

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Scalable Methods for Detecting Microlensing Black Holes with the Zwicky Transient Facility

Permalink

<https://escholarship.org/uc/item/2v4173vq>

Author

Medford, Michael Seth

Publication Date

2021

Peer reviewed|Thesis/dissertation

Scalable Methods for Detecting Microlensing Black Holes with the Zwicky Transient
Facility

by

Michael S. Medford

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Astrophysics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Associate Professor Jessica Lu, Co-chair

Doctor Peter Nugent, Co-chair

Professor Joshua S. Bloom

Summer 2021

Scalable Methods for Detecting Microlensing Black Holes with the Zwicky Transient
Facility

Copyright 2021
by
Michael S. Medford

Abstract

Scalable Methods for Detecting Microlensing Black Holes with the Zwicky Transient Facility

by

Michael S. Medford

Doctor of Philosophy in Astrophysics

University of California, Berkeley

Associate Professor Jessica Lu, Co-chair

Doctor Peter Nugent, Co-chair

Despite models predicting the presence of millions of isolated black holes in the Milky Way galaxy, not a single one has been observed to date. Microlensing of background stars caused by these non-luminous black holes is theoretically the only observational method for detecting them. However black hole candidates obtained through photometric microlensing for astrometric follow-up have been scarce due to the limited number of surveys looking for microlensing events.

This thesis presents the first wide-field untargeted microlensing survey designed and executed on a multi-purpose synoptic instrument. The Zwicky Transient Facility (ZTF) has taken observations covering the visible Northern night sky in multiple filters every few nights for several years. This work makes predictions for the number of observable microlensing events contained within the ZTF surveys, implements significant improvements to ZTF long-wavelength optical calibration necessary to characterize these events, and carries out an all-sky search for microlensing events in the facility's first several years of observations. A list of microlensing events containing black hole candidates is produced that cover regions of the Galactic plane and field stars not previously covered by microlensing campaigns.

ZTF is a predecessor to upcoming synoptic surveys such as the Vera C. Rubin Observatory and the Nancy Grace Roman Space Observatory. Challenges exist in processing the massive heterogeneous datasets these instruments will generate as they attempt to further many different scientific pursuits. The computational and statistical techniques developed within this thesis can help realize the ambitions to conduct microlensing surveys on these instruments.

To my grandparents Jerry, Doris, Joe and Janet, for each in their own way demonstrating to me the joy and responsibility of loving those around you.

To my parents Nancy and David, for teaching me I could accomplish whatever I set my mind to with patience, determination, and a belief that it is possible.

To the love of my life Brittany for giving me the greatest gift I could ever ask for, her unwavering commitment to building a life together. For seeing in me the best version of myself and for helping me to grow toward it. For showing me everyday how to truly love someone. This is for you, for me and for us.

Contents

List of Figures	v
List of Tables	vi
Acknowledgements	viii
1 Introduction	1
1.1 Gravitational Microlensing	3
1.1.1 Formalism	3
1.1.2 Observation	4
1.1.3 Black Holes	5
1.2 The Zwicky Transient Facility	7
1.2.1 Surveys and Data Products	7
1.2.2 Data Analysis Challenges	7
1.3 Thesis Outline	8
2 Gravitational Microlensing Event Statistics for the Zwicky Transient Facility	10
2.1 Introduction	11
2.2 The Zwicky Transient Facility	13
2.3 ZTF Microlensing Estimate	15
2.3.1 Event Rate: Γ_{obs}	16
2.3.2 Number of Stars: $N_{\text{stars}}^{\text{ZTF}}$, Survey Duration: T_{obs}	20
2.3.3 Interpolation Across the Galactic Plane	21
2.3.4 Results: ZTF Microlensing Event Statistics	21
2.4 Microlensing Population Properties in the Outer Galaxy ($\ell \geq 10^\circ$)	23
2.4.1 Distance Ratio	23
2.4.2 Einstein Crossing Time, Einstein Radius, and Relative Proper Motion	23
2.4.3 Extinction	25
2.4.4 Contribution of Neighbors to Blended Light	27
2.4.5 Implications for Outer Galaxy Microlensing	28
2.5 Example Outer Galaxy Microlensing Event Analysis	29

2.5.1	Event Selection	29
2.5.2	Event Analysis	30
2.6	Discussion	35
3	Removing Atmospheric Fringes from Zwicky Transient Facility i-Band Images using Principal Component Analysis	42
3.1	Introduction	43
3.2	The Zwicky Transient Facility Instrument	44
3.3	Method for Removing Atmospheric Fringes	45
3.3.1	Generating PCA Fringe Models	46
3.3.2	Removing Fringes from Single-Epoch Images	47
3.4	Measuring Improved Photometric Precision	51
3.4.1	The Uniform Background Indicator	51
3.4.2	Photometric Error Due to Fringes	53
3.4.3	Effects on Fake Sources: Single Epoch	54
3.4.4	Effects on Fake Sources: Multi Epoch	56
3.5	Discussion	57
4	Transformations from Pan-STARRS1 and UBV Filters into ZTF Filters	66
4.1	Introduction	66
4.2	Pan-STARRS1 to ZTF Photometric Conversions	67
4.3	UBV to ZTF Photometric Conversions	67
5	128 Microlensing Events from the Three Years of Zwicky Transient Facility Phase One	71
5.1	Introduction	72
5.2	The Zwicky Transient Facility	74
5.2.1	Surveys and Data Releases	74
5.2.2	Object Lightcurves	75
5.3	Software	78
5.3.1	zort: ZTF Object Reader Tool	78
5.3.2	PUZLE: Pipeline Utility for ZTF Lensing Events	80
5.3.3	Upgrading to a New Public Data Release	82
5.4	Detection Pipeline	83
5.4.1	Cutting on the von Neumann ratio, Star Catalogs and a Four-Parameter Microlensing Model	83
5.4.2	Simulated Microlensing Events	88
5.4.3	Cutting on a Seven-Parameter Microlensing Model	90
5.4.4	Cutting on a Bayesian Microlensing Model	91
5.4.5	Manual Lightcurve Inspection	93
5.5	Results	93
5.6	Discussion	97

5.7	PUZLE Database Tables	102
5.8	NERSC	103
5.8.1	Python Environment	103
5.8.2	uLens Group	103
5.8.3	Database Access	103
5.9	PUZLE Website	103
5.9.1	Website Features	104
5.9.2	Hosting and Updating	104
6	Conclusion	120
	Bibliography	121

List of Figures

1.1	Example of Microlensing Lightcurve	2
1.2	Limits on MACHO Dark Matter Fraction	5
1.3	Einstein Crossing Times of Primordial Black Holes	6
1.4	Atmospheric Fringing on ZTF i-Band Images	8
2.1	ZTF Public Data Release 1 Lightcurve Statistics	14
2.2	Observable Microlensing Event Map	19
2.3	Observable Microlensing Event Rates	20
2.4	Galactic Plane Events: Distance Ratio	24
2.5	Galactic Plane Events: Einstein Crossing Time	25
2.6	Galactic Plane Events: Einstein Radii	26
2.7	Galactic Plane Events: Extinction	27
2.8	Galactic Plane Events: Neighbor Flux Fraction	28
2.9	ZTF18abhxjmj: Lightcurve	31
2.10	Galactic Plane Events: Lens-Source-Flux Ratio Approximation	33
2.11	ZTF18abhxjmj: Color-Magnitude Diagram	39
2.12	ZTF18abhxjmj: Absolute Color-Magnitude Diagram	40
2.13	Detectability of Black Holes via Galactic Plane Microlensing	41
3.1	<code>fringe</code> Schematic	45
3.2	Example Images in the Stages of Fringe Removal	49
3.3	Average of Training Fringe Maps: Correlation with CCD Thickness	50
3.4	Explained Variance of Fringe Model PCA Components	59
3.5	Uniform Background Indicator: Example Images	60
3.6	Uniform Background Indicator: Distributions	61
3.7	Uniform Background Indicator: Correlation with Airmass	62
3.8	Uniform Background Indicator: Correlation with Fringe Photometric Error	63
3.9	Improved Photometry in Single Epoch Images	64
3.10	Improved Photometry in Mutli-Epoch Co-additions	65
4.1	PS1 and UBV Photometric Transformations to ZTF Filters	69
5.1	PUZLE Diagram	77
5.2	Processing Job Cells	81

5.3	Example Lightcurves with η Values	85
5.4	Example Simulated Lightcurves with η Values	89
5.5	The η - η_{residual} Plane	90
5.6	Seven-Parameter Microlensing Model	92
5.7	Example Label Lightcurves	94
5.8	PUZLE Website	95
5.9	Level 6 Catalog: Sky Location	96
5.10	Level 6 Catalog: Sample Lightcurves	97
5.11	Level 6 Catalog: Einstein Crossing Time	98
5.12	Level 6 Catalog: Luminosity Function	107
5.13	Level 6 Catalog: Source Flux Fraction	108
5.14	Level 6 Catalog: t_E - π_E Space	109
5.15	Level 4 Ongoing Candidates	111

List of Tables

2.1	PopSyCLE Observational Cuts	18
2.2	Description of Fiducial Microlensing Simulations	22
2.3	Microlensing Parameters of ZTF18abhxjmj	37
2.4	Model Magnitudes of ZTF18abhxjmj	37
3.1	fringe _z Math Symbols	46
5.1	Object Statistics for Public Data Release 5	76
5.2	Boundaries for Four-Parameter Microlensing Model	87
5.3	PUZLE Pipeline Cuts	105
5.4	Manual Label Descriptions	106
5.5	Exported Candidate Schema	110
5.6	PUZLE: <code>source</code>	112
5.7	PUZLE: <code>star</code>	112
5.8	PUZLE: <code>source_ingest_job</code>	113
5.9	PUZLE: <code>star_ingest_job</code>	113
5.10	PUZLE: <code>star_process_job</code>	114
5.11	PUZLE: <code>candidate_level2</code>	115
5.12	PUZLE: <code>candidate_level3</code>	116
5.13	PUZLE: <code>candidate_level4</code> 1	117

5.14 PUZLE: candidate_level4 2	118
5.15 PUZLE: candidate_level4 3	119

Acknowledgments

The work contained within this thesis is the culmination of decades of teaching by thoughtful, persistent, and enthusiastic educators under which I have had the good fortune to study. They opened my eyes to the beauty of the Universe and the richly rewarding experience of studying its phenomena. The transfer of understanding and appreciation from one mind to another is one of the most generous gifts a person can give. I owe any successes I have had to my brilliant teachers.

