 0.1 ps increments and ends at $t_{pre-end}$ (I usually set $t_{pre-end}$ to -1 ps). It then automatically goes to $t = 0$ ps. Next, you set start and end times for the linear portion, as well as the number of time delays to measure. It takes that many linearly spaced points between $t_{lin-start}$ and $t_{lin-end}$. Finally you set the end point of the logarithmic section, as well as the number of delays to measure. It logarithmically spaces that many points between $t_{lin-end}$ and $t_{log-end}$. “User Defined:” you manually enter time delays or positions. “File:” provide a file path to a data file containing a list of time delays.
- B15 Scan settings for polarization. For each polarization setting you’d like to scan at, manually enter pump polarization (not currently in use), probe polarization, and probe half-wave plate. Also enter the name and the number of scans. Note that even if you’re not scanning polarization you still need to put something in this array otherwise the scan won’t start. Just enter 0 for each value, and as long as you don’t press the green button, the rotation mounts won’t actually move from their current positions. A more common option is to use “Randomized Wprobe” in which the probe half-wave plate randomly explores several settings. You enter the probe polarizations you want (the actual half-wave plate setting will be one-half of these values), as well as the number of scans at each polarization and the number of rounds. One round visits each polarization in a random order.
- B16 Displays a spectral slice of scan data. The settings are fixed using the buttons below (see B18, B19, B20, B21).

- B17 Display a non-spectral slice of scan data. The settings are fixed using the buttons below (see B18, B19, B20, B21, B22). The independent variable is chosen with the tabs on top, and can be X , Y , Z , D , T , or a two-dimensional plane XY , XZ , YZ (I think XZ and YZ aren't working properly). If running a long scan, go to "None (Speedy)". This will bypass the code that determines what to plot, making the scan run faster. It would be nice to be able to use "Scan number" as the independent variable, but the current data storage structure makes that difficult.
- B18 Settings for what to plot. For each of X , Y , Z , scan number, and T you can either select a specific value using the arrows or press the square green button to average over all values.
- B19 Select whether to plot ΔOD or transmission.
- B20 Folder where data to plot is being stored. By default this will be the folder of the current scan, but if you are running a scan with multiple polarizations they are all saved in separate folders. Use this menu to look at data collected earlier in the scan. It would be nice to be able to plot polarization as the independent variable, but right now you are limited to selecting one polarization.
- B21 Press this button to enforce fixed limits on the plots. This is often useful so that you can click between different plots and compare without the axes changing, or look at ΔOD without the pump scatter spike taking up the whole screen.
- B22 Select a wavelength for the plot in B17. This wavelength is not selected directly, instead you move the cursor on plot B16 to select a wavelength. Use "Wavelength Width" to select how many pixels to average over. If it is 0, then data will be displayed for a single pixel. If it is 5, then the center pixel and 5 on either side will be averaged together.
- B23 Pause the scan, stop the scan, or exit the program. Stopping a scan can sometimes take a few seconds.
- B24 Press "Pump Scatter" if you have a pump scatter signal to set the timing with. Otherwise the camera doesn't know which pulses correspond to pump on and which to pump off, and your ΔOD signal can flip sign. Enter the wavelength of the pump scatter signal in the box, and the program will make sure that ΔOD at that wavelength is always negative. Note that you enter the wavelength corresponding to the pixel where the pump scatter is striking, not the actual physical wavelength of the pump.
- B25 Minimum and maximum wavelength for which to collect data. This is useful because only a fraction of the pixels are generally being used to measure the probe, and by throwing away the rest we save both processing time and memory.
- B26 Estimates the time remaining in the scan, if the green button is pressed. It stopped working correctly a few months ago, and I haven't looked in to why.

B27 Saving and copying scan data. For saving you can take either the plot, which is an image, or all of the data. The first time you press save you must navigate to the directory and provide a name. That name will actually be used to construct a file name. In the case of a full scan the various files will be put in directories automatically, and for saving a figure (either an image or the data from the plot B17 or the spectrum B16) a descriptive file name will be automatically constructed. If you then proceed to save more plots or spectra it will once again ask you for a name, but completely disregard what you say. All subsequent saves (until you run a new scan) will be saved together with descriptive file names. It would be nice if it didn't ask you for a prompt after the first save, but I never got that to work.

Wavelength Control:

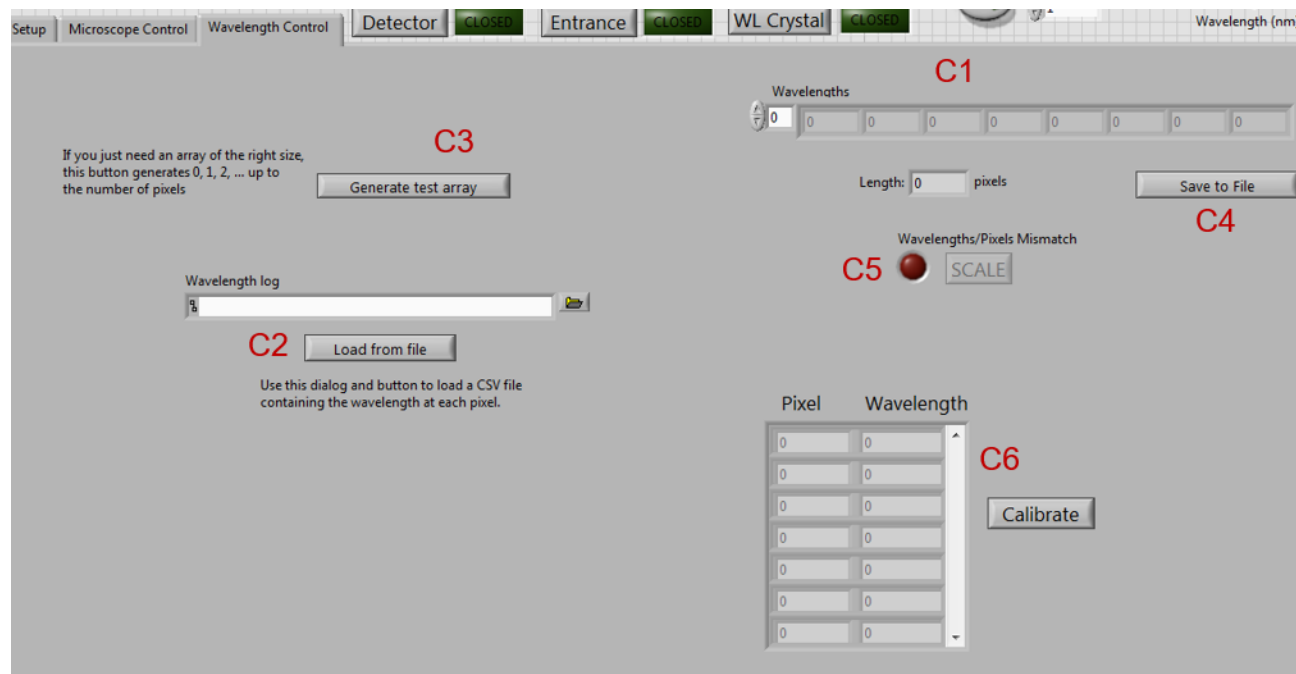


Figure B.5: “Wavelength Control” screen

- C1 Array of the wavelength (in nm) at each camera pixel. These numbers will be used for plotting data, saving data, and identifying pump scatter.
- C2 Load a wavelength file. Use the folder icon to select the file path and press “Load from file” to load.
- C3 Generate a basic array of wavelengths: 0, 1, 2, 3, 4, This is useful if you’re just getting things aligned and don’t care about the exact wavelength yet.
- C4 Save the current wavelength array to a file.
- C5 Lights up if there is a mismatch between the length of the wavelength array and the number of pixels on the camera. One situation where this can occur if if you have decided to do horizontal binning (which I don’t recommend). Pressing “Scale” will expand/contract the wavelength array to match the number of pixels, keeping the extremal values fixed.
- C6 Use several known wavelength positions to calibrate the wavelength array. In practice, you can rotate several bandpass filters into position in front of the camera. In the left column write down the pixel where the transmission is centered, and in the right column

write down the wavelength of light. After entering several such pairs press “Calibrate” and it will linearly interpolate to all pixels. In practice this can be challenging because the bandpass filters are broad and the white light spectrum is not-flat. It’s best to save all spectra (including a background spectrum) as you work, and then use that data later on to do a more precise calibration. However this quick procedure will make the spectrum accurate to within a few nanometers, and allows you to get a real-time sense for the data.

Appendix C

Code for Simulating Rubrene Crystallization

C.1 Main Simulation Code

The Jupyter notebook below includes the primary functions that run the Monte Carlo simulation, as well as cells for initializing and running a batch of simulations.

The easiest way to run the code is to install the Anaconda distribution, and then open Jupyter. Create a notebook file, as well as .py files in the same folder to hold the helper functions. All necessary Python packages are included in Anaconda, but if you are running Windows then there are problems with Cython. God help you if that's the case.

Cython is an extension of Python that allows you to turn Python code into compiled C modules. By compiling before run-time, you can achieve much faster execution—in this case I increased the speed by more than a factor of 10. However, the software to compile into C doesn't seem to play nicely on a Windows machine. I spent a day Googling and found many proposed methods to solve this problem. Ultimately I got it to work by installing the correct version of Microsoft Visual Studio, but exactly which version is “correct” depends on which version of Python you are running. Best of luck!

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3 import random
4 import matplotlib.colors as colors
5 import time
6 import os
```

```
1 %run Functions_Initialization.py
2 %run Functions_Visualization.py
3 %run Functions_Write.py
```

```
1 #Loads Cython
2 import setuptools
```



```
3 %reload_ext cython
```

Calling this functions runs an complete simulation for a given set of parameters. It initializes the simulation and then loops, running MC_Step() many times.