Thank you to Carl Sagan, the first science educator to myself and millions of others, for translating the enigmatic secrets of the Cosmos into a language we could all understand. Thank you to my acting teachers Ann Woodworth and David Catlin for each teaching me how to express myself to the world using someone else's words. Thank you to Michael Smutko for awarding an astrophysics research position to an acting major, patiently introducing me to scientific research, and encouraging me to take ownership over and have pride in my work. Thank you to Arthur Schmidt for hiring me into my first teaching position and giving me the opportunity to learn how to combine my loves of science and performance. Thank you to Vicky Kalogera for believing in my academic potential and providing an actor the opportunity to pursue a scientific research career. Thank you to all of the students I have ever had the pleasure to teach, who helped me to learn how difficult it is to truly comprehend something and how rewarding it is once you do.

I would not have made it through the ups and downs of graduate school without the unflinching generosity and support of Peter Nugent. Always guiding his students to what is in their own best interest, Peter has created more opportunities and removed more obstacles for me than I can count. Thank you to Peter for teaching me how to practically answer a technical question, how to find my own "unfair advantage", and how to never lose sight of what is really important. He is an excellent advisor and a true friend.

One is quite lucky to find a great research advisor, and I have been truly blessed to have had two. I would have left graduate school disappointed and without completing my PhD if not for my work with Jessica Lu. Jessica taught me how to see a project through to its end and not get stuck "polishing the cannonball". She gave me the freedom to develop independence while simultaneously creating the structure I needed to learn discipline. Thank you to Jessica for illuminating the many different pathways to accomplishing meaningful research and for opening my eyes to opportunities beyond academia.

The love and support of my friends and family made all that I have accomplished possible. Thank you to Josh, Liz, Nate, Saul, Will, Abbey, Wren, Alex, and Lee for being the kindest and most inspiring friends anyone could ask for. Thank you to Allison, Rachel, and Daniel for being not only great siblings but dear friends. Thank you to my parents Nancy and David for miraculously supporting their son whether he is pursuing acting, astrophysics, or anything else. And thank you most of all to Brittany for being the love of my life and my very best friend. We did it Beedy!

Chapter 1

Introduction

In 1915, Albert Einstein published his theory of general relativity, revolutionizing our conception of gravity for the first time in over two hundred years. This theory transformed our description of gravity from a force between two objects into a set of rules describing the warping of a four-dimensional spacetime. However it was several years earlier in 1912 when he jotted down equations in his personal notebook describing how the gravitational influence of objects could, by bending their surrounding spacetime, focus the light of background stars like a lens (Remn 1997). Twenty-four years later he shared the final equations of this work in a note in the journal *Science* entitled *Lens-Like Action of a Star by the Deviation of Light in the Gravitational Field* (Einstein 1936). In these pages Einstein invented the field of gravitational microlensing. However he failed to foresee the future advances in experimental techniques and incorrectly concluded that this phenomenon would never be observable. Refsdal & Bondi (1964) formalized the mathematics of gravitational microlensing some years later and asserted that it could be observed through the apparent amplification of a background star. B. Paczyński (1986) outlined the modern methods still used today for finding microlensing events in the digital area.

Today thousands of microlensing events are observed each year by telescopes across the globe (Sumi et al. 2011; Wyrzykowski et al. 2015; Mróz et al. 2017). These targeted microlensing surveys (and others) monitor millions of Milky Way stars nearly every night in dense stellar fields in an attempt to capture this rare astrophysical event. The resulting gravitational microlensing detections have contributed to many significant discoveries in different areas of astrophysics. They have informed our models of the Milky Way Galaxy (Aubourg et al. 1993; Kerins 1995; Navarro et al. 2020), discovered exoplanets around distant stars (Cassan et al. 2012; Gaudi 2012), and placed constraints on the types of matter contributing to the dark matter halo of our galaxy (Alcock et al. 2001; Afonso et al. 2003; Wyrzykowski et al. 2011a; Niikura et al. 2019). The search for isolated black holes has been conducted using microlensing by observing individual events (Alcock et al. 2001; Afonso et al. 2003; Wyrzykowski et al. 2011a; Lu et al. 2016; Niikura et al. 2019; Arredondo et al. 2019; Abdurrahman et al. 2021) and by statistical analysis of event populations (Lu et al. 2019; Wyrzykowski et al. 2020; Golovich et al. 2020). Large-scale synoptic surveys conducted on

future instruments such as the Vera C. Rubin Observatory and the Nancy Grace Roman Space Telescope have the potential to increase the number of microlensing detected by an order of magnitude (Sajadian & Poleski 2019; Novati et al. 2014; Penny et al. 2019).

However these next-generation surveys will face significant challenges when attempting to observe microlensing. Synoptic surveys are designed with multiple science goals in mind, not specifically designed for microlensing. Synoptic data sets that are heterogeneous in cadence and pointed at areas of the sky where microlensing events are not typically observed. They will observe to deeper magnitudes, creating a problem with crowding in dense galactic fields that may make it easier to discover microlensing events at longer Galactic latitudes previously thought too sparsely populated to be worth observing. The massive amounts of data generated by these surveys further necessitates automated production pipelines with sophisticated statistical methods for characterizing and rejecting outliers when searching for rare microlensing signals. And while these new surveys cover an impressive range of wavelengths, valuable for microlensing due to its achromaticity, long wavelength optical images suffer from atmospheric fringing that reduces the signal-to-noise for faint sources. In this thesis I will lay out the path forward for transforming these general purpose synoptic surveys into microlensing discovery machines, including the first all-sky microlensing survey to ever be conducted on a synoptic instrument.

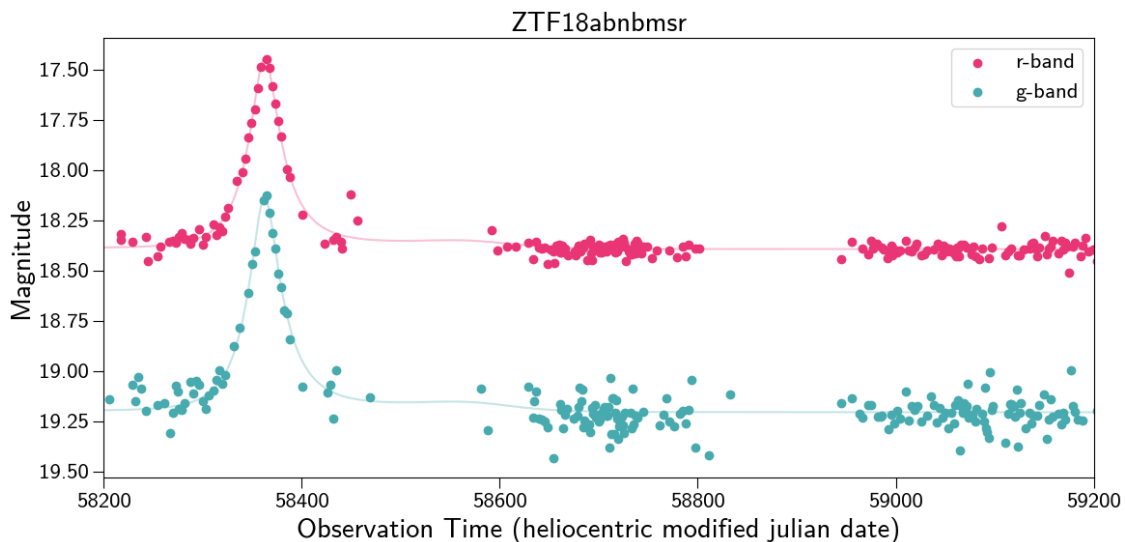


Figure 1.1: Example of a microlensing event found by Mróz et al. (2020a) on the Zwicky Transient Facility. The observations taken in multiple filters are well fit by a microlensing model. The achromaticity (in the absence of blending from neighboring stars) and near-symmetry (in cases with small annual parallax) are two of the markers that distinguishes microlensing from other forms of stellar variability.

1.1 Gravitational Microlensing

Gravitational microlensing occurs when two objects appear to align in the night sky despite being at very different distances to an observer such as a telescope. The mass of the closer object bends the spacetime around it and directs the light from the background star along not straight lines but curved geodesics. This creates multiple images of the background star that are unresolvable to the observer and therefore appear as a single amplified source. The lightcurve of ZTF18abnbmsr in Figure 1.1 demonstrates what a microlensing event looks like when observed photometrically.

This phenomenon requires three essential properties of the two objects that appear aligned. First, the more distant object, or source, must be luminous so that there is light to be amplified. Second, the closer object, or lens, must possess mass so that it can sufficiently bend the spacetime around it. Third, the relative distances to the lens and source must be approximately the same order of magnitude for the apparent magnification of the source to be detectable. Given the vast emptiness that occupies most of the night sky, the chance alignment of two objects in this configuration is a decidedly rare event.

1.1.1 Formalism

In the ideal case where a source appears exactly behind a lens, the source's light will be focused into a ring with an Einstein radius (θ_E), determined by the mass of the lens (M_{lens}) and distances to the source (d_S) and lens (d_L)

$$\theta_E = \sqrt{\frac{4GM_{\text{lens}}}{c^2} \frac{d_S - d_L}{d_S d_L}}. \quad (1.1)$$

This is commonly condensed into an abbreviated form by defining the constant $\kappa = 4G/c^2$ and the lens-source relative parallax $\pi_{\text{rel}} = (1/d_L) - (1/d_S)$ resulting in

$$\theta_E = \sqrt{\kappa M_{\text{lens}} \pi_{\text{rel}}}. \quad (1.2)$$

When the source and object are not perfectly aligned, the Einstein radius defines the approximate angular scale of the microlensing effect. π_{rel} can also be scaled by the Einstein radius to get the Einstein parallax

$$\pi_E = \frac{\pi_{\text{rel}}}{\theta_E}. \quad (1.3)$$

The easiest-to-measure property of a microlensing event is the timescale of the source amplification, or the Einstein crossing time (t_E)

$$t_E = \frac{\theta_E}{\mu_{\text{rel}}}, \quad (1.4)$$

where μ_{rel} is the relative proper motions between the source and lens. This is approximately the time for which the luminous source will be amplified due to the magnification of the lens. The value μ_{rel} is often unknown although it can be statistically approximated or measured after years or even decades have passed for the source and lens to become individually resolvable.