```
1 def run_MC_sim(out_path, side_length, nucleus_radius, activation_prob_min,
2   activation_prob_max, theta_crit, binding_energy, nucleus_scaling,
3   temperature=0, eps2=0):
4   """
5   run_MC_sim runs an entire Monte Carlo simulation and outputs the
6   results to files
7   out_path is a string containing the path where all files are saved
8   side_length is the number of cells on a side of the simulation area
9   nucleus_radius is the radius (in cells) of the nucleus
10  actiation_prob_min and activation_prob_max are are the minimum and
11  maximum probabilities of activating a neighbor (corresponding to an
12  orthogonal and parallel jump, respectively)
13  theta_crit is the critical angle beyond which there is no further
14  misorientation penalty
15  binding_energy is the binding energy of two perfectly aligned
16  crystalline neighbors
17  nucleus_scaling is the scaling factor (less than 1) that weakens bonds
18  with the nucleus
19  temperature is the dimensionless temperature, which determines the
20  ability of the system to move up in energy
21  eps2 is the coefficient of the quadratic term in the free energy
22  """
23  orientations, is_nucleus = Initialize_Nucleus(side_length,
24  nucleus_radius) # create nucleus and orientations matrices
25  orientations = np.mod(orientations, np.pi) # all orientations should
26  be mod pi
27
28  is_crystalline = np.copy(is_nucleus) # create matrix for crystalline
29  area
30  is_active, active_positions = Initialize_Active_Sites(is_nucleus) #
31  create matrix and list of active sides
32
33  # ensure out_path is a proper path, with the last character equal to "
34  /". If not, add "/"
35  if not out_path[-1] == "/":
36      out_path = out_path + "/"
37
38  # check to make sure path exists, and if not, create it
39  if not os.path.isdir(out_path):
40      os.makedirs(out_path)
41
42  # write the initial configuration
43  initial_config_fname = out_path + "initial_config.txt"
44  Write_Output(initial_config_fname, orientations, is_crystalline,
45  is_nucleus)
```

```

32     # write the log file
33     log_fname = out_path + "log.txt"
34     log_file = open(log_fname, "w")
35     log_file.write("side length: {0}\n".format(side_length))
36     log_file.write("nucleus radius: {0}\n".format(nucleus_radius))
37     log_file.write("critical angle: {0}\n".format(theta_crit))
38     log_file.write("binding energy: {0}\n".format(binding_energy))
39     log_file.write("nucleus scaling: {0}\n".format(nucleus_scaling))
40     log_file.write("activation probability, min: {0}\n".format(
activation_prob_min))
41     log_file.write("activation probability, max: {0}\n".format(
activation_prob_max))
42     log_file.write("Temperature (eV): {0}\n".format(temperature))
43     log_file.write("Epsilon 2: {0}\n".format(eps2))
44     log_file.close()
45
46     # create a blank file to hold each move
47     moves_fname = out_path + "moves.txt"
48     moves_file = open(moves_fname, "w")
49     moves_file.close()
50
51     # record the start time
52     start_time = time.time()
53
54     # figure out how many steps to run. We want the simulation to fill the
space, so calculate roughly how many steps that takes
55     # average probability of advancing to the next cell
56     activation_prob_mean = (activation_prob_min + activation_prob_max)/2
57     # distance from the edge of the nucleus to the edge of the simulation
area
58     distance = (side_length - 1)/2 - nucleus_radius
59     # divide those two to get the expected number of steps
60     num_steps = int(distance/activation_prob_mean/2)
61
62     # cycle through the number of steps
63     for i in range(num_steps):
64         moves_file = open(moves_fname, "a") # open the moves_file
65         moves_file.write("#{0}\n".format(i)) # write a header to indicate
a new step
66         # perform the step
67         (orientations, is_active, active_positions, is_crystalline) =
MC_Step(orientations, is_active, active_positions, is_crystalline,
is_nucleus, moves_file, activation_prob_min, activation_prob_max,
theta_crit, binding_energy, nucleus_scaling, temperature, eps2)
68         moves_file.close() # close the log file
69
70     # record the simulation run time in the log file
71     end_time = time.time()
72     log_file = open(log_fname, "a")
73     log_file.write("Simulation time (seconds): {0:.1f}\n".format(end_time

```

```

- start_time))
74     log_file.close()
75
76     # write the final configuration to a file
77     final_config_fname = out_path + "final_config.txt"
78     Write_Output(final_config_fname, orientations, is_crystalline,
is_nucleus)
79
80     # show the final configuration
81     Show_Orientations(orientations, is_crystalline, is_nucleus)

```

```

1  def MC_Step(orientations, is_active, active_positions, is_crystalline,
is_nucleus, file, activation_prob_min, activation_prob_max, theta_crit,
binding_energy, nucleus_scaling, temperature=0, eps2=0):
2      """
3      Perform a Monte Carlo step on the current configuration
4      The number of moves is randomly chosen as a Poisson process with an
average equal to the number of active sites.
5      For each move a random pixel is chosen and a move is made, which may
or may not be kept
6      Amorphous regions may become activated if they are next to a
crystalline region
7      orientaions: array of the orientation of each cell
8      is_active: boolean array of whether or not a cell is active
9      active_positions: list of all active positions, so they can be more
easily iterated through
10     is_crystalline: boolean array of whether or not a cell is crystalline
11     file: file containing a log of all moves
12     activation_prob_min: minimum probability of a crystalline cell
activating a neighbor (when the propagation direction is orthogonal to
the orientation)
13     activation_prob_max: maximum probability of a crystalline cell
activating a neighbor (when the propagation direction is parallel to
the orientation)
14     theta_crit: critical angle beyond which there is no penalty for
further misorientation
15     binding_energy: energetic reward when two adjacent crystalline cels
have the same orientation
16     nucleus_scaling: scaling factor (less than 1) that weakens bonds with
the nucleus
17     """
18
19     dTheta = temperature/binding_energy # amount by which orientations can
fluctuate
20     num_active = len(active_positions) # number of active positions at the
start of the step.
21     num_moves = np.random.poisson(num_active) # randomly select the number
of moves to do
22
23     for i in range(num_moves):

```

```

24     # select a random active site and get its position
25     index = np.random.randint(num_active)
26     pos = active_positions[index]
27
28     # Record the orientation and energy of this position
29     old_orientation = orientations[pos]
30     if not is_crystalline[pos]: # If this position is currently
amorphous, energy is 0
31         old_energy = 0
32     else:
33         old_energy = Get_Energy(pos, orientations, is_crystalline,
is_nucleus, binding_energy, theta_crit, nucleus_scaling, eps2)
34
35     # generate a new orientation and calculate the new energy
36     new_orientation = Generate_Orientation_Matching_Cython(pos,
orientations, is_crystalline) # Generate a new orientation
37     new_orientation = Thermalize_Orientation(new_orientation, dTheta)
38     orientations[pos] = new_orientation # Set the orientation of the
cell to new_orientation
39     new_energy = Get_Energy(pos, orientations, is_crystalline,
is_nucleus, binding_energy, theta_crit, nucleus_scaling, eps2) #
Calculate the new energy
40
41     # Decide whether or not to accept the step
42     if Accept_Step_Cython(old_energy, new_energy):
43         # Write the new position to a file, unless the cell was
already crystalline and its orientation has not changed
44         if not (is_crystalline[pos] and new_orientation ==
old_orientation):
45             is_crystalline[pos] = True
46             file.write("{0}, {1}, {2}\n".format(pos[0], pos[1],
new_orientation))
47         else:
48             # if not accepting the step, revert to old orientation
49             orientations[pos] = old_orientation # reset the orientation
50
51     # activate nearby sites
52     is_active, active_positions = Activate_Sites(pos, new_orientation,
is_active, active_positions, is_nucleus, activation_prob_min,
activation_prob_max)
53
54     return orientations, is_active, active_positions, is_crystalline

```

```

1 def Thermalize_Orientation(theta, dTheta):
2     # Thermally scrambles an orientation, theta, by adding a random number
between -dTheta and +dTheta
3     #theta = theta + np.random.uniform(low = -dTheta, high = dTheta)
4     theta = theta + np.random.normal(scale=dTheta)
5     return theta%np.pi

```

```
1 %%cython -a
2 import cython
3 import numpy as np
4 cimport numpy as np #cimport allows the C-api to access numpy functions
5 cimport cython # so we can use cython decorators
6 from cpython cimport bool # type annotation for boolean
7 from libc.math cimport sin, cos, atan2 #trig functions in C
8
9 #This code provides only a modest speed-up, going from 12.6 us to generate
   an orientation using pure python
10 #to 9.6 us with Cython. I believe the basic problem is that I'm still
   reliant on numpy arrays and did not
11 #transition over to C data types (such as C arrays). The overhead in
   putting data into numpy arrays mostly
12 #cancels out the benefits of doing the arithmetic in C and using static
   typing.
13 #I had a lot of trouble, even just using memory views. One big problem is
   that the code works by whittling down an array
14 #First it gets all of the neighbors, then it selects the orientations of
   those that are crystalline, then it generates
15 #a random subset. C requires you to know the size of an array before
   initializing it.
16 #I tried to figure out the size beforehand and then assign that size to a
   variable and use that variable to set the
17 #size of the array, but that created type errors. Using constants created
   type errors. Using memory views created
18 #type errors. Using lists and trying to dynamically assign values crashed
   the kernel.
19 #Basically, making this faster requires more than just some static typing.
   I would have to have a much better grasp
20 #of C as a language, and think about things in a different way.
21
22 # disable index bounds checking and negative indexing for speedups
23 @cython.wraparound(False)
24 @cython.boundscheck(False)
25 def Generate_Orientation_Matching_Cython((int,int) pos, np.ndarray[np.
   float64_t, ndim=2] orientations, np.ndarray[np.int32_t, ndim=2]
   is_crystalline):
26     #Generate_Orientation_matching generates a plausible orientation for
   the space given by pos
27     #A plausible orientation is one that shares some alignment with the
   spaces around it
28
29     #break out orientation.shape and cast it into a tuple (orientations.
   shape is actually a C array, so I have to)
30     #pluck out each element individually, otherwise there's a type
   mismatch
31     cdef (int, int) grid_size = (orientations.shape[0], orientations.shape
   [1])
```