1.1.2 Observation

The challenges with observing microlensing events are numerous. Microlensing events are very rare events occurring less than once in every million or so stars at any given time (B. Paczyński 1986). The average timescale for a microlensing event is around 20 days and must therefore be observed with nearly daily coverage to provide useful constraints. Variable stars often show variations in brightness that can appear similar to microlensing especially over the timescales of most microlensing surveys. Microlensing is most likely to occur in dense stellar fields where precision photometry is difficult to achieve.

Ideally photometric microlensing is achromatic because the bending of light by space-time is agnostic to wavelength. This can be used as a tool to discern microlensing from other variable stars. However stars in dense stellar fields often have neighboring stars with overlapping point spread functions. These stars that are neither the source nor the lens of the microlensing event contribute additional flux that reduced the magnitude of the observable amplification. The fraction of the observable flux originating from the source is the blend-source-flux-fraction

$$b_{\text{sff}} = \frac{F_S}{F_S + F_L + F_N}. \quad (1.5)$$

A significantly blended event, where a large fraction of the light in an observational aperture does not originate from the source, can appear similar to an event where the distance between the source and lens at closest approach is relatively far apart. Blending makes it challenging to measure even the simplest microlensing event parameters.

Work is also underway to measure the astrometric effects of microlensing Lu et al. (2016). The multiple images created during a microlensing event are asymmetrically located relative to the flux-weighted average location of the source and lens. This causes a wobble in the measured centroid location that when measured places constraints on the Einstein radius of the microlensing event. The effect is extremely small, often on the order of miliarcseconds occurring over months to years, making it only possible to measure with multiple epochs of high-precision astrometry using adaptive optics. This measurement, along with photometric measurements of the Einstein crossing time and the annual parallax, can measure the mass of the lens (Miyamoto & Yoshi 1995).

Chapter 2 includes, along with an estimate for the number of microlensing events that ZTF could observe, descriptions of how the observable population statistics for microlensing events differs from the more commonly observed Galactic bulge to the less observed Galactic plane. One of the advantages of executing an all-sky microlensing survey is the reduced

blend-source-flux-fraction found in the less dense Galactic plane as compared to the bulge. This work also outlines several additional challenges, such as uncertainty about the distance to sources in these plane fields that cannot be assumed to be located at approximately the same distance, as they can be when looking toward the bulge.

1.1.3 Black Holes

B. Paczyński (1986) proposed a method for detecting dark matter in the form of massive halo compact objects (MACHOs) using microlensing. He concluded that if dark matter is significantly comprised of MACHOs then one of every million or so star in the Magellanic Clouds could be lensed by a dark lens in a given moment. Surveys such as MACHO (Alcock et al. 2000), EROS-2 (Tisserand et al. 2007) and others have set limits on the amount of the Galactic Halo that could be found in MACHOs by following Paczyński’s suggestions. Their work has provided yet another probe in the search for the origin of dark matter, as shown in Figure 1.2.




Figure 1.2: Limits on the amount of dark matter in the galactic halo made of MACHOs, reproduced from Wyrzykowski et al. (2011a). The OGLE, EROS and MACHO surveys place these 95% confidence limits by converting the number of expected long-duration events compared with the number observed using Poisson statistics. The absence of events caused by lenses with large masses is evidence suggesting that MACHOs are not significant fractions of the dark matter halo. However recent work challenges these constraints after improving mass models (García-Bellido & Clesse 2017; Calcino et al. 2018).

Stellar mass black holes are known to be biased toward longer Einstein crossing times due to their larger masses (B. Paczyński 1986), albeit with significant challenges in performing

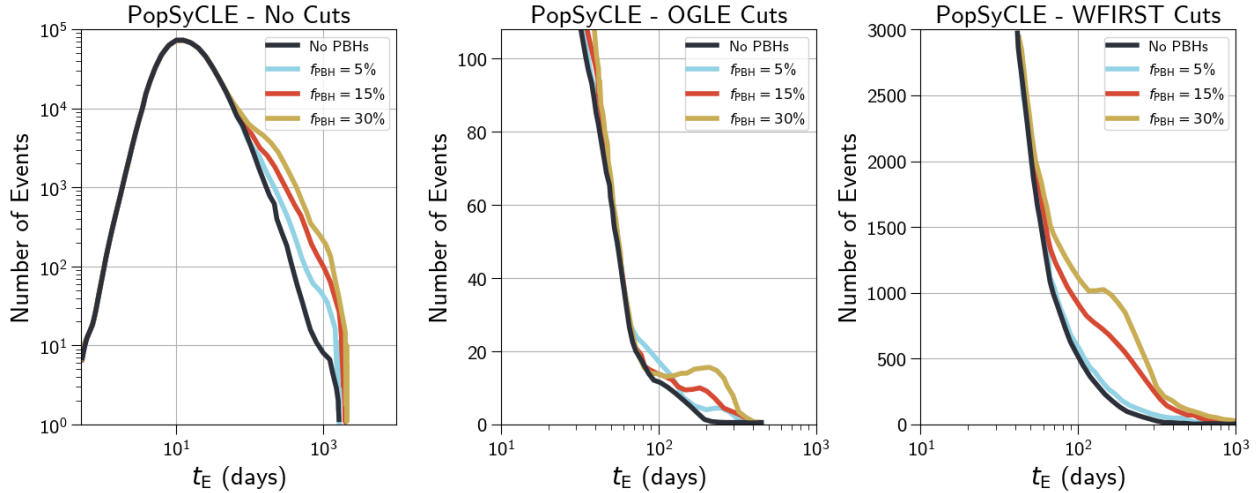


Figure 1.3: The distribution of Einstein crossing times has a measurable increase for long-duration events in the presence of black holes, reproduced from [Lu et al. \(2019\)](#). In this way black holes can be statistically detected in a survey before astrometric follow-up has been conducted.

this measurements ([Mao & Paczyński 1996](#)). This has lead surveys to search for evidence of black holes by looking for events with the largest Einstein crossing times ([Agol et al. 2002](#); [Wyrzykowski et al. 2016](#)). In order to constrain the microlensing parallax it had been thought that surveys should focus on those events with large amounts of parallax that could be well measured ([Bennett et al. 2002](#)). However more recent simulations have argued against this approach, showing that black hole events are biased toward small microlensing parallax ([Lam et al. 2020](#)). Therefore recent searches have looked for black holes in the region of long t_E - small π_E ([Golovich et al. 2020](#)).

Confirming a candidate black hole identified by photometric microlensing requires measuring the lens mass and placing an upper limit on the lens flux. This requires measuring either the source-lens relative proper motion ([Abdurrahman et al. 2021](#)) or the Einstein radius directly ([Lu et al. 2016](#); [Arredondo et al. 2019](#)) along with the easier to measure parameters such as the Einstein crossing time. However these techniques are extremely expensive requiring years of high-precision astrometric measurements on the world’s most precise telescopes. The astrometric signal is strongest in the months before and after the source star’s peak magnification. Services such as the OGLE Early Warning System ([Udalski et al. 2015b](#)) provide the microlensing community with lists of ongoing events in the Galactic bulge from which black hole candidates can be selected for real-time astrometric follow-up.

Chapter 5 outlines the steps taken to produce a list of black hole microlensing candidates from across the entire sky using the Zwicky Transient Facility. The population statistics of these events such as their Einstein crossing times can present evidence for the presence of black holes (Figure 1.3). These events can also be followed up with astrometric observation

near their peak amplification to measure their Einstein radius and therefore constrain their lens mass.

1.2 The Zwicky Transient Facility

The Zwicky Transient Facility (ZTF) is an all-sky optical synoptic survey operated on the 48-inch Samuel Oschin Telescope at Palomar Observatory (Bellm et al. 2019c; Graham et al. 2019). ZTF images are taken on a 47 square degree camera averaging $\sim 2.0''$ FWHM on a plate scale of $1.01'' \text{ pixel}^{-1}$. This extremely large camera enables ZTF to image all visible fields in the Northern Night sky every few nights, creating a massive optical history covering two and a half years of Phase-I operations (lasting until September 2020). The funding for this instrument is split between the National Science Foundation’s Mid-Scale Innovations Program (MSIP) and a private collaboration of international academic partners. This hybrid funding model results in some data products being served to the public immediately, while others are retained for a proprietary period before public release.

1.2.1 Surveys and Data Products

As a collaborative synoptic survey, ZTF is executing several observing surveys simultaneously. Three surveys from the past several years best suited for observing microlensing events are the Northern Sky Survey and the Galactic Plane Survey. These surveys covered a mixture of Galactic plane and field stars with cadences varying from one to three nights. Other surveys focusing on extra-galactic targets, coincidence with the Transiting Exoplanet Survey Satellite, and higher cadence observations in dense stellar fields were also conducted during these times. This produced a collection of observations with varying cadences and depths.

Every several months the ZTF partnership releases many of these observations as Public Data Releases. The releases include single epoch images, reference images, photometric catalogs and time series lightcurves for objects identified in reference image photometry. The latest Public Data Release 5 included over 22 million good-quality exposures covering over 97% of the sky in at least two filters. The lightcurves in this release contained over 1.7 billion sources with at least 20 observations taken in either the ZTF g-band, r-band or i-band filters. Data releases of this size are becoming commonplace in astronomy as the transition to “big data” has accelerated in recent years. This requires new approaches to finding rare astrophysical objects that can be applied at these enormous scales.

1.2.2 Data Analysis Challenges

Traditional methods for detecting microlensing events have relied on datasets designed to optimize their detection. The many different surveys simultaneously run on ZTF creates data with cadences varying from multiple images per hour to gaps of days or weeks. Exposures

within the datasets range from extremely dense Galactic bulge fields to sparse fields outside of the Galactic plane. Photometric catalogs are not optimized for any particular science case and can therefore perform poorly under these varying conditions. The statistical methods developed in Chapter 5 are flexible enough to handle these uncertainties when searching for microlensing events.




Figure 1.4: Long-wavelength *i*-band imagery on the Zwicky Transient Facility is contaminated with noise from atmospheric fringing. This correlated background noise is not present in *g*-band or *r*-band images. This requires a data reduction process tailored to removing these fringes while retaining a Gaussian background. Chapter 3 outlines an algorithm for removing these fringes that is now a standard stage in the ZTF data reduction pipeline.