```

32     cdef list neighbors = Get_Neighbor_Positions_Cython(pos, grid_size) #
    get list of neighbors
33
34     #select out the orientations of those neighbors that are crystalline
35     cdef list neighbor_orientations_list = [orientations[neighbor] for
neighbor in neighbors if is_crystalline[neighbor]]
36
37     #convert the above list to a numpy array, so that I can later make
calls like neighbor_orientations[some_boolean_array]
38     cdef np.ndarray[np.float64_t, ndim=1] neighbor_orientations = np.array
(neighbor_orientations_list)
39
40     cdef Py_ssize_t num_neighbors = neighbor_orientations.shape[0] #
number of crystalline neighbors
41     # if there's only one crystalline neighbor, adopt that orientation
42     if num_neighbors == 1:
43         return neighbor_orientations[0]
44
45     #create empty boolean numpy array of length num_neighbors
46     cdef np.ndarray[np.uint8_t, cast=True] align_neighbors = np.empty(
num_neighbors, dtype=np.uint8).astype(bool)
47     cdef Py_ssize_t i = 0 #iterator variable
48
49     #generate num_neighbors random numbers (either 0 or 1) to decide which
neighbor orientations to consider
50     align_neighbors = np.random.randint(2, size = num_neighbors).astype(
bool)
51     return Get_Angle_Mean(neighbor_orientations[align_neighbors]) #Get
average of this subset of crystalline neighbor orientations
52
53 @cython.wraparound(False)
54 @cython.boundscheck(False)
55 def Get_Neighbor_Positions_Cython((int, int) position, (int, int)
grid_size):
56     #Given an initial position (tuple and y and x coordinates) returns a
list of the positions of the 4 neighbors
57     #x_length and y_length are the side lengths of the simulation space (
usually the same, but don't want to assume)
58     #This is useful because we want to enforce periodice boundary
conditions, so everything is taken modulo these numbers
59     cdef int x, y, x_length, y_length
60     (y, x) = position #break position out in to y and x positions
61     (y_length, x_length) = grid_size
62
63     cdef (int, int) left_pos = (y%y_length, (x-1)%x_length) #left
positions is (y, x-1), but have to take everything modulo y_length &
x_length
64     cdef (int, int) right_pos = (y%y_length, (x+1)%x_length) #right pos
65     cdef (int, int) up_pos = ((y-1)%y_length, x%x_length) #up pos
66     cdef (int, int) down_pos = ((y+1)%y_length, x%x_length) #down pos

```

```

67     cdef list neighbor_pos = [up_pos, down_pos, left_pos, right_pos]
68     return neighbor_pos
69
70
71 @cython.wraparound(False)
72 @cython.boundscheck(False)
73 def Get_Angle_Mean(np.ndarray[np.float64_t, ndim=1] theta_list):
74     # Get_Angle_Mean averages a numpy array of angles, theta_list
75     # This is difficult because you have to deal with circular
76     discontinuities, so instead of doing an arithmetic mean,
77     # convert each angle to a unit vector, average the cartesian positions
78     of those vectors, and take the angle of the result
79     # Furthermore, we have to take two-fold symmetry into account, because
80     181 degrees = 1 degree
81     # So we modulo all angles to [0, pi), then double all angles, then use
82     the averaging procedure,
83     # then take that result modulo [0, 2pi), then halve the result
84
85     cdef Py_ssize_t num_theta = theta_list.shape[0] #number of angles
86
87     #if no angles were passed, then just generate a random number from 0
88     to pi
89     if num_theta == 0:
90         return np.random.random()*3.14159
91
92     #if one angle was passed, align with it
93     if num_theta == 1:
94         return theta_list[0]
95
96     cdef Py_ssize_t i = 0 #iterates through angles
97     cdef double theta = 0.0 #variable to hold angle
98     cdef double x = 0.0 #x position of vector sum
99     cdef double y = 0.0 #y position of vector sum
100
101     for i in range(num_theta): #iterate through angles
102         theta = (theta_list[i]%3.14159)*2 #take mod pi and multiply by 2
103         #add sine and cosine to y and x variables
104         y += sin(theta)
105         x += cos(theta)
106
107     cdef double mean_theta = atan2(y, x) #find angle of sum vector
108     mean_theta = mean_theta%6.28318 #take mod 2*pi
109     return mean_theta/2 #divide by 2 and return

```

```

1 def Get_Energy(pos, orientations, is_crystalline, is_nucleus,
2 binding_energy, theta_crit, nucleus_scaling, eps2):
3     #Get_Energy figures out the energy of the molecules at a certain
4     position
5     #pos is the coordinates of the space under consideration
6     energy = 0

```

```

5  neighbors = Get_Neighbor_Positions_Cython(pos, orientations.shape) #
   get the four neighbors of pos
6
7  for neighbor in neighbors: #cycle through the neighbors
8      if is_crystalline[neighbor]: #if the neighbor is crystalline,
   calculate the interaction energy and add to energy
9          energy += Calc_Energy(orientations[pos], orientations[neighbor
   ], is_nucleus[neighbor]*1, binding_energy, theta_crit, nucleus_scaling,
   eps2)
10
11  return energy

```

```

1  %%cython -a
2  import cython
3
4  #using cython reduced run-time from 160 ns to 95 ns, which is nice,
5  #but getting neighbor direction was not a significant bottleneck
6
7  def Get_Direction_Neighbor_Cython((int, int) pos1, (int, int) pos2):
8      # Get_Direction_Neighbor returns the angle of the vector from pos1 to
   pos2
9      # pos1 and pos2 are tuples (y,x)
10     # It is assumed that pos1 and pos2 are neighbors
11
12     cdef int x1, y1, x2, y2
13
14     #split out into x and y coordinates
15     (y1, x1) = pos1
16     (y2, x2) = pos2
17
18     #calculate the difference in x
19     cdef int dx = x2 - x1
20
21     if dx == 0:
22         #if dx = 0 then this is a vertical move, and the direction is pi/2
23         return 1.571
24     return 0 #otherwise it's a horizontal move, and the direction is 0

```

This is where Cython provides the biggest speed-up.

```

1  %%cython -a
2  import cython
3  from libc.math cimport cos #cosine function in C
4  import random
5
6  @cython.cdivision(True) #hypothetically saves some time by not checking
   for division by 0 errors
7  def Calc_Energy(float orientation1, float orientation2, int nucleus, float
   binding_energy, float theta_crit, float nucleus_scaling, float eps2):
8      # Calc_Energy calculates the energy between two adjacent spaces

```



```

9      # orientation1 and orientation2 are the orientation angles of the two
      spaces
10     # nucleus is true if one of the spaces is part of the nucleus
11
12     # Calculate the minimum angle between the two orientations.
13     cdef float delta_theta
14     delta_theta = Get_Angle_Difference(orientation1, orientation2) #
      delta_theta should be between 0 and pi/2
15
16     cdef float energy
17
18     #Energy function: -binding_energy for oriented interfaces,
      quadratically increasing to 0 at theta_crit, then 0 for larger angles
19     if delta_theta > theta_crit:
20         return 0
21     energy = (binding_energy/theta_crit - eps2*theta_crit)*delta_theta -
      binding_energy + eps2*delta_theta**2
22
23     # option to make attraction weaker if it is with nucleus
24     if nucleus == 1:
25         energy = energy*nucleus_scaling
26
27     return energy
28
29 def Get_Angle_Difference(float theta1, float theta2):
30     #Get_Angle_Difference returns the smallest difference between angles
      theta1 and theta2
31     #Need to consider going both ways around the circle (theta1 - theta2
      and theta2 - theta1)
32     #and then take that difference modulo pi. The lower number is the
      actual minimum angle
33     #The smallest difference is always between 0 and pi/2
34
35     cdef float dtheta1, dtheta2
36     cdef float pi = 3.14159
37
38     dtheta1 = (theta1 - theta2)%pi
39     dtheta2 = (theta2 - theta1)%pi
40
41     return min(dtheta1, dtheta2)
42
43 def Decide_Activate_Neighbor_Cython(float theta, float neighbor_direction,
      float min_prob, float max_prob):
44     # Decide_Activate_Neighbor decides whether or not a neighboring site
      is to become active (eligible for crystallization)
45     # theta is the orientation of the initial site
46     # neighbor_direction is the angle from the initial site to its
      neighbor
47     # min_prob and max_prob are the minimum and maximum probabilities of
      activating a neighbor (corresponding to an orthogonal and parallel jump

```

```

, respectively)
48
49     cdef float d_theta = Get_Angle_Difference(theta, neighbor_direction)
50
51     cdef float threshold = min_prob + (max_prob - min_prob)*cos(d_theta)**
52     2
53     return random.random() < threshold

```

```

1  %%cython -a
2  import cython
3  from libc.math cimport exp #exponential function in C
4  from libc.stdlib cimport rand, RAND_MAX #random numbers in C
5
6  @cython.cdivision(True) #hypothetically saves some time by not checking
   for division by 0 errors
7  def Accept_Step_Cython(float old_energy, float new_energy, float
   temperature=0):
8
9
10     if new_energy <= old_energy:
11         return True
12
13     if temperature == 0:
14         return False
15
16     cdef float acceptance_prob = exp((old_energy - new_energy)/temperature
   )
17     cdef float random_num = rand()/(RAND_MAX*1.0)
18
19     return random_num < acceptance_prob

```

```

1  def Activate_Sites(pos, orientation, is_active, active_positions,
   is_nucleus, activation_prob_min, activation_prob_max):
2     # Activate_Sites around a newly crystallizes site
3     # pos is the position of the newly crystallized site
4     # activation_prob_min and activation_prob_max are the minimum and
   maximum probabilities of activating a neighbor (corresponding to an
   orthogonal and parallel jump, respectively)
5
6     neighbors = Get_Neighbor_Positions_Cython(pos, is_active.shape) #get
   the four neighbors of pos
7     for neighbor in neighbors:
8         if not is_active[neighbor] and not is_nucleus[neighbor]:
9             neighbor_direction = Get_Direction_Neighbor_Cython(pos,
   neighbor)
10            if Decide_Activate_Neighbor_Cython(orientation,
   neighbor_direction, activation_prob_min, activation_prob_max):
11                is_active[neighbor] = True
12                active_positions.append(neighbor)
13

```

```
14     return is_active, active_positions
```

The following cell demonstrates how to run the code. It fixes most parameters, makes arrays of temperature and ϵ_2 , and loops through parameter space, performing each simulation and saving the results in a different folder.

```
1 side_length = 151
2 nucleus_radius = 15
3 binding_energy = 50
4 activation_rate = 0.01
5 theta_crit = 0.1
6 nucleus_scaling = 0.3
7 temperature_array = [0.03, 0.05, 0.08]
8 eps2_array = [10, 100, 1000]
9 for eps2 in eps2_array:
10     for temperature in temperature_array:
11         path = "eps2 {0} temp {1}/".format(str(eps2).replace('.', '_'),
12         str(temperature).replace('.', '_'))
13         out_path = "Monte Carlo Results/test folder/" + path
14         run_MC_sim(out_path, side_length, nucleus_radius, activation_rate,
15         activation_rate, theta_crit, binding_energy, nucleus_scaling,
16         temperature, eps2)
```

C.2 Helper Functions

Functions Initialization.py

```
1 import numpy as np
2 import random
3
4 def Initialize_Nucleus(side_length, nucleus_radius):
5     #Initialize_Nucleus creates 2D arrays holding the position of
6     the nucleus and the orientation of each space
7     center_pos = np.floor(side_length/2) #find the center position
8     of the side length
9     x_vec = np.arange(side_length) - center_pos #create x and y
10    arrays that go from roughly -side_length/2 to +side_length/2
11    y_vec = np.arange(side_length) - center_pos
12
13    X, Y = np.meshgrid(x_vec, y_vec)
14    Theta = np.arctan2(-Y, X) #calculate the angle of a line from
15    each space to the center point
16    Radius = np.sqrt(np.power(X, 2) + np.power(Y, 2)) #calculate
17    distance from the center point
```

```
14     is_nucleus = (Radius <= nucleus_radius)*1 #Points that are a
15     distance less than nucleus_radius are part of the nucleus
16     orientations = Theta*is_nucleus #orientations are given by Theta
17     , but only for spaces that are part of the nucleus
18     #(everything else is amorphous
19     for now)
20
21     return orientations, is_nucleus
22
23 def Initialize_Active_Sites(is_nucleus):
24     #Initialize_Active_Sites finds those positions that are active -
25     - that is, they can be addressed by the Monte Carlo sim
26     #Addressable spaces are those that are not in the nucleus and
27     are either crystalline or adjacent to something crystalline
28     #At this point nothing outside the nucleus is crystalline, so we
29     just find all amorphous spaces that border the nucleus
30     active_positions = [] #This will be a list of tuples holding
31     positions of active spaces
32     (y_len, x_len) = is_nucleus.shape #Get the shape of the
33     simulation space
34     is_active = np.full(is_nucleus.shape, False) #Initialize an
35     array of False to hold which spaces are active
36
37     for y in np.arange(y_len): #Cycle through the grid
38         for x in np.arange(x_len):
39             pos = (y,x) #Pack y and x into a position tuple
40             if not is_nucleus[pos] and np.any([is_nucleus[pos] for
41             pos in Get_Neighbor_Positions(pos, is_nucleus.shape)]):
42                 #This if statement determines whether or not a space
43                 is to be declared "active"
44                 #The first part checks to see if it's part of the
45                 nucleus -- if it is, then this space is disqualified
46                 #Next we get a list of the positions of all
47                 neighbors and check to see if they are in the nucleus
48                 #If that is true for any of the neighbors, then this
49                 space is declared active
50                 active_positions.append(pos)
51                 is_active[pos] = True
52
53     return is_active, active_positions
54
55 def Get_Neighbor_Positions(position, grid_size):
56     #Given an initial position (tuple and y and x coordinates)
57     returns a list of the positions of the 4 neighbors
```

```

43     #x_length and y_length are the side lengths of the simulation
    space (usually the same, but don't want to assume)
44     #This is useful because we want to enforce periodic boundary
    conditions, so everything is taken modulo these numbers
45     (y, x) = position #break position out in to y and x positions
46     (y_length, x_length) = grid_size
47
48     left_pos = (y%y_length, (x-1)%x_length) #left positions is (y, x
    -1), but have to take everything modulo y_length & x_length
49     right_pos = (y%y_length, (x+1)%x_length)
50     up_pos = ((y-1)%y_length, x%x_length)
51     down_pos = ((y+1)%y_length, x%x_length)
52
53     return [up_pos, down_pos, left_pos, right_pos] #package the four
    positions into a list

```

Functions_Visualization.py

```

1  import matplotlib.pyplot as plt
2  import matplotlib.colors as colors
3  import numpy as np
4
5  def Show_Orientations(orientations, is_crystalline, is_nucleus):
6      #Show_Orientations visualizes the orientation map of the crystal
7      vis_array = np.copy(orientations) #make a copy of the
    orientations_array so we can mess with it
8      vis_array = np.mod(vis_array, np.pi) #Take everything mod pi
9
10     vis_array = vis_array - 5*is_nucleus #Everything that's the
    nucleus is made negative
11     vis_array = vis_array + 5*(1 - is_crystalline) #Everything that's
    s amorphous is made bigger than pi
12
13     map_name = 'hsv' #This is the map to use (hsv is good because it
    cycles)
14     num_segments = 50 #How finely to discretize it
15     new_cmap = Make_Saturated_Cmap(map_name, num_segments) #Add
    saturation values to the ends of this color map
16
17     #Want the first value of the color map (black) to be just below
    0, the second value (red) to be zero, the next to last value (red
    )
18     #to be pi, and the last value (white) to be just above pi. This
    can be achieved with the interval pi*[-1/num, 1+1/num)
19     minval = -np.pi/num_segments

```

```

20     maxval = np.pi*(1 + 1/num_segments)
21     fig = plt.imshow(vis_array, cmap=new_cmap, vmin=minval, vmax=
maxval)
22     plt.colorbar()
23     plt.show()
24
25 def Make_Saturated_Cmap(map_name, num_segments):
26     #Make_Saturated_Cmap constructs a color map that is normal, but
has fixed colors to show negative and positive saturation
27     #For example I might take the hsv color map, which both starts
and ends on red, and then add black at the top and white on the
bottom
28     #This way I can specific scale limits for my figure, and
everything within those limits will be colored using hsv while
29     #everything outside will be black or white
30     #map_name is the name of the starting map (as a string)
31     #num_segments is how finely to discretize the map
32     cmap = plt.get_cmap(map_name) #get the color map
33     cmap_array = cmap(np.linspace(0, 1, num_segments)) #Turn it into
an array of num_segments colors
34
35     low_color = [0, 0, 0, 1] #define the low color (black)
36     high_color = [1, 1, 1, 1] #define the high color (white)
37     new_cmap_array = np.vstack((low_color, cmap_array, high_color))
#sandwich the existing array between the low and high colors
38
39     new_cmap = colors.LinearSegmentedColormap.from_list('my_map',
new_cmap_array) #use this new array to make a new colormap
40     return new_cmap
41
42 def Linecut_angular(orientations, r, thickness):
43     # Linecut_angular creates a circular linecut of the orientation
map around the center of the simulation area
44     # r is the radius of the linecut, thickness is the thickness (in
pixels) of the region to average
45
46     if thickness < 1:
47         thickness = 1
48
49     r = np.abs(r)
50     if r < thickness:
51         r = thickness
52
53     (side_y, side_x) = orientations.shape

```

```

54     # find the center positions of the sides
55     center_x = np.floor(side_x/2)
56     center_y = np.floor(side_y/2)
57     # create x and y arrays that go from roughly -side_length/2 to +
side_length/2
58     x_vec = np.arange(side_x) - center_x
59     y_vec = np.arange(side_y) - center_y
60
61     X, Y = np.meshgrid(x_vec, y_vec)
62     Theta = np.arctan2(-Y, X) #calculate the angle of a line from
each space to the center point
63     Theta = np.mod(Theta, 2*np.pi)
64     Radius = np.sqrt(np.power(X, 2) + np.power(Y, 2)) #calculate
distance from the center point
65
66     num_points = int(2*np.pi*r*2/3)
67     theta_array = np.linspace(0, 2*np.pi, num_points)
68     d_theta = 2*np.pi/num_points
69     crystal_angle = np.zeros(num_points)
70
71     for i, theta_i in enumerate(theta_array):
72         mask1 = np.abs(Radius - r) <= thickness/2
73         mask2 = np.mod(Theta - theta_i, 2*np.pi) <= d_theta
74         mask = np.logical_and(mask1, mask2)
75         crystal_angle[i] = np.sum(orientations[mask])/np.sum(mask*1)
76
77     fig = plt.plot(theta_array, crystal_angle)
78     plt.xlabel('Position angle')
79     plt.ylabel('Crystal angle')
80     plt.show()

```

Functions_Write.py

```

1 import numpy as np
2
3 def Write_Output(fname, orientations, is_crystalline, is_nucleus):
4     """
5     Write_Output writes the state of the system to an output file
6     fname is the file name
7     orientations is a numpy array of the orientation at each point
8     is_crystalline is a boolean array indicating whether or not a
specific point is crystalline
9     is_nucleus is a boolean array indicating whether or not a
specific point is part of the nucleus
10    """

```

```
11
12     full_array = np.copy(orientations) #make a copy of the
orientations_array so we can mess with it
13     full_array = np.mod(full_array, np.pi) #Take everything mod pi
14
15     full_array = full_array - 5*is_nucleus #Everything that's the
nucleus is made negative
16     full_array = full_array + 5*(1 - is_crystalline) #Everything
that's amorphous is made bigger than pi
17
18     np.savetxt(fname, full_array, fmt='%.3f')
```


Appendix D

Code for Identifying Rubrene Morphology

D.1 Jupyter Notebooks

Jupyter notebooks to train and apply the model. The notebook files must be in the same directory as folders containing the images. All helper functions can be found in the next Section.

Training the Model

The following cells train a random forest model to classify pixels in crossed polarizer images of rubrene thin films.