One of the characteristics of microlensing events that can be used to constrain them are fitting microlensing models in multiple filters. However the *i*-band images in ZTF suffer from atmospheric fringes common to long wavelength optical imagery. Figure 1.4 demonstrates how a typical *i*-band image is contaminated with significant additional noise due to atmospheric fringes as compared to a *g*-band or *r*-band image. Chapter 3 describes a method for removing these fringes from images using Principal Component Analysis. This method significantly improves upon previous methods for removing fringes by creating single-epoch fringe bias images generated from pre-trained per-pixel models. This algorithm has been adopted as a fiducial step in the ZTF calibration pipeline and is run on all *i*-band images ever taken by the facility.

1.3 Thesis Outline

In this thesis, I outline the steps taken to transform the Zwicky Transient Facility into a microlensing machine capable of detecting black holes. This includes simulating the number of microlensing events that ZTF could potentially find, improving the data quality of ZTF images that can be used to find these events, and executing the first robust microlensing search pipeline built for a multi-purpose observing facility. ZTF is a predecessor to future

synoptic surveys such as the Vera C. Rubin Observatory and the Nancy Grace Roman Space Telescope. The scalable methods developed in this thesis have addressed many of the challenges that will arise when constructing microlensing pipelines on the massive datasets generated by these instruments.

In Chapter 2, I predict the number of microlensing events detectable by ZTF running resolved simulations throughout the Galactic plane and modelling the observational biases of the instrument. These simulations are also used to forecast the unique characteristics of microlensing events outside of the more commonly searched Galactic bulge. In Chapter 3, I significantly improve the photometric quality of ZTF i-band images by introducing a method for removing atmospheric fringes that otherwise mask the observability of faint sources. In Chapter 4, I calculate a 2-dimensional transformation for converting ZTF filters into a more common filters system to enable cross-calibrating microlensing events with historical measurements. In Chapter 5, I construct and execute the first microlensing detection pipeline ever conducted on a multi-year survey not originally designed for microlensing. This work serves as a blueprint for conducting future microlensing searches on large scale synoptic surveys of the future. In Chapter 6, I conclude the thesis.

Chapter 2

Gravitational Microlensing Event Statistics for the Zwicky Transient Facility

This is a replication of the article published as [Medford, Lu, Dawson, Lam, Golovich, Schlafly, & Nugent 2020a](#), ApJ, 897, 2, pg. 144.

Abstract

Microlensing surveys have discovered thousands of events with almost all events discovered within the Galactic bulge or toward the Magellanic clouds. The Zwicky Transient Facility (ZTF), while not designed to be a microlensing campaign, is an optical time-domain survey that observes the entire northern sky every few nights including the Galactic plane. ZTF observes $\sim 10^9$ stars in g-band and r-band and can significantly contribute to the observed microlensing population. We predict that ZTF will observe ~ 1100 microlensing events in three years of observing within 10° degrees latitude of the Galactic plane, with ~ 500 events in the outer Galaxy ($\ell \geq 10^\circ$). This yield increases to ~ 1400 (~ 800) events by combining every three ZTF exposures, ~ 1800 (~ 900) events if ZTF observes for a total of five years, and ~ 2400 (~ 1300) events for a five year survey with post-processing image stacking. Using the microlensing modeling software PopSyCLE, we compare the microlensing populations in the Galactic bulge and the outer Galaxy. We also present an analysis of the microlensing event ZTF18abhxmj to demonstrate how to leverage these population statistics in event modeling. ZTF will constrain Galactic structure, stellar populations, and primordial black holes through photometric microlensing.

2.1 Introduction

First proposed by [Einstein \(1936\)](#), gravitational lensing occurs when a massive object intersects the line of sight between an observer and a luminous source. The gravitational field of the intermediate object bends spacetime, acting as a lens and causing the appearance of multiple closely spaced images to an observer along this line of sight. When the massive lens and the luminous source are both stars, the multiple images of the source are separated by only microarcseconds. They are thus unresolvable and are therefore called microlensing ([Refsdal & Bondi 1964](#)). The photometric effect of these multiple images is an apparent amplification of the source’s brightness while the source crosses behind the lens. This phenomenon is called photometric microlensing.

Microlensing possess several distinct signatures unique among astrophysical transients that aid in their discovery. If the lens and source are assumed to be point sources and the observer remains approximately stationary, the photometric light curve is a rise in brightness followed by a symmetric fall in brightness of the same timescale ([B. Paczyński 1986](#); [Paczynski 1996](#)). This simple model is complicated by the motion of the Earth around the Sun which produces a parallax effect that perturbs the magnification depending on the time of the year that the event is observed and the location of the event in the sky ([Gould 1992](#)). Microlensing is ideally achromatic; however additional sources of light in the photometric aperture, or blending, can introduce differential color changes into the transient signal ([Stefano & Esin 1995](#)). Still, images taken in multiple filters containing an approximately equal increase in brightness serve as a key piece of evidence for claiming a microlensing detection.

Observable microlensing events occur almost entirely between two stars in the Milky Way (or a nearby galaxy) as the sources and lenses rotate around the center of the galaxy. The size of the apparent ring formed by the lensed source during a theoretical perfect alignment is called the Einstein radius, given by

$$\theta_E = \sqrt{\frac{4GM_L}{c^2} \left(\frac{1}{d_L} - \frac{1}{d_S} \right)}, \quad (2.1)$$

where M_L is the mass of the lens and d_L and d_S are the distance between the Sun and the lens and source, respectively. The Einstein radius is the approximate angular scale of a microlensing event in the case of a more realistic imperfect alignment between the source, lens, and observer. The centroid of the aperture’s flux will perturb during a microlensing event on a scale approximately equal to the Einstein radius. This effect, known as astrometric microlensing, is extremely difficult to measure. For a typical microlensing event in the Milky Way bulge, with a source located at eight kiloparsecs (near the center of the galaxy) and a lens halfway between the Earth and the source, a one solar mass star would produce an Einstein radius and astrometric perturbation of approximately one milliarcsecond.

The time for the luminous source to pass across the Einstein radius in the reference frame of the gravitational lens is the Einstein crossing time, given by

$$t_E = \frac{\theta_E}{\mu_{\text{rel}}}, \quad (2.2)$$

where μ_{rel} is the relative proper motion between the source and lens as seen by the observer. This observable can be measured by fitting a photometric lightcurve with a microlensing model and identifying the timescale over which the magnification of the signal increases and then subsequently decreases. A typical microlensing event in the Milky Way bulge has an Einstein crossing time of approximately 20 days (Sumi et al. 2011; Wyrzykowski et al. 2015; Mróz et al. 2017).

Microlensing detections have resulted in many significant discoveries in the past few decades. Galactic models have been constrained by looking at the population statistics of microlensing events including spatial and Einstein crossing times distributions (Aubourg et al. 1993; Kerins 1995; Wyrzykowski et al. 2015; Navarro et al. 2020). Microlensing has been used to discover and constrain exoplanet populations (Cassan et al. (2012); See Gaudi 2012 for review) and the Nancy Grace Roman Space Telescope (formally named the Wide Field Infrared Survey Telescope) aims to significantly increase the number of exoplanets found through microlensing by ~ 1400 (Calchi Novati 2018; Penny et al. 2019). Looking for dark matter in the Milky Way halo using microlensing was originally proposed by B. Paczyński (1986), with constraints on the contribution of primordial and astrophysical black holes to the dark matter mass halo successfully executed in the years since (Alcock et al. 2001; Afonso et al. 2003; Wyrzykowski et al. 2011a; Niikura et al. 2019). More recent work proposes detecting free floating black holes through photometric microlensing alone (Lu et al. 2019), as well as combining these observations with astrometric measurements (Lu et al. 2016; Kains et al. 2016; Rybicki et al. 2018).

Microlensing has been traditionally dominated by surveys conducted in the Galactic bulge (Sumi et al. 2013; Udalski et al. 2015a; Navarro et al. 2017; Kim et al. 2018a; Mróz et al. 2019) to maximize their yields, as well as the Magellanic clouds (Alcock et al. 2000; Tisserand et al. 2007; Wyrzykowski et al. 2011a) and M31 (Novati et al. 2009, 2014) to increase the relative probability of detecting a dark matter lens relative to a stellar lens. The microlensing event rate is proportional to the number of luminous sources in the field of view and the mass density along the line of sight (Calcino et al. 2018), favoring pointing towards the Galactic bulge over other lines of sight in the Galaxy if attempting to maximize the microlensing event rate. The measurement of optical depths to microlensing by EROS-2 (Hamadache et al. 2006), optical depth and event rate by both MOA-II (Sumi et al. 2013) and OGLE-IV (Mróz et al. 2019), and the study on Galactic longitude dependence by VVV (Navarro et al. 2020) are all calculated in the bulge, containing fields entirely located within Galactic longitudes of $-10^\circ < \ell < 10^\circ$. The EROS-2 spiral arm surveys (Derue et al. 2001; Rahal et al. 2009) searched for microlensing at Galactic longitudes $|\ell| > 10^\circ$ but were only able to find 27 microlensing events among the 12.9 million stars observed over seven years. Synoptic surveys (those with large footprints and wide fields of view that repeatedly observe the same fields over long stretches of time) will discover more microlensing events outside of the Galactic bulge in the outer Galaxy, and even outside of the Galactic plane, than ever before. Sajadian & Poleski (2019) predicts that the Vera C. Rubin Observatory (previously referred to as the Large Synoptic Survey Telescope) could observe anywhere from 34,000 microlensing events in its first year to 795 events per year over ten years depending on the observing strategy,

showcasing the potential for an all sky survey to significantly grow the total population of microlensing events depending on the observing strategy that is implemented. Mróz et al. (2020b) lists 30 likely microlensing events discovered in the first year of the Zwicky Transient Facility’s Galactic Plane Survey, and our work suggests that there remain many more events still to be discovered. Photometric filters which focus on efficiency and scale (Price-Whelan et al. 2014) or introduce novel machine learning techniques that can easily scale (Godines et al. 2019) may be the keys to discovering these additional events.

In this paper we present the Zwicky Transient Facility’s opportunity to conduct the first all sky microlensing survey and the potential scientific contributions such a survey could enable. In Section 2.2, we describe the Zwicky Transient Facility instrument and data. In Section 2.3, we estimate the total number of microlensing events that ZTF could discover in its first three years and methods for increasing these yields. In Section 2.4, we explore the difference in population statistics for microlensing events in the outer Galaxy as compared to the Galactic bulge. In Section 2.5, we demonstrate a proof of principle for how to use the microlensing simulation software PopSyCLE (Lam et al. 2020) to model events in the outer Galaxy and we conclude in Section 3.5.