```
1 %run Functions_Regions.py
2 %run Functions_Feature_Extraction.py
3 %run Functions_Visualization.py
4 %run Functions_Training.py
```

```
1 from skimage import io
2 from skimage import transform
3 import matplotlib.pyplot as plt
4 import random
5 import pickle
6 import numpy as np
```

```
1 #file paths to images to be used for training data set
2 train_file_paths = []
3 train_file_names = []
4 train_file_paths.append("JD10/"); train_file_names.append("Y1_5 X12_0.bmp"
5     )
6 ...
7 ...
```

```

7 #SEE SOURCE FOR COMPLETE LIST OF TRAINING IMAGES

1 #Training data. Each data set is a region of an image and a classification
  (amorphous, spherulite, etc.)
2 #Each data set requires six numbers -- the min and max x and y of the
  region, the class, and the image number
3 #(number is given by the order in which they are imported, in the above
  cell)
4 #Classes are amorphous (1), platelet (2), spherulite (3), and junk (4)
5 train_ymin = []
6 train_ymax = []
7 train_xmin = []
8 train_xmax = []
9 train_image_num = []
10 train_class = []
11
12 #JD10/Y1_5 X12_0.bmp
13 train_ymin.append(0); train_ymax.append(75)
14 train_xmin.append(75); train_xmax.append(159)
15 train_image_num.append(0); train_class.append(1)
16
17 ...
18 ...
19 #SEE SOURCE FOR COMPLETE TRAINING DATA
20
21 train_ymin = np.asarray(train_ymin, dtype=int)
22 train_ymax = np.asarray(train_ymax, dtype=int)
23 train_xmin = np.asarray(train_xmin, dtype=int)
24 train_xmax = np.asarray(train_xmax, dtype=int)
25 train_image_num = np.asarray(train_image_num, dtype=int)
26 train_class = np.asarray(train_class, dtype=int)
27
28 train_data = (train_ymin, train_ymax, train_xmin, train_xmax,
  train_image_num, train_class)

1 #import images
2 images = []
3 for i, (path, name) in enumerate(zip(train_file_paths, train_file_names)):
4     img = io.imread(path + name).astype(int)
5     background = io.imread(path + "Background.bmp").astype(int)
6     images.append(NormalizeImage(img, background))

1 #find the regions
2 labels_list = FormRegions(images)

1 #Extract all of the features for each pixel of each image
2 #the "local" region is a square of length 2d+1
3 d = 4
4 X_all = MakeFeaturesMatrix(images, labels_list, d)

```

```

5 classifications = WriteTrainingClassification(images, train_data)
6
7 #build the random forest classifier with 80 trees
8 #(chosen by comparing several values)
9 n_trees = 80
10 classifier = TrainForest(X_all, classifications, n_trees)

```

```

1 #Use model to predict classes of all pixels in training images
2 #Plot predictions overlaid on images
3 Y_pred = classifier.predict(X_all)
4 images_predicted_class = np.reshape(Y_pred, classifications.shape)
5 plt.figure()
6 ShowAllClassification(images, images_predicted_class)

```

```

1 #Save model to file
2 filename = 'RF models/full_rf.sav'
3 pickle.dump(classifier, open(filename, 'wb'))

```

```

1 #constants used to train lots of models on subsets of training data
2 num_train_regions = len(train_data[0])
3 num_subsample_regions = 26
4 num_forests = 200

```

```

1 #train lots of models, each on a randomly generated data subset
2 indices = np.arange(num_train_regions) #create an array from 0 to
   num_train_regions -1. This will be used to select subsets
3 # loop through the number of forests
4 for i in range(num_forests):
5     np.random.shuffle(indices) # shuffle the indices and take the first
   num_subsample_regions
6     train_indices = indices[0:num_subsample_regions]
7     remaining_data, train_data_subset = SplitTrainTest(train_data,
   train_indices) # split into training data and remaining, unused data
8     train_matrix = WriteTrainingClassification(images, train_data_subset)
   #create training data matrix
9     classifier = TrainForest(X_all, train_matrix, n_trees) # train random
   forest classifier
10    filename = "RF models/model{0}.sav".format(i)
11    pickle.dump(classifier, open(filename, 'wb')) # write model to a file

```

```

1 #load the full model and extract the relative importance of each feature
2 filename = 'RF models/full_rf.sav'
3 classifier_full = pickle.load(open(filename, 'rb'))
4 features_imp = classifier_full.feature_importances_
5 num_features = len(features_imp)
6 print("The relative importance of the features are: {0}".format(
   features_imp))

```

Applying the Model

The following cells apply the model to the images of a specific film (in this case, labelled “JD #10”).

```
1 %run Functions_Regions.py
2 %run Functions_Feature_Extraction.py
3 %run Functions_Visualization.py
4 %run Functions_Training.py
```

```
1 import matplotlib.pyplot as plt
2 import pickle
3 import glob
```

```
1 Path = "JD10/"
2 backgroundName = "Background.bmp"
3 background = io.imread(Path + backgroundName).astype(int)
4 n_rows = 256
5 n_cols = 320
```

```
1 #ITO
2 file_paths = glob.glob(Path + "Y5_5*.bmp") #import files at Y = 5.5
3 n_images_ITO = len(file_paths)
4 images = []
5 for file_path in file_paths:
6     image_this = io.imread(file_path).astype(int)
7     images.append(NormalizeImage(image_this, background))
8
9 images_ITO = images
```

```
1 #glass
2 file_paths = glob.glob(Path + "Y1_5*.bmp") #import files at Y = 1.5
3 n_images_glass = len(file_paths)
4 images = []
5 for file_path in file_paths:
6     image_this = io.imread(file_path).astype(int)
7     images.append(NormalizeImage(image_this, background))
8
9 images_glass = images
```

Note: to save space, only the code for analyzing images above ITO is reproduced below.

```
1 #find the regions
2 labels_list_ITO = FormRegions(images_ITO)
```

```
1 #Calculate features at each pixel
2 d = 4
3 X_all_ITO = MakeFeaturesMatrix(images_ITO, labels_list_ITO, d)
```

```

1 #Load and apply the model
2 filename = 'RF models/full_rf.sav'
3 classifier = pickle.load(open(filename, 'rb'))
4 Y_pred_ITO = classifier.predict(X_all_ITO)

```

```

1 #Visualize the results
2 images_pred_ITO = np.reshape(Y_pred_ITO, (n_rows, n_cols, n_images_ITO))
3 plt.figure()
4 ShowAllClassification(images_ITO, images_pred_ITO)

```

```

1 # calculate proportion of each type: amorphous, platelet, spherulite, and
   junk (for ITO)
2 total_pix = len(Y_pred_ITO)
3 p_amorphous = 1 - np.count_nonzero(Y_pred_ITO - 1)/total_pix
4 p_platelet = 1 - np.count_nonzero(Y_pred_ITO - 2)/total_pix
5 p_spherulite = 1 - np.count_nonzero(Y_pred_ITO - 3)/total_pix
6 p_junk = 1 - np.count_nonzero(Y_pred_ITO - 4)/total_pix
7 print(f'Probability of amorphous: {p_amorphous}')
8 print(f'Probability of platelet: {p_platelet}')
9 print(f'Probability of spherulite: {p_spherulite}')
10 print(f'Probability of junk: {p_junk}')

```

```

1 #Apply each of the models that were trained on a subset of training data
2 num_models = 200
3 # create arrays to hold the fraction of each type (amorphous, platelet,
   spherulite, junk)
4 # for each of the models
5 fraction_ITO = np.zeros((4, num_models))
6 fraction_glass = np.zeros((4, num_models))
7 for i in range(num_models):
8     filename = "RF models/model{0}.sav".format(i)
9     classifier = pickle.load(open(filename, 'rb')) # load the model
10    Y_pred_ITO = classifier.predict(X_all_ITO) # predict for ITO
11    Y_pred_glass = classifier.predict(X_all_glass) # predict for glass
12    #add up the results and calculate the fraction of each type (for glass
   )
13    total_pix = len(Y_pred_ITO)
14    fraction_ITO[0,i] = 1 - np.count_nonzero(Y_pred_ITO - 1)/total_pix
15    fraction_ITO[1,i] = 1 - np.count_nonzero(Y_pred_ITO - 2)/total_pix
16    fraction_ITO[2,i] = 1 - np.count_nonzero(Y_pred_ITO - 3)/total_pix
17    fraction_ITO[3,i] = 1 - np.count_nonzero(Y_pred_ITO - 4)/total_pix

```

```

1 #Save results
2 filename = Path + "RF results/results_ITO.txt"
3 np.savetxt(filename, fraction_ITO.T, fmt='%.6f', delimiter=',')

```

```

1 #Visualize results with histograms
2 f, axs = plt.subplots(2,2, figsize = (8,8))

```

```

3 axs[0,0].hist(fraction_ITO[0,:], bins=np.linspace(0,1,51))
4 axs[0,0].set_title("amorphous")
5 axs[0,1].hist(fraction_ITO[1,:], bins=np.linspace(0,1,51))
6 axs[0,1].set_title("platelet")
7 axs[1,0].hist(fraction_ITO[2,:], bins=np.linspace(0,1,51))
8 axs[1,0].set_title("spherulite")
9 axs[1,1].hist(fraction_ITO[3,:], bins=np.linspace(0,1,51))
10 axs[1,1].set_title("junk")

```

D.2 Helper Functions

All files must be in the same directory as the Jupyter notebooks, above.

Functions_Regions.py

The following functions calculate and display the clustered regions for an image.

```

1 import numpy as np
2 import scipy as sp
3 from scipy import ndimage as ndi
4 from skimage import transform
5 from skimage import io
6 from skimage import color
7 from skimage import morphology
8 import matplotlib.pyplot as plt
9 from sklearn.feature_extraction import image
10 from sklearn.cluster import spectral_clustering
11 from sklearn.feature_extraction.image import grid_to_graph
12 from sklearn.cluster import AgglomerativeClustering
13
14 def FormRegions(images):
15     #Cycles through an array of images, and for each one breaks it
16     #up into regions using a clustering algorithm
17     #beta and eps are constants that affect how strongly edges in
18     #the image tend to be region boundaries
19     beta = 5.8
20     eps = 1e-6
21     N_REGIONS = 15 #Number of regions in each image. Need to figure
22     #out a way to do this programatically.
23
24     #empty lists to hold the graphs and the labels
25     graph_list = []
26     labels_list = []

```

```

25     for img in images: #cycle through the images
26         graph = image.img_to_graph(img) #form a graph network
linking each pixel to the others
27         graph.data = np.exp(-beta*graph.data/graph.data.std()) + eps
#exponentially suppress weak connections, pulling out edges
28         graph_list.append(graph) #stick it on the list of graphs
29
30         #perform the spectral clustering algorithm. amg is a fast
eigen solver
31         cluster_success = False #keeps track of whether or not the
clustering algorithms has worked without error
32         N_REGIONS = 15 #Number of regions in each image
33         while not cluster_success:
34             try:
35                 labels = spectral_clustering(graph, n_clusters=
N_REGIONS, assign_labels='kmeans', random_state=2, eigen_solver =
'amg')
36                 labels = labels.reshape(img.shape) #Now there's an
array and each pixel is assigned a label corresponding to its
37                 #region. Reshape it so this
array is a matrix with the same dimensions as the image
38                 labels = CloseRegions(labels) #Apply morphological
closing to remove tiny holes in a region
39                 labels = CleanRegions(labels) #Make sure the region
labeling doesn't skip a number
40                 #(could happen if
closing wiped out a region entirely)
41                 labels_list.append(labels)
42                 cluster_success = True
43             except (np.linalg.LinAlgError, ValueError): #If an error
was raised
44                 N_REGIONS -= 3 #Reduce the number of regions, and
try again
45                 cluster_success = False
46                 if N_REGIONS < 1: #If somehow we're all the way down
to 0 regions, just label everything as region 1
47                     labels_list.append(np.ones(img.shape))
48
49         return labels_list
50
51 def ShowRegions(img, labels, alpha = 0.65):
52     #Given an image img (2d matrix) and an overlapping matrix of
labels, plots the image with each region shaded a different color
. Alpha sets the transparency of the overlay

```