2.2 The Zwicky Transient Facility

The Zwicky Transient Facility (ZTF) is an optical time-domain survey that has been operating on the 48-inch Samuel Oschin Telescope at Palomar Observatory since March 2018 (Bellm et al. 2018). ZTF’s camera covers 47 square degrees in a single exposure, enabling coverage of the entire visible Northern sky every few nights in ZTF g-band, r-band and i-band filters with an average 2.0" FWHM on a plate scale of 1.01" pixel⁻¹. ZTF produces a real-time alert stream triggered by transient event detections on difference images processed by the IPAC facility (Patterson et al. 2018). In addition to these alerts, the ZTF collaboration routinely produces public data releases which contain, among other data products, lightcurves assembled from single image point spread function (PSF) photometry for every star in the northern sky which appears in a deep co-added reference image (Masci et al. 2018). Reference images are ideally constructed from 40 individual exposures resulting in an approximate r-band limiting magnitude of 22.6, although weather and visibility produces variable results. ZTF’s observing time is split between public observations (funded by the National Science Foundation’s Mid-Scale Innovations Program or MSIP) and partnership observations, which are held in a proprietary period for collaboration members of the survey. The i-band filter is used only for partnership observations and is thus absent from this analysis.

ZTF has several observing surveys covering the northern sky in r-band to a five-sigma depth of approximately $m_{lim,r} = 20.6$ and g-band to a depth of approximately $m_{lim,g} = 20.8$ every few nights (Bellm et al. 2018, 2019a). The Northern Sky Survey observes the entire visible sky north of -31° declination in both g-band and r-band with a three night cadence and has been executed since 2018 March. The Galactic Plane Survey (Prince & Zwicky




Figure 2.1: ZTF Public Data Release 1 contains 1.7×10^9 lightcurves assembled from sources in 3.4×10^6 single-exposure images taken in g-band and r-band. Top: The number of lightcurves in each filter containing a given number of epochs, as well as the total for the two filters combined. Most lightcurves in the catalog are in fact single source detections with no subsequent observations most likely resulting from optical artifacts, moving solar system objects or faint transient sources. Bottom: The number of lightcurves with observations more than the threshold number of epochs, as well as the total for the two filters combined. For example, there are 7.8×10^7 r-band lightcurves and 1.4×10^8 g-band lightcurves with more than 60 observations, for a total of 2.2×10^8 lightcurves. Computational costs effect how many lightcurves can be searched for microlensing events and determines the minimum number of observations a lightcurve must contain. It should be noted that the ZTF data reduction pipeline treats sources detected at the same location in the sky but in the different filters as separate sources.

[Transient Facility Project Team 2018](#)) observes all Galactic plane fields ($-7^\circ < b < 7^\circ$) visible on a given night in both bands when the Galaxy is visible from Palomar Observatory. In total ZTF observes over 2000 square degrees in a combination of g-band, r-band, and i-band exposures every night.

The Northern Sky Survey and the Galactic Plane Survey are public surveys producing a real-time alert stream generated by transient detections on difference images (Patterson et al. 2019). Science images of these observations are released at regular intervals, as well as a variety of data products including lightcurves assembled from single epoch photometry. These surveys generate well sampled lightcurves for hundreds of millions of stars with non-uniform sampling due to field visibility and weather losses. Additionally, the ZTF partnership conducts a high cadence survey in the Galactic plane with 30 second images taken on the same fields for several weeks that are released on a more infrequent basis. All of these surveys provide excellent datasets for observing microlensing events due to short cadences and images taken in multiple filters.

On 2019 May 8, ZTF released Public Data Release 1 (DR1) containing 1.7×10^9 lightcurves assembled from sources in 3.4×10^6 single-exposure images taken in g-band and r-band for observations taken between 2018 March 17 and 2018 December 31 . To generate these lightcurves, ZTF ran PSF photometry on both individual exposures and reference images constructed from co-adding science exposures. Sources which appeared in the reference image catalogs were used as seeds for the construction of lightcurves. Sources which appeared in the photometric catalogs of individual science images at the location of a source from the reference catalog were appended to their respective lightcurves. The lightcurve catalogs from DR1 contain over 8.1×10^8 lightcurves with $N_{obs} \geq 20$ from across the northern sky (Figure 2.1). Both releases also include science images, reference images, subtraction images, photometric catalogs and other data products.

2.3 ZTF Microlensing Estimate

ZTF can be used to detect a significant number of microlensing events due to its large sky coverage, multiple filters, and repeated observations. What follows is an approximation for the number of events that ZTF could discover in its three years of operations. Here we calculate the approximate number of microlensing events (N_{events}) through combining the duration of the ZTF survey in years (T_{obs}), the number of sources ZTF observes (N_{stars}^{ZTF}), and the observable microlensing event rate per star per year (Γ_{obs}):

$$N_{events} = \Gamma_{obs} \cdot N_{stars}^{ZTF} \cdot T_{obs} . \quad (2.3)$$

The number of sources is counted from ZTF reference image photometric catalogs; however, the microlensing event rate must be estimated from simulations.

We utilize PopSyCLE to estimate microlensing event rates at different Galactic latitudes and longitudes. PopSyCLE, or Population Synthesis for Compact object Lensing Events (Lam et al. 2020), is a recently released open-source code that uses galaxy modeling and stellar population synthesis to generate realistic microlensing populations, including compact object sources and lenses. These simulations are generated along specified lines of sight in the Galaxy using stars from Galactic models (Robin et al. 2003) produced by Galaxia (Sharma et al. 2011) and compact objects determined by initial-final mass relations (Kalirai et al.

2008; Sukhbold et al. 2016; Raithel et al. 2018) calculated in `PyPopStar` (Jr. et al. 2020). Estimating event rates with `PopSyCLE` provides us more physical insight into the populations of stars and compact objects undergoing microlensing than would be deduced from using analytic expressions. Simulations were run using the `PopSyCLE` v3 galaxy model (Lam et al. (2020): Appendix A, a model which is demonstrated to accurately produce event rates in various bulge fields when compared to (Mróz et al. 2019). We note that `PopSyCLE` v3 galaxy model does not reproduce observed stellar densities in the Galactic field. However our paper adopts a relative stellar density fraction (See Section 2.3.2) that corrects for this discrepancy between observed stellar densities and modelled stellar densities. This ensures that our estimate of the microlensing event rate per star are accurate.

Executing a `PopSyCLE` simulation, especially in the high stellar densities of the Galactic bulge, incurs significant computational cost and cannot therefore be performed at every ZTF field location across our estimate’s footprint. The accuracy of our estimate is limited by the discrete number of simulations carried out across the Galactic plane over which we interpolate the observable event rate. Bootstrapping of the discrete simulations indicate that the precision of our event rate estimates at each location vary by approximately 10%. The accuracy of the predicted event rate is also limited by systematic errors in the Galactic model implemented in `PopSyCLE` that we did not explore, which are known to contribute to errors in Galactic microlensing modelling (Evans & Belokurov 2002).

2.3.1 Event Rate: Γ_{obs}

The event rate in this estimate, Γ_{obs} , is

$$\Gamma_{\text{obs}} = \frac{N_{\text{events, detected}}^{\text{PopSyCLE}}}{\frac{N_{\text{stars}}^{\text{ZTF}}}{N_{\text{stars}}^{\text{PopSyCLE}} \Big|_{\text{area}}} \cdot T_{\text{obs}} \cdot N_{\text{stars}}^{\text{PopSyCLE}}} \cdot f_{\text{visibility}}. \quad (2.4)$$

The event rate is found at each sky location by dividing the number of simulated events detected $N_{\text{events, detected}}^{\text{PopSyCLE}}$ by the total number of stars in our `PopSyCLE` simulation $N_{\text{stars}}^{\text{PopSyCLE}}$ and the simulated survey duration T_{obs} . In order to account for observational effects that aren’t simulated by `PopSyCLE`, such as blending and incompleteness in the number of stars, we then apply a correction factor $N_{\text{stars}}^{\text{ZTF}}/N_{\text{stars}}^{\text{PopSyCLE}} \Big|_{\text{area}}$ that is the ratio of stellar densities in `PopSyCLE` and on-sky from ZTF. This ratio is less than one across most of the Galactic plane where ZTF sees fewer stars than `PopSyCLE` due to these effects. However at the smallest galactic latitudes the ratio can be larger than one if the extinction is overestimated and there are more ZTF stars than the model predicts. However these are locations where our event rate is near zero and does not largely effect our final estimates. The rate is then corrected by a visibility completeness term $f_{\text{visibility}}$ that down-weights the number of microlensing events from fields proportional to their visibility by ZTF. Both the relative stellar density fraction and the visibility completeness are discussed in more detail below. We note that our predicted event rate is specifically for those events that are observable by ZTF. This would

be equivalent to observational event rates reported before the completeness correction often applied by other work (Sumi et al. 2013; Wyrzykowski et al. 2015; Mróz et al. 2019).

The number of events detected ($N_{\text{events, detected}}^{\text{PopSyCLE}}$) is calculated by implementing observational cuts similar to Sumi et al. (2011); Mróz et al. (2017) in the manner outlined in Lam et al. (2020). However PopSyCLE, which does not create and sample individual lightcurves, must artificially calculate some of the observational criteria of surveys. For example, when analyzing millions of lightcurves, microlensing surveys must determine whether an increase in flux is significant. Significant bumps in flux are with three consecutive measurements are above 3σ of the baseline flux (e.g. Mróz et al. (2017) Extended Data Table 3, Sumi et al. (2011) Table S2). A microlensing events in PopSyCLE is deemed to have a significant bump in flux if

$$F_{\text{peak}} - F_{\text{base}} > 3\sigma_{\text{base}} \approx 3\sqrt{F_{\text{base}}},$$

where F_{peak} and F_{base} are the peak and baseline flux, respectively. Calculations on non-variable ZTF lightcurves of $\sqrt{F_{\text{base}}}$ found it to be equal to or larger than σ_{base} on almost all objects, making this version of the significant bump requirement a conservative estimator. To calculate the error on the peak and baseline flux, knowledge of the zero point magnitude m_{ZP} is needed. m_{ZP} is the magnitude that corresponds to a single count in the detector $F_{\text{ZP}} = 1$. Thus the flux-magnitude relation can be written

$$m - m_{\text{ZP}} = -2.5 \cdot \log_{10}(F).$$

m_{ZP} is calculated for each simulated filter and the fluxes are assumed to have Poisson errors.