```

53     img = img - np.min(img)
54     img = img.astype(float)/np.max(img) #Get all image gray values
between 0 and 1
55     rows, cols = img.shape #size of image
56
57     rows2, cols2 = labels.shape
58     if rows != rows2 or cols != cols2:
59         print("Error: image and labels should have the same
dimensions")
60         return
61
62     img_color = np.dstack((img, img, img)) #create an RGB version of
the grayscale image
63     img_hsv = color.rgb2hsv(img_color) #convert image to HSV color
space
64     color_mask = np.zeros((rows, cols, 3)) #create an empty array to
hold the color mask
65
66     num_regions = np.max(labels) + 1
67     for i in range(num_regions):
68         color_mask[labels == i] = plt.cm.jet(i/num_regions)[:3]
69     color_mask_hsv = color.rgb2hsv(color_mask)
70
71     img_hsv[..., 0] = color_mask_hsv[...,0]
72     img_hsv[..., 1] = color_mask_hsv[..., 1]*alpha
73
74     img_masked = color.hsv2rgb(img_hsv)
75
76     f, (ax0, ax1, ax2) = plt.subplots(1, 3, figsize = (12,8))
77     ax0.imshow(img, cmap=plt.cm.gray)
78     ax1.imshow(color_mask)
79     ax2.imshow(img_masked)
80     plt.show()
81
82     return
83
84 def ShowAllRegions(images, labels_list, alpha=0.65):
85     #Visualizes the regions for all images in list images
86     #labels_list is the list of labels, alpha is the transparency
87     for img, labels in zip(images, labels_list):
88         ShowRegions(img, labels, alpha=alpha)
89     return
90
91 def CloseRegions(labels):

```



```

92     #Takes a set of labels for an image and closes any tiny missing
    dots, making the regions more compact
93     num_regions = np.max(labels) + 1 #figure out how many regions
    are in the image
94
95     for i in range(num_regions): #cycle through the regions
96         labels_subset = labels*(labels == i) #get a binary array of
    pixels that belong to region i
97         labels_subset_closed = morphology.binary_closing(
    labels_subset, morphology.square(3))
98         #close that array. Use a 3x3 square
99         labels[labels_subset_closed == 1] = i #Take the newly closed
    region and assign its number (i) to relevant pixels
100
101     return labels
102
103 def CleanRegions(labels):
104     #Takes a set of labels for an image and makes sure the labels
    are contiguous from 0 to num_regions-1
105     index = 0
106     while index < np.max(labels):
107         if not np.any(labels == index):
108             current_max = np.max(labels)
109             labels[labels == current_max] = index
110             index += 1
111
112     return labels

```

Functions_Feature_Extraction.py

The following functions calculate the features for each pixel.

```

1 import numpy as np
2 from scipy import ndimage as ndi
3 from skimage import transform
4
5 def NormalizeImage(img, background):
6     # Takes an image (2d array) and background (2d array of same
    size)
7     # and normalizes the image. This involves subtracting the
    background
8     # and then dividing by the average background signal, so any
    pixel that
9     # is the brightness of the background will have value 0,
    something that

```

```

10     # is much brighter will have value 1 (or higher), much darker -1
11     img = img - background
12     img = img.astype(float)
13     background_mean = np.mean(background)
14     img = img/background_mean
15     img = transform.rescale(img, 0.25)
16     return img
17
18 def MakeFeaturesMatrix(images, labels_list, d):
19     # Calculates several features for each pixel and turns it into a
20     # matrix
21     # images is a list of images, labels_list is a list of labels
22     # for each image
23     # (segmenting it into regions), d is the size of the box around
24     # each pixel
25     # which gets investigated (specifically, a square of side length
26     # 2d+1 is used)
27     #images_data = MakeImagesData(images) #I don't think this is
28     #doing anything.....
29     images_region_mean = MakeRegionMean(images, labels_list)
30     images_region_var = MakeRegionVar(images, labels_list)
31     images_region_var_norm = images_region_var/images_region_mean
32
33     images_local_mean = MakeLocalMean(images, labels_list, d)
34     images_local_var = MakeLocalVar(images, labels_list, d)
35     images_local_var_norm = images_local_var/images_local_mean
36
37     images_local_grad = MakeLocalGrad(images, labels_list, d)
38     images_local_grad_norm = images_local_grad/images_local_mean
39
40     X_all = np.stack((images_region_mean.flatten(),
41     images_region_var.flatten(), images_region_var_norm.flatten(),\
42     images_local_mean.flatten(), images_local_var.\
43     flatten(), images_local_var_norm.flatten(),\
44     images_local_grad.flatten(),
45     images_local_grad_norm.flatten()), axis=1)
46     #Stack all features into one matrix of size (nsamples, nfeatures
47     ), where nsamples is the total number of pixels in all images
48     return X_all
49
50 def MakeImagesData(images):
51     #Takes a list of images and smushes their data into a 3d array (
52     #rows by cols by number of images)
53     if not isinstance(images, list):

```

```

44     images = [images]
45
46     num_images = len(images)
47     (rows, cols) = images[0].shape
48     image_data = np.empty((rows, cols, num_images))
49     for i in range(num_images):
50         image_data[:, :, i] = images[i]
51
52     return image_data
53
54 def MakeRegionMean(images, labels_list):
55     #Take a list of images and a corresponding list of labels
56     matrices, and compute the regional mean for each one
57     if not isinstance(images, list):
58         images = [images]
59         labels_list = [labels_list]
60
61     num_images = len(images)
62     (rows, cols) = images[0].shape
63     images_region_mean = np.empty((rows, cols, num_images)) #make an
64     empty 3d array to hold data
65     for i, (img, labels) in enumerate(zip(images, labels_list)): #
66     cycle through the images and labels
67         images_region_mean[:, :, i] = ComputeRegionMean(img, labels) #
68     use ComputeRegionMean to do calculation for one image
69     return images_region_mean
70
71 def ComputeRegionMean(img, labels):
72     #For an image img and labels for some regions, compute the
73     average pixel value within each region
74     num_regions = np.max(labels) + 1
75     region_mean = np.empty(img.shape) #make empty matrix to hold
76     mean regional value at each pixel
77
78     for i in range(num_regions): #cycle through regions
79         region_mean[labels == i] = np.mean(img[labels == i]) #img[
80     labels == i] is the subset of the image for region i
81         #compute the mean within that region,
82     and assign that value to all the pixels for which labels == i
83     return region_mean
84
85 def MakeRegionVar(images, labels_list):
86     #Take a list of images and a corresponding list of labels
87     matrices, and compute the regional variance for each one

```

```

79     if not isinstance(images, list):
80         images = [images]
81         labels_list = [labels_list]
82
83     num_images = len(images)
84     (rows, cols) = images[0].shape
85     images_region_var = np.empty((rows, cols, num_images))
86     for i, (img, labels) in enumerate(zip(images, labels_list)):
87         images_region_var[:, :, i] = ComputeRegionVar(img, labels)
88     return images_region_var
89
90 def ComputeRegionVar(img, labels):
91     #For an image img and labels for some regions, compute the
92     variance of pixel value within each region
93     num_regions = np.max(labels) + 1
94     region_var = np.empty(img.shape) #make empty matrix to hold mean
95     regional value at each pixel
96
97     for i in range(num_regions): #cycle through regions
98         region_var[labels == i] = np.var(img[labels == i]) #img[
99         labels == i] is the subset of the image for region i
100         #compute the variance within that region
101         , and assign that value to all the pixels for which labels == i
102     return region_var
103
104 def MakeLocalMean(images, labels_list, d):
105     #Take a list of images and a corresponding list of labels
106     matrices, and compute the local mean for each one
107     if not isinstance(images, list):
108         images = [images]
109         labels_list = [labels_list]
110
111     num_images = len(images)
112     (rows, cols) = images[0].shape
113     images_local_mean = np.empty((rows, cols, num_images))
114     for i, (img, labels) in enumerate(zip(images, labels_list)):
115         images_local_mean[:, :, i] = ComputeLocalMean(img, labels, d)
116     return images_local_mean
117
118 def ComputeLocalMean(img, labels, d):
119     #For image img and regional labels, compute the mean of the area
120     around each pixel, respecting the region boundaries
121     #The local area is a square of side length 2d+1
122     (rows, cols) = img.shape

```

```

117     local_mean = np.empty((rows, cols))
118
119     #I apologize for the inelegance of this solution. It feels like
120     there should be something involving convolutions,
121     #but between wanting to only compute the mean within the masked
122     region and not wanting to go outside the image boundary,
123     #I couldn't come up with anything better than these crude loops
124     for row in range(rows):
125         for col in range(cols):
126             #Compute the boundaries of the box, not going outside
127             the boundaries of the image
128             min_x = max(col - d, 0)
129             max_x = min(col + d + 1, cols)
130             min_y = max(row - d, 0)
131             max_y = min(row + d + 1, rows)
132
133             #Create sub-matrices for labels and img, just including
134             the area around this pixel
135             labels_sub = labels[min_y:max_y, min_x:max_x]
136             img_sub = img[min_y:max_y, min_x:max_x]
137             #compute the mean within the pixels that are actually in
138             the same region as the pixel of interest
139             local_mean[row, col] = np.mean(img_sub[labels_sub ==
140             labels[row, col]])
141
142     return local_mean
143
144 def MakeLocalVar(images, labels_list, d):
145     #Take a list of images and a corresponding list of labels
146     matrices, and compute the local variance for each one
147     if not isinstance(images, list):
148         images = [images]
149         labels_list = [labels_list]
150
151     num_images = len(images)
152     (rows, cols) = images[0].shape
153     images_local_var = np.empty((rows, cols, num_images))
154     for i, (img, labels) in enumerate(zip(images, labels_list)):
155         images_local_var[:, :, i] = ComputeLocalVar(img, labels, d)
156     return images_local_var
157
158 def ComputeLocalVar(img, labels, d):
159     #For image img and regional labels, compute the variance of the
160     area around each pixel, respecting the region boundaries

```

```

153     #The local area is a square of side length 2d+1
154     (rows, cols) = img.shape
155     local_var = np.empty((rows, cols))
156
157     for row in range(rows):
158         for col in range(cols):
159             #Compute the boundaries of the box, not going outside
the boundaries of the image
160                 min_x = max(col - d, 0)
161                 max_x = min(col + d + 1, cols)
162                 min_y = max(row - d, 0)
163                 max_y = min(row + d + 1, rows)
164
165                 #Create sub-matrices for labels and img, just including
the area around this pixel
166                 labels_sub = labels[min_y:max_y, min_x:max_x]
167                 img_sub = img[min_y:max_y, min_x:max_x]
168                 #compute the mean within the pixels that are actually in
the same region as the pixel of interest
169                 local_var[row, col] = np.var(img_sub[labels_sub ==
labels[row, col]])
170
171     return local_var
172
173 def MakeLocalGrad(images, labels_list, d):
174     #Take a list of images and a corresponding list of labels
matrices, and compute the local gradient magnitude for each one
175     if not isinstance(images, list):
176         images = [images]
177         labels_list = [labels_list]
178
179     num_images = len(images)
180     (rows, cols) = images[0].shape
181     images_local_grad = np.empty((rows, cols, num_images))
182
183     for i, (img, labels) in enumerate(zip(images, labels_list)):
184         images_local_grad[:, :, i] = ComputeLocalGrad(img, labels, d)
185     return images_local_grad
186
187 def ComputeLocalGrad(img, labels, d):
188     #For image img and regional labels, compute the average gradient
magnitude of the area around each pixel, respecting the
189     #region boundaries. The local area is a square of side length 2d
+1

```

```

190     img = ndi.gaussian_filter(img, sigma=2)
191     img_grad_all = np.gradient(img)
192     img_grad = np.sqrt(img_grad_all[0]**2 + img_grad_all[1]**2)
193
194     (rows, cols) = img.shape
195     local_grad = np.empty((rows, cols))
196
197     for row in range(rows):
198         for col in range(cols):
199             #Compute the boundaries of the box, not going outside
the boundaries of the image
200                 min_x = max(col - d, 0)
201                 max_x = min(col + d + 1, cols)
202                 min_y = max(row - d, 0)
203                 max_y = min(row + d + 1, rows)
204
205                 #Create sub-matrices for labels and img, just including
the area around this pixel
206                 labels_sub = labels[min_y:max_y, min_x:max_x]
207                 img_grad_sub = img_grad[min_y:max_y, min_x:max_x]
208                 #compute the mean within the pixels that are actually in
the same region as the pixel of interest
209                 local_grad[row, col] = np.mean(img_grad_sub[labels_sub =
= labels[row, col]])
210
211     return local_grad

```

Functions_Visualization.py

The following functions overlay an image with its classifications.

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3 from skimage import color
4
5 def ShowClassification(img, img_class):
6     #Given an image img (2d matrix) and an overlapping matrix of
classifications, plots the image
7     #with each region shaded a different color. Similar to
ShowRegions except the color for each classification type
8     #is specified below
9     alpha = 0.95 #Transparency for the overlap
10    img = img - np.min(img)
11    img = img.astype(float)/np.max(img) #Get all image gray values
between 0 and 1

```

```
12     rows, cols = img.shape #size of image
13
14     rows2, cols2 = img_class.shape
15     if rows != rows2 or cols != cols2:
16         print("Error: image and labels should have the same
17 dimensions")
18         return
19
20     img_color = np.dstack((img, img, img)) #create an RGB version of
the grayscale image
21     img_hsv = color.rgb2hsv(img_color) #convert image to HSV color
space
22     color_mask = np.zeros((rows, cols, 3)) #create an empty array to
hold the color mask
23
24     num_classes = np.max(img_class) + 1
25     for i in range(1, num_classes):
26         color_mask[img_class == i] = GetClassColor(i)
27
28     color_mask_hsv = color.rgb2hsv(color_mask)
29     img_hsv[..., 0] = color_mask_hsv[...,0]
30     img_hsv[..., 1] = color_mask_hsv[..., 1]*alpha
31     img_masked = color.hsv2rgb(img_hsv)
32
33     #f, (ax0, ax1, ax2) = plt.subplots(1, 3, figsize = (12,8))
34     #ax0.imshow(img, cmap=plt.cm.gray)
35     #ax1.imshow(color_mask)
36     #ax2.imshow(img_masked)
37     fig = plt.imshow(img_masked)
38     plt.show()
39
40     return
41 def GetClassColor(i):
42     if i == 0:
43         return (1, 1, 1) #unknown: white
44     if i == 1:
45         return (1, 0, 0) #amorphous: red
46     if i == 2:
47         return (0, 1, 0) #platelet: green
48     if i == 3:
49         return (0, 0, 1) #spherulite: blue
50     if i == 4:
51         return (1, 1, 0) #junk: yellow
```



```

52     return (1, 0, 1) #other: purple
53
54
55 def ShowAllClassification(images, classifications):
56     #Calls ShowClassification for an array of images and a three
57     dimensional classifications matrix
58     #(the last dimension corresponds to iamge number)
59     for i, image in enumerate(images):
60         ShowClassification(image, classifications[:, :, i])
61
62     return

```

Functions_Training.py

The following functions train, test, and evaluate a random forest classifier.

```

1 import numpy as np
2 import matplotlib.pyplot as plt
3 import itertools
4 import os
5 import time
6 from sklearn.ensemble import RandomForestClassifier
7 from sklearn.metrics import confusion_matrix
8 from Functions_Visualization import ShowAllClassification
9 from Functions_Feature_Extraction import MakeFeaturesMatrix
10
11 def WriteTrainingClassification(images, train_data):
12     #Takes a the encoded training data and writes it to a 3D matrix
13     matching the dimension of images,
14     #so that each pixel is classified
15     num_images = len(images)
16     (rows, cols) = images[0].shape
17     classifications = np.zeros((rows, cols, num_images)) #
18     classifications has the same dimensions as images, there's a
19     #one-to-one correspondence between a pixel
20     and its class
21
22     #Break out train_data into its constituent parts
23     train_ymin = train_data[0]
24     train_ymax = train_data[1]
25     train_xmin = train_data[2]
26     train_xmax = train_data[3]
27     train_img_num = train_data[4]
28     train_class = train_data[5]

```

```
27     num_sets = len(train_ymin)
28
29     if not all (num_sets == len(subdata) for subdata in train_data):
30         print("Warning: you've got some mismatched lengths in your
31 training data")
32
33     #Iterate through the constituents parts of train_data (i.e. take
34     the i'th entry of each of the 6 arrays)
35     for (ymin, ymax, xmin, xmax, img_num, pxl_class) in zip(
36 train_ymin, train_ymax, \
37         train_xmin, train_xmax, train_img_num
38 , train_class):
39         classifications[ymin:ymax, xmin:xmax, img_num] = pxl_class #
40 Assign the appropriate block of pixels to pxl_class
41
42     return classifications.astype(int)
43
44 def KfoldRandomForest(images, X_all, full_train_data, n_trees=50, k=
45 5):
46     #Takes a series of images, region labels, and already-classified
47     training regions, and performs k-fold validation
48     #using a Random Forest Classifier.
49     #It splits the training data into k subsets, and iteratively
50     sets aside one to be the test set. It uses the other k-1
51     #subsets to train the random forest model, and then tests the
52     results on the test set. These results are compiled
53     #to create a confusion matrix
54     #Images is an array of images, each of which is a 2D array of
55     greyscale values from 0 to 255
56     #Xall is a 2D features array of size nsamples*nfeatures, where
57     nsamples is the total number of pixels in all images, and
58     nfeatures is the number of features we use to classify the pixel
59     #full_train_data is list full of information necessary to
60     classify certain training pixels. It contains 6 arrays: ymin,
61     ymax, xmin, xmax, img_num, and class. The elements of each array
62     are used to define a rectangle on a specific image and classify
63     it as being a certain type
64     #n_trees is the number of estimators to use in the forest, and k
65     is the number of sets to split training data into
66
67     #Split the training data into k subsets
68     #There are thousands of pixels in the training set, but we want
69     to split randomly by region, not pixel.
```

```
52     #That's because the pixels in a given region are often quite
53     similar to each other, so if the model has access to
54     #most of the pixels in a given region, it's not impressive to
55     get the rest of the pixels in that region correct.
56     #It's more representative to test the model on regions it has
57     never seen before
58
59     num_train_regions = len(full_train_data[0]) #Number of training
60     regions
61     train_regions_split = KfoldSplit(num_train_regions, k) #Split
62     training regions into k sets
63
64     #Perform k-fold validation
65     cm_sum = TestRandomForest(images, X_all, full_train_data,
66     train_regions_split, n_trees)
67
68     #Plot resulting confusion matrix and accuracy
69     class_names = ["amorphous", "platelet", "spherulite", "junk"]
70     plt.figure()
71     PlotConfusionMatrix(cm_sum, classes=class_names, title='Total
72     Confusion matrix')
73
74     plt.figure()
75     PlotModelAccuracy(cm_sum, classes=class_names, title="Total
76     Accuracy")
77     plt.show()
78
79     #Now train the model on all of the data
80     classifications = WriteTrainingClassification(images,
81     full_train_data)
82     classifier = TrainForest(X_all, classifications, n_trees)
83
84     #Run the model on all pixels
85     Y_pred = classifier.predict(X_all)
86     images_predicted_class = np.reshape(Y_pred, classifications.
87     shape)
88     plt.figure()
89     ShowAllClassification(images, images_predicted_class)
90
91     return classifier, cm_sum
92
93 def CompareRandomForests(images, labels_list, full_train_data,
94     subdir, d_list, n_trees_list, k=5):
```

```
84     #CompareRandomForests performs k-fold cross validation on
85     several random forest classifiers, each with different parameters
86     #images is an array of images, each of which is a 2D array of
87     greyscale values from 0 to 255
88     #labels_list is a list of arrays, each of which labels every
89     point within a given image as belonging to a particularly region
90     #full_train_data encodes all the training data (several regions
91     each of which are already classified as amorphous, platelet,
92     spherulite, or junk)
93     #subdir is the name of a directory in which to save the results
94     #d_list is a list of d values to test. d controls how large an
95     area a given pixel samples in order to determine its local
96     landscape
97     #n_trees_list is a list of n_trees values to test. n_trees is
98     the number of estimators in the forest model
99     #k is how many subsets to split the total training data into for
100     doing k-fold validation
101
102     num_train_regions = len(full_train_data[0]) #Number of training
103     regions
104     train_regions_split = KfoldSplit(num_train_regions, k) #Split
105     training regions into k sets
106
107     num_d = len(d_list)
108     num_n_trees = len(n_trees_list)
109     class_names = ["amorphous", "platelet", "spherulite", "junk"]
110     num_classes = len(class_names)
111
112     #cm = np.array([[10482,      1,  1189,      0], [ 473,  6004,
113     226,    21], [ 9730,  9610, 19348,   291], [   0,    96,   548,
114     801]])
115     x = np.arange(num_classes)
116     width = 0.9
117     cmap = plt.cm.plasma
118
119     plt.figure(figsize=(3*num_n_trees, 2*num_d))
120     plt.suptitle(f'Accuracy of Random Forest Classifiers for Various
121     Parameters (k = {k})')
122     subplot_num = 1
123
124     for d in d_list:
125         X_all = MakeFeaturesMatrix(images, labels_list, d)
126         for n_trees in n_trees_list:
127             t1 = time.time()
```

```

114         cm = TestRandomForest(images, X_all, full_train_data,
115         train_regions_split, n_trees)
116         t2 = time.time()
117
118         plt.subplot(num_d, num_n_trees, subplot_num)
119
120         accuracy = np.diag(cm)/np.sum(cm, axis=1)
121         plt.bar(x, accuracy, width, color=cmap(accuracy))
122
123         if d == d_list[0]:
124             plt.title(f'{n_trees} trees')
125         if d == d_list[-1]:
126             plt.xticks(x, class_names, rotation=45)
127         else:
128             plt.xticks(x, ['', ]*num_classes)
129         if n_trees == n_trees_list[0]:
130             plt.yticks([0, .5, 1])
131             plt.ylabel(f'd = {d}')
132         else:
133             plt.yticks([0, .5, 1], ['', ', ', ', '])
134         subplot_num += 1
135
136         WriteRandomForestResults(subdir, k, d, n_trees, t2-t1,
137         cm)
138
139     return
140
141 def TestRandomForest(images, X_all, full_train_data,
142 train_regions_split, n_trees):
143     #TestRandomForest cycles through the k folds, for each one holds
144     it back as a test set, builds the model, and runs it
145     #The results are compiled into a total confusion matrix
146     #X_all is a matrix of features
147     #full_train_data is the information required to construct a
148     classifications vector
149     #train_regions_split tells the function how to split the
150     training data into k subsets. It is an array of arrays,
151     #each of which selects 1/k'th of the training data
152     #n_trees is the number of trees to use in the forest
153
154     cm_all = [] #will hold all of the confusion matrices from each
155     of the k validations