Table 2.1 contains the complete list of our selection criteria. Both the survey duration ($T_{\text{obs}} \in [1, 3, 5]$ years) and minimum baseline magnitude ($19 \text{ mag} < m_{\text{lim,r}} < 22 \text{ mag}$) selection criteria are calculated for the stated range of values. The choice to calculate our estimate for multiple survey durations is discussed in Section 2.3.2. Section 2.3.4 discusses applying post-processing image stacking to increase the total number of observable microlensing effects. We calculate this effect by increasing the minimum baseline magnitude accordingly. Events are required to have an Einstein crossing time, source flux fraction, and impact parameters within the limits of the stated values. The magnitude amplification Δm is calculated by subtracting baseline magnitude from the source, lens and all neighboring stars from the magnitude at maximum amplification and must also be greater than the stated cutoff value. All of our calculations are performed with the ZTF r-band filter by transforming PopSyCLE’s UBv photometry into the ZTF filter system (Medford et al. 2020c).

The observational cuts in Table 2.1 are chosen to result in a conservative estimate for the number of detectable microlensing events. While the average full-width half-maximum of ZTF is closer to $1.5''$, we set the seeing disk radius to the confusion limit measured in our densest fields. Setting the seeing disk radius $\theta_{\text{blend}} = 2.25''$ places more neighboring stars into the observational aperture and therefore increases the baseline flux of a microlensing event in a field with high stellar density. This makes the event less likely to be detected because (1) a larger baseline flux requires a larger peak flux in order to have a significant bump, (2) an event with a larger baseline flux will have smaller magnitude amplification,

Table 2.1: PopSyCLE Observational Cuts

Parameter/Criteria	Value
Filter	ZTF r-band
Seeing disk radius, θ_{blend} [arcsecond]	2.25
Minimum Einstein Crossing Time, t_E [days]	≥ 3
Minimum baseline magnitude, m_{base} [mag]	$19 < m < 22$
Maximum impact parameter, u_0	≤ 1
Removal of low-amplitude events, Δm [mag]	≥ 0.1
Removal of highly blended events, $b_{\text{ff,r}}$	≥ 0.1
Survey duration, T_{obs} [yrs]	1, 3, 5
Significance of bump, $F_{\text{peak}} - F_{\text{base}}$	$> 3\sigma_{\text{base}}$

Observational cuts applied to PopSyCLE microlensing candidate catalogs to simulate the ZTF survey, including choosing a filter and seeing disk radius to match the instrument. Limiting magnitudes are set to a range of values to determine the effect of post-processing stacking on the final event rate. Survey durations are set to one, three and five years to measure the effect of extending the ZTF survey. See [Lam et al. \(2020\)](#) for more details on the implementation of each cut.

and (3) a larger neighbor flux decreases the source-flux-fraction. All of these effects lower the observable event rate in the Galactic bulge where more crowding occurs due to higher stellar densities.

The number of stars in the simulation ($N_{\text{stars}}^{\text{PopSyCLE}}$) results from the simulation’s line of sight and the area of each simulation, which ranged from between 0.33 square degrees to 10 square degrees. There must also be a relative stellar density fraction ($N_{\text{stars}}^{\text{ZTF}}/N_{\text{stars}}^{\text{PopSyCLE}}|_{\text{area}}$) applied to the number of PopSyCLE stars to account for blending and the discrepancies between the PopSyCLE Galactic model and our observations. PopSyCLE generates many faint stars that appear in a ZTF aperture as a single source. Failing to account for this effect would result in an artificially low event rate by over-counting the total number of observable stars. We therefore calculated the ratio of ZTF stars from reference images and PopSyCLE stars that overlap in the same area on the sky for each magnitude in our range of minimum baseline magnitudes.

One might note that the number of ZTF stars ($N_{\text{stars}}^{\text{ZTF}}$) and the number of PopSyCLE stars ($N_{\text{stars}}^{\text{PopSyCLE}}$) both appear twice in Equations 2.3 and 2.4 and conclude that these terms can both be cancelled. If simulations were able to be carried out at all locations across the Galactic plane this would be correct because the number of events detected ($N_{\text{events, detected}}^{\text{PopSyCLE}}$) is itself an accurate measure of the number of events ZTF could detect toward that line of sight. However our strategy of constructing an interpolated grid of event rates requires that we convert the number of events detected into a rate per star. This allows us to multiply the interpolated event rate density ($\text{star}^{-1} \text{year}^{-1}$) by the stellar density (deg^{-2}) to calculate




Figure 2.2: The observable microlensing event rate density (top), stellar density (middle), and event density (bottom), for a three year survey of standard 30 second exposures. The event rate here is a detectable event rate, calculated by applying observational cuts to the PopSyCLE catalogs and scaling the number of sources in PopSyCLE to ZTF reference images. The finite grid of PopSyCLE runs, shown with their respective areas as red circles in the first subplot, creates lower resolution in the event rate density than the stellar density and results in an event density that maintains some of these lower resolution features. The gray area of the Galactic plane are regions which are not sufficiently visible to ZTF to render an estimate. The r-band limiting magnitude for this estimate was set at $m_{lim,r} = 20.6$ magnitudes.

the event density ($\text{year}^{-1} \text{deg}^{-2}$).

The visibility completeness ($f_{\text{visibility}}$) is determined for each field by simulating observation of that field throughout the year and calculating the fraction of nights per year that the field is visible for more than 30 minutes at an airmass less than 2.1. The event rate for a field is down-weighted by this fraction because only events that are observed during peak would be detected as microlensing events. The ZTF Northern Sky Survey and Galactic Plane survey ensure that a Galactic plane field that is visible will be observed and therefore this simulated fraction accurately represents the relative fraction of microlensing events that will be observed to peak within the survey duration of ZTF.

2.3.2 Number of Stars: $N_{\text{stars}}^{\text{ZTF}}$, Survey Duration: T_{obs}

The ZTF DR1 contains reference images photometric catalogs constructed from deep co-additions. We count the number of sources in each field, using the range of minimum baseline magnitudes as a limiting magnitude cut on the catalog. For each of these magnitude cuts, we generate an interpolated stellar density map.

ZTF has a planned operation timeline of three years with almost two years of operations already completed. Longer surveys are able to observe events with longer Einstein crossing times, creating a non-linear increase in the number of observable events with increasing survey duration. Our estimate was performed with a T_{obs} equal to one, three, and five years in order to demonstrate the increased yields in future ZTF data releases, as well as the benefit of continuing operations beyond the planned operation timeline. The PopSyCLE simulated survey duration was set to the same time in order to remove long duration microlensing events from the observable event rate that could not be detected in the duration of the survey.




Figure 2.3: The total number of microlensing events observable by ZTF at different limiting magnitudes for one year (red), three years (yellow), and five years (blue) in the visible Galactic plane (top) and the outer Galaxy (bottom). ZTF will observe ~ 1100 events over three years of operation at a r-band limiting magnitude of 20.6 (vertical black), with ~ 500 of these events occurring in the outer Galaxy ($l \geq 10^\circ$). If every three images are stacked together before generating photometric catalogs, the limiting magnitude would increase to 21.2 magnitudes (vertical dashed black) and would increase the yield to ~ 1400 events over three years, with ~ 800 events in the outer Galaxy. This stacking strategy would result in a cadence of three to five days. The total number of events observed would increase to ~ 2400 if the ZTF survey were extended to five years and this image stacking procedure were implemented, with ~ 1300 events in the outer Galaxy.

2.3.3 Interpolation Across the Galactic Plane

We ran PopSyCLE simulations and calculated stellar counts from ZTF reference images for fields in the Galactic plane visible to ZTF, at galactic longitudes from $100^\circ > \ell > 0^\circ$ and galactic latitudes from $-10^\circ < b < 10^\circ$ (Figure 2.2). Preliminary investigation suggested that extending the search to $|b| > 10^\circ$ and $\ell > 100^\circ$ would not significantly increase the predicted yield of microlensing events, although ZTF will observe these fields. The locations of our PopSyCLE simulations roughly cover the morphology of the Galactic plane and were used to create a linear interpolation of the event rate density ($\text{star}^{-1} \text{year}^{-1}$) and stellar density (deg^{-2}). PopSyCLE simulations were run at different sizes depending on their sky location in order to strike a balance between computational runtime and statistically significant numbers of microlensing events. Simulations away from the Galactic bulge were run on patches ranging from 1 deg^2 to 10 deg^2 , making the observable microlensing event rates at these locations an average over the simulation’s field of view. Simulations in the Galactic bulge were executed with an area of 0.33 deg^2 . Interpolating over the Galactic plane required choosing a scheme that accurately reflected the dynamic range of the stellar density, which we expect to be an approximate tracer of the event rate. We therefore choose to apply a linear interpolation and nearest extrapolation to our grid of event rates. Our sparse sampling is subject to interpolation errors that could effect our final results by up to a factor of two.

The location of Mount Palomar in the northern hemisphere limits the visibility to fields in the Galactic bulge closest to the Galactic center. The lack of data in these fields prevents us from making a measurement of the number of stars because too few exposures were taken in these fields to generate reference images. However, individual images of these fields have been taken by ZTF and some of fields are expected to have reference images by the end of the telescope’s three year lifespan. Microlensing predictions and searched can be recalculated after the completion of ZTF to increase their accuracy and yields.

2.3.4 Results: ZTF Microlensing Event Statistics

ZTF will observe ~ 1100 events over its fiducial three years of operation, assuming an r-band 5σ limiting magnitude of $m_{\text{lim,r}} = 20.6$ (Figure 2.3). ~ 600 events occur in the Galactic bulge ($\ell < 10^\circ$) where both the event rate and stellar density are large. This appears to validate the observing strategy taken by most microlensing

Table 2.2: Description of Fiducial Microlensing Simulations

Property	Inner Galactic Bulge	Outer Galactic Bulge	Outer Galaxy
Galactic Longitude ℓ	2.0°	6.0°	45.2°
Galactic Latitude b	1.0°	3.0°	4.9°
PopSyCLE Area	0.33 deg ²	0.33 deg ²	10 deg ²
PopSyCLE Extinction in ZTF r-band at 8 kpc	6.6 mag	2.4 mag	1.8 mag
ZTF Stellar Density at $m_{\text{lim},r} = 20.6$ mags	2.76×10^7 deg ⁻²	5.01×10^6 deg ⁻²	1.76×10^5 deg ⁻²

The three fields were chosen to demonstrate the differences in microlensing populations between the Galactic bulge and the outer Galaxy. The Galactic bulge fields represent the range of typical observations in the bulge with significantly higher stellar densities and extinctions than a field in the outer Galaxy. The Galactic bulge fields are smaller in order to be computationally tractable, while the outer Galaxy field is larger to generate a statistically significant numbers of microlensing events.

campaigns to observe in the Galactic bulge where the event rate is highest. However ~ 500 events occur throughout the outer Galaxy ($\ell > 10^\circ$) despite the drop-off in event rate and stellar density at larger Galactic longitudes. This is driven by the increased efficiency at detecting events further out in the plane (Sajadian & Poleski 2019) where reductions in crowding and consequently less confusion from neighboring stars in the photometric aperture make it easier to detect events relative to the bulge. The yields in the outer Galaxy are also increased due to the much larger footprint it covers compared to the bulge. Few microlensing events have been found at such large Galactic latitudes (Nucita et al. 2018; Dong et al. 2019; Wyrzykowski et al. 2020). Extending the survey duration of ZTF to five years would increase the number of detectable events to ~ 1800 , with ~ 900 events occurring in the outer Galaxy. Increasing the lifetime of the survey captures more of the long duration events particularly present at the larger Galactic longitudes, as well as increasing the number of short duration events across the entire plane.

The ZTF Northern Sky Survey and Galactic Plane Survey take 30 second exposures with a cadence of one to three days across the Galactic plane. The shift in the distribution of Einstein crossing times discussed in Section 2.4 means that most microlensing events would still be observable with a cadence of three to five days. Surveys such as the ZTF Uniform Depth Survey (Goldstein et. al. in prep) are creating photometric catalogs from co-additions of science images that simulate a deeper and longer cadence survey than ZTF’s current operations. Combining every three observations on ZTF would increase the r-band limiting magnitude to 21.2 magnitudes, increasing the three year yield to ~ 1400 events (~ 800 events in the outer Galaxy), with ~ 2400 (~ 1300) microlensing events observable if ZTF were extended to five years.

We stress here that the majority of these microlensing events will occur outside of the Galactic bulge and therefore beyond the footprint of most previously conducted microlensing campaigns. This presents the opportunity to constrain Galactic models and measure stellar population statistics in ways previously not possible with gravitational microlensing. While

our method does not make extremely precise predictions, it does demonstrate that executing a microlensing survey with ZTF will yield significant numbers of microlensing events through the less explored Galactic plane.

2.4 Microlensing Population Properties in the Outer Galaxy ($\ell \geq 10^\circ$)

Simulations of microlensing generated by PopSyCLE at these larger Galactic longitudes predict significant differences in the population distributions as compared to microlensing events the Galactic bulge. To highlight some of the difference in the microlensing populations at these different locations, we selected several fiducial fields to compare against each other. Analysis was performed in (1) the inner Galactic bulge, (2) the outer Galactic bulge, and (3) the outer Galaxy. Details of the characteristics of these fields can be found in Table 2.2. These fields are not meant to serve as representative of the Galactic bulge or outer Galaxy in their entirety, but were chosen in order to highlight the significant differences that can be found between the microlensing populations at different locations in the Galaxy. Such differences must be examined in order to properly model microlensing events and measure the physical parameters of a microlensing event. We demonstrate these effects on modeling in Section 2.5.

2.4.1 Distance Ratio

The distance ratio between the lenses and sources (d_L/d_S) is largely determined by the mass density along a given line of sight in the galaxy. Therefore it should not be surprising that the distribution of the distance ratio is different along different lines of sight (Figure 2.4). The average distance ratio towards the Galactic bulge is approximately 0.8, with sources in the inner bulge appearing at slightly larger distances. The distribution of distances to sources and lenses toward the outer Galaxy is significantly different, with a distance ratio peaking at approximately 0.25. The difference in these two distributions is driven by the different distance distributions of both the sources and the lenses. Sources and lenses towards the Galactic bulge are almost entirely located in the bulge ($\sim 6\text{--}11$ kiloparsecs away), while the number of sources in the outer Galaxy increases approximately linearly at further distances.

2.4.2 Einstein Crossing Time, Einstein Radius, and Relative Proper Motion

A commonly noted difference between microlensing populations in the Galactic bulge and the outer Galaxy is the distribution of Einstein crossing times (Sajadian & Poleski 2019) and the trend toward longer Einstein crossing times at larger Galactic longitudes (Mróz et al. 2019). We find a similar trend, with lines of sight further out along the outer Galaxy having




Figure 2.4: The location of microlensing sources in the galaxy against their distance ratio in the inner bulge (green), outer bulge (red) and in the outer Galaxy (purple), with histograms on both axes showing the marginalized distributions of the parameters. Events in the direction of the Galactic bulge have lenses and sources almost entirely located in the bulge ($\sim 6\text{--}11$ kiloparsecs away). The outer Galaxy events are more evenly spread out in source distance, with an approximately linear increase in sources at further distances. This results in an overall lower average distance ratio that must be appropriately used as a prior for any microlensing modeling in the outer Galaxy.

larger crossing times (Figure 2.5), averaging approximately 25 days in the bulge and almost 80 days in the outer Galaxy. This divergence is driven by the difference in relative proper motions and the Einstein radii between the two populations (Figure 2.6). The events in the bulge have mostly small Einstein radii and large relative proper motions, both pushing the Einstein crossing time toward smaller values (Equation 2.2). The opposite is found in the outer Galaxy, where lenses with large Einstein radii are crossed by luminous sources at relatively slower speeds.

Events in the Galactic bulge are difficult to measure astrometrically due to their smaller Einstein radii caused by the relatively similar distances to their sources and lenses as compared to the outer Galaxy (Equation 2.1). However microlensing events in the outer Galaxy will be easier to measure astrometrically due to their larger Einstein radii, with a significant number of events having radii larger than one milliarcsecond. Astrometric measurement is




Figure 2.5: The distribution of the Einstein crossing time in the inner bulge (green), outer bulge (red) and in the outer Galaxy (purple), with histograms on both axes showing the marginalized distributions of the parameters. Both Galactic bulge fields have an average Einstein crossing time of approximately 25 days, in alignment with previous work. However the outer Galaxy distribution averages around 80 days and stretches out beyond 1000 days in far excess of the Galactic bulge fields, with almost no events having an Einstein crossing time shorter than 10 days. Surveys can afford a longer observational cadence when searching for microlensing in the outer Galaxy due to this shift in the Einstein crossing time distribution.

a key method for breaking the mass-distance degeneracy that often plagues microlensing modeling. Lens masses will be better able to be constrained in the galactic plane because of these larger Einstein radii. It should also be noted that the decrease in relative proper motion will make it harder to observe these events with high resolution follow-up, which can determine the contribution to the aperture flux originating from neighbors and possibly observe source-lens separation after long periods of time.

2.4.3 Extinction

In order to infer the absolute magnitude and therefore spectral type of a microlensing source and lens we require an estimate of the extinction to both. This is difficult in the inner bulge due to large amounts of extinction that can be significantly different between the source and the lens. PopSyCLE uses the color excess values from the [Schlegel et al. \(1998\)](#) 3-D dust maps and the [Damineli et al. \(2016\)](#) reddening law to calculate interstellar extinction. [Lam et al. \(2020\)](#) Appendix B outlines how this results in accurate magnitudes and colors for stars throughout the bulge and greater Galactic plane.

Figure 2.7 shows the r-band extinction to sources and lenses in our three fields with significantly more extinction occurring in the inner Galactic bulge than the other fields as is expected. Extinction toward the inner Galactic bulge varies between five and nine

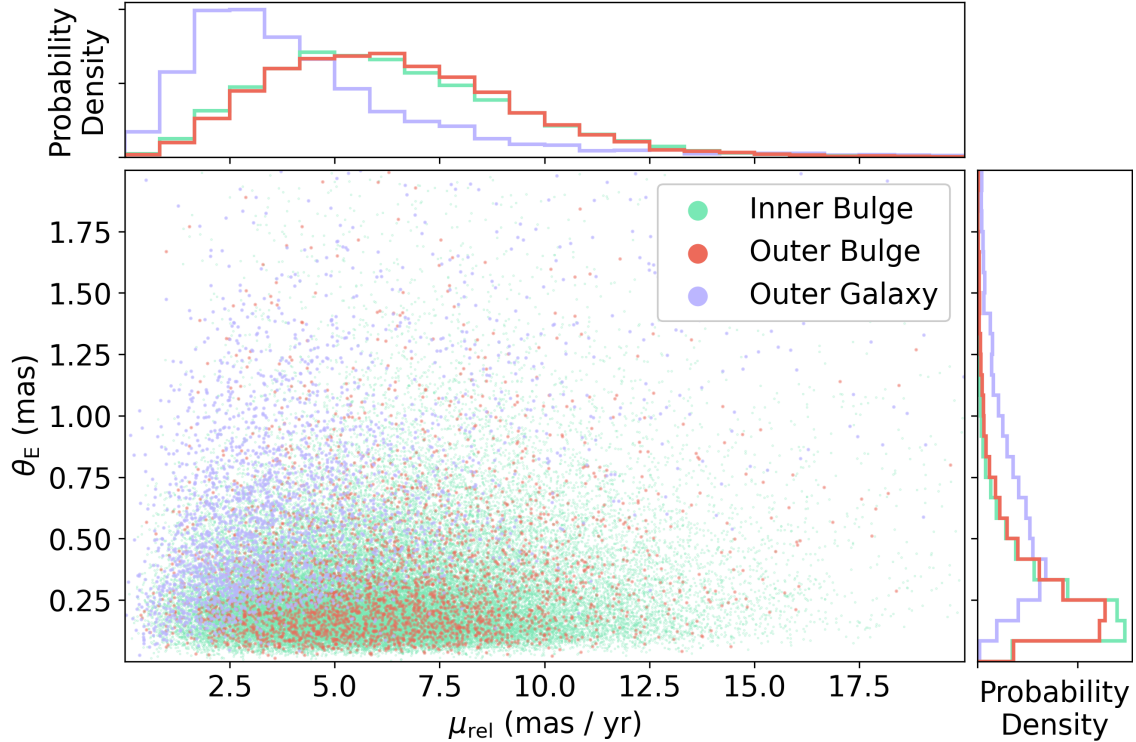


Figure 2.6: The size of the Einstein lens radii against the relative proper motions between the sources and lenses in the inner bulge (green), outer bulge (red) and in the outer Galaxy (purple), with histograms on both axes showing the marginalized distributions of the parameters. Microlensing events in the outer Galaxy have longer Einstein crossing times than those in the bulge due to their shorter relative proper motions and larger Einstein radii. The increased Einstein radii of outer Galaxy events makes them easier to follow up astrometrically in order to break the microlensing mass-distance degeneracy. However their slower relative proper motions results in a longer time before sources and lenses are resolvable on the sky due to separation.

magnitudes, with sources and lens having a difference of zero to four magnitudes despite their relatively equal distances. The outer Galactic bulge and outer Galaxy fields are more similar, each having less than three magnitudes of extinction to their sources and averaging approximately 0.1 magnitudes difference between the source and lens. We use the tightness of this distribution in the outer Galaxy in our estimate of the source and lens stellar types in Section 2.5.