```

```
150     for train_regions_sub in train_regions_split: #Cycle through the
151         k-fold subsets of train_regions_split
152             train_data, test_data = SplitTrainTest(full_train_data,
153             train_regions_sub)
154
155             #Create training/testing matrices by assigning
156             classifications
157             train_matrix = WriteTrainingClassification(images,
158             train_data)
159             test_matrix = WriteTrainingClassification(images, test_data)
160
161             #Use train_matrix and features X_all to create and train
162             Random Forest Classifier, then apply to test data
163             classifier = TrainForest(X_all, train_matrix, n_trees)
164             cm = CheckForest(classifier, X_all, test_matrix) #returns
165             confusion matrix
166
167             #Add to list of all confusion matrices
168             cm_all.append(cm)
169
170             #Create and plot the total confusion matrix
171             cm_sum = SumConfusionMatrices(cm_all)
172
173             return cm_sum
174
175 def SumConfusionMatrices(cm_all):
176     #Takes a list of confusion matrices (although they could really
177     be any matrices) and sums them all together
178     cm1 = cm_all[0] #take first matrix
179     cm_shape = cm1.shape #get the shape
180     cm_sum = np.zeros(cm_shape, dtype=int) #use that to create an
181     appropriately-sized array of zeros
182
183     for cm in cm_all: #loop through matrices and sum them
184         cm_sum += cm
185
186     return cm_sum
187
188 def PlotConfusionMatrix(cm, classes, normalize=False, title='
189 Confusion matrix', cmap=plt.cm.Blues):
190     #This function prints and plots the confusion matrix.
191     #Normalization can be applied by setting 'normalize=True'.
192     #courtesy of scikit-learn.org
```

```

185     if normalize:
186         cm = cm.astype('float') / cm.sum(axis=1)[:, np.newaxis]
187
188     plt.imshow(cm, interpolation='nearest', cmap=cmap)
189     plt.title(title)
190     plt.colorbar()
191     tick_marks = np.arange(len(classes))
192     plt.xticks(tick_marks, classes, rotation=45)
193     plt.yticks(tick_marks, classes)
194
195     fmt = '.2f' if normalize else 'd'
196     thresh = cm.max() / 2.
197     for i, j in itertools.product(range(cm.shape[0]), range(cm.shape
198 [1])):
199         plt.text(j, i, format(cm[i, j], fmt),
200                 horizontalalignment="center",
201                 color="white" if cm[i, j] > thresh else "black")
202
203     plt.tight_layout()
204     plt.ylabel('True label')
205     plt.xlabel('Predicted label')
206
207 def PlotModelAccuracy(cm, classes, title='Accuracy', cmap=plt.cm.
208 plasma):
209     #Uses the confusion matrix to plot the accuracy of the model.
210     #That is, for each category what fraction of pixels did it
211     classify correctly
212
213     #Calculate accuracy for each class as the number of correctly
214     identified pixels (diagonal) divided by the
215     #total number of that class (sum along columns)
216     accuracy = np.diag(cm)/np.sum(cm, axis=1)
217
218     #Create x positions and width of bars
219     x = np.arange(len(accuracy))
220     width = 0.9
221
222     #Make plot, coloring each bar by its height. Set title and ticks
223     plt.bar(x, accuracy, width, color=cmap(accuracy))
224     plt.title(title)
225     plt.xticks(x, classes, rotation=45)
226     plt.yticks([0, .5, 1])
227
228     #plt.show()

```

```
225     return
226
227
228 def KfoldSplit(n, k):
229     #KfoldSplit performs the math to randomly splits n objects into
230     k roughly equal sized chunks
231     #It does this abstractly, but creating an array 0, 1, ..., n-1,
232     shuffling that, and splitting it
233     #Each element of the resulting array of arrays can then be used
234     to select the actual objects of interest
235     indices = np.arange(n) #Create array from 0 to n-1
236     np.random.shuffle(indices)#shuffle this array so the split is
237     random
238     return np.array_split(indices, k) #Split into k similarly-sized
239     sections
240
241 def SplitTrainTest(full_data, test_indices):
242     #SplitTrainTest takes some full training data and splits it into
243     training and testing subsets
244     #It does this abstractly -- full_data is just a list of arrays,
245     each of which has the same length.
246     #In practice those arrays encode the corners, image numbers, and
247     classes for a bunch of training regions,
248     #but this function does not concern itself with such details.
249     #It just uses the test_indices to select some subset of the full
250     training data
251     #For example if test_indices = [3, 7, 10] then for each array in
252     the list full_data it will pull out the 3rd, 7th, and 10th
253     #elements and use those to make a new list of arrays (the test
254     data). What remains is the train data.
255
256     train_data, test_data = [], []
257     for subarray in full_data:
258         test_data.append(subarray[test_indices])
259         train_data.append(np.delete(subarray, test_indices))
260
261     return train_data, test_data
262
263 def TrainForest(X_all, classifications, n_trees):
264     #TrainForest takes in a matrix of features (X_all), a matrix of
265     classes (classifications), and trains a
266     #Random Forest model with number of trees equal to n_trees
267     Y_train = classifications.flatten() #flatten the training
268     classes matrix
```



```
256     #only keep features and classes that have been classified
already (i.e. where Y_train is not 0)
257     X_train = X_all[Y_train != 0, :]
258     Y_train = Y_train[Y_train != 0]
259
260     #create and fit the model
261     classifier = RandomForestClassifier(n_estimators=n_trees)
262     classifier.fit(X_train, Y_train)
263
264     return classifier
265
266 def CheckForest(classifier, X_all, test_matrix):
267     #CheckForest takes an already-trained forest and applies it to
some test data, returning the confusion matrix
268     Y_test = test_matrix.flatten() #flatten the text classes matrix
269     #only keep features and classes that have been classified
already (i.e. where Y_test is not 0)
270     X_test = X_all[Y_test != 0, :]
271     Y_test = Y_test[Y_test != 0]
272
273     Y_pred = classifier.predict(X_test) #run the model on the test
data
274     cm = confusion_matrix(Y_test, Y_pred, labels=[1, 2, 3, 4]) #
create confusion matrix
275
276     return cm
277
278 def WriteRandomForestResults(subdir, k, d, n_trees, dt, cm):
279     #Writes the results of doing k-fold validation on a given
training set to an output file
280     #This file records the number of folds (k), the model
parameters d and n_trees, the time it took (dt),
281     #and the results of the model as summarized with a confusion
matrix
282
283     try:
284         os.mkdir(subdir)
285     except Exception:
286         pass
287
288     filename = f'trees_{n_trees} d_{d}.txt'
289     path = os.path.join(subdir, filename)
290
291     f = open(path, 'w')
```

```
292     f.write(f'k = {k} (Number of validation folds)\n')
293     f.write(f'd = {d} (size of local area probed around each pixel)\n')
294     f.write(f'n')
295     f.write(f'trees = {n_trees}\n')
296     f.write(f'Time = {dt:.2f} seconds\n')
297
298     f.write('Confusion matrix:\n')
299     np.savetxt(f, cm, fmt='%d')
300
301     accuracy = np.diag(cm)/np.sum(cm, axis=1)
302     f.write('Accuracy:\n')
303     np.savetxt(f, accuracy, fmt='%.4f')
304
305     f.close()
306
307     return
```