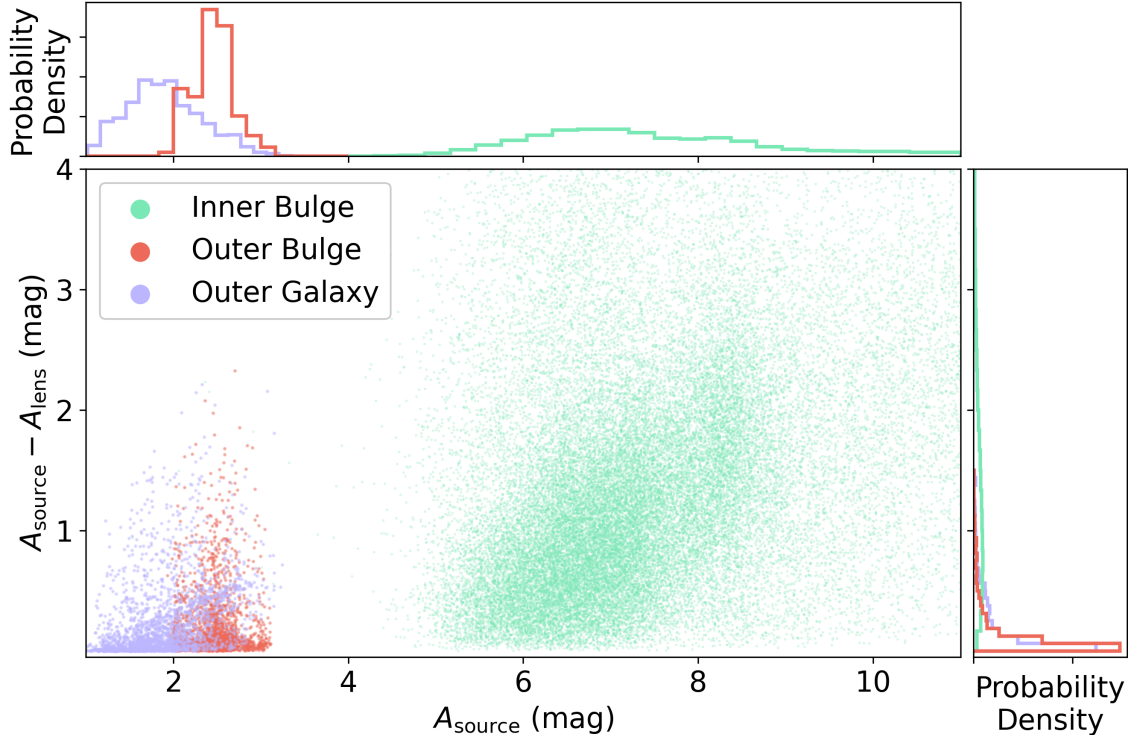


Figure 2.7: Extinction in the r-band to the sources and lenses of microlensing events in the inner bulge (green), outer bulge (red) and in the outer Galaxy (purple). Extinction towards the inner bulge is much larger than towards the outer Galaxy and even towards the outer bulge. This makes estimating the difference in extinction toward microlensing sources and lenses much more difficult in the inner bulge (up to 4 magnitudes) than in the outer bulge and outer Galaxy (less than 1 magnitude), despite the tighter constraints on both the distance to the sources and the distance ratio in the bulge (see Figure 2.4).

2.4.4 Contribution of Neighbors to Blended Light

The source flux fraction, b_{sff} , is

$$b_{\text{sff}} = \frac{f_{\text{S}}}{(f_{\text{S}} + f_{\text{L}} + f_{\text{N}})},$$

or the flux from the source f_{S} divided by the sum of the fluxes from the source, lens f_{L} and any neighbors that reside within the observational PSF f_{N} . The source flux fraction is often dominated by the presence of neighbors (stars that fall in the aperture but are neither the source nor the lens) in crowded Galactic fields. Figure 2.8 shows the contribution of flux from neighboring stars in an aperture of radius of $\theta_{\text{blend}} = 1.0''$ to simulate high quality seeing conditions on ZTF and an aperture of radius $\theta_{\text{blend}} = 2.25''$ to match the conservative estimate used throughout this analysis. Decreasing the observational aperture and the surrounding stellar density both reduce the fraction of flux originating from neighbors. Over half of the

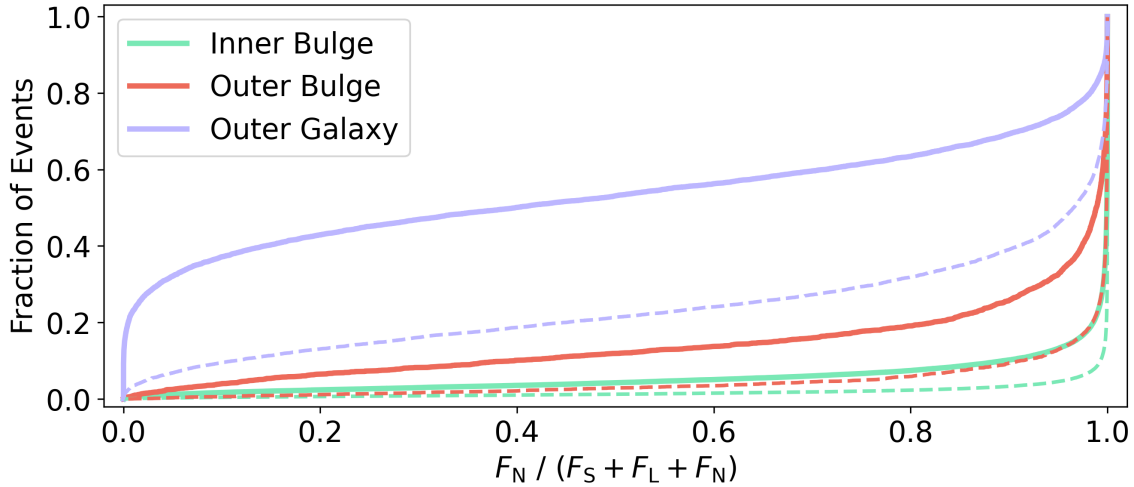


Figure 2.8: Fractional contribution of the flux from neighboring stars in a $\theta_{\text{blend}} = 1.0''$ observational aperture (solid) for the inner bulge (green), outer bulge (red) and outer Galaxy (purple), as well as a larger $\theta_{\text{blend}} = 2.25''$ aperture (dashed) for the outer Galaxy. Increasing the size of the observational aperture has a small effect on bulge fields where even the smaller aperture is dominated by the presence of neighbor flux. However improved seeing conditions in the outer Galaxy minimizes the contamination from neighbor flux to microlensing events, making these events easier to model.

events have in excess of 99% of their flux originating from neighbors in all fields observed with the larger aperture. This causes the source flux fraction to shift towards zero in these microlensing populations. However observing the outer Galaxy field with a smaller aperture results in half of the microlensing events having less than 40% neighbor flux. This makes modelling the source flux fraction of microlensing events along these lines of sight easier because one can reasonably use a strong prior that assumes a only a small amount of neighbor flux present, assuming that the event has been observed with high quality photometry. The population of events in the outer Galaxy is almost entirely devoid of neighbor flux due to lower stellar densities. This makes modelling the source flux fraction of microlensing events along these lines of sight easier because one can reasonably use a strong prior that assumes little to no neighbor flux present.

2.4.5 Implications for Outer Galaxy Microlensing

Future microlensing searches with ZTF must consider how the distribution of microlensing parameters across the outer Galaxy differs from those distributions in the Galactic bulge. While the shift in Einstein crossing times to larger values at these Galactic longitudes have been predicted, other microlensing parameters also change at these plane locations and must be considered to properly model events, measure properties of stellar populations and constrain galactic structure.

Modeling microlensing events with a Bayesian analysis requires selecting priors that are physically motivated by population statistics. The differences between the statistics of Galactic bulge and outer Galaxy populations should be noted as both an opportunity and a warning. Priors for microlensing populations that are appropriate for the Galactic bulge cannot be extended to analysis conducted in the outer Galaxy, and instead probabilistic priors should be derived from microlensing simulations performed at the location of microlensing events. We have made the catalogs of our fiducial microlensing populations available for public download at https://portal.nersc.gov/project/uLens/Galactic_Microlensing_Distributions/ following the data structure outlined in the PopSyCLE documentation. Future work will include releasing the full set of catalogs generated by our grid of PopSyCLE simulations.

2.5 Example Outer Galaxy Microlensing Event Analysis

The different microlensing population distributions in the outer Galaxy open the door to new opportunities for how to fit microlensing events. We here present an example ZTF microlensing event analysis to demonstrate how modelling outer Galaxy microlensing events can take advantage of these population statistics.

2.5.1 Event Selection

Price-Whelan et al. (2014) investigates statistical methods for detecting microlensing events in non-uniformly spaced time domain surveys that cover large areas of the sky. The heterogeneous time sampling and increased number of lightcurves in such a survey makes it challenging to adapt detection methods optimized from searches in the Galactic bulge to searches across the outer Galaxy. A method for finding microlensing events in surveys with a larger footprint must be extremely inexpensive to calculate for each lightcurve in order to scale efficiently. Price-Whelan et al. (2014) concludes that the von Neumann ratio (the mean square successive difference divided by the sample variance) works well as a statistic for filtering microlensing events that is inexpensive enough to be calculated for many lightcurves while discerning enough to avoid many of the false positives that other statistics routinely produce.

We calculated the von Neumann ratio on all lightcurves in the ZTF DR1 with $N_{\text{obs}} \geq 100$, totalling approximately 1.25×10^8 lightcurves. We removed all lightcurves with more than one cluster of consecutive observations more than 3σ above the median brightness of the source. This left 136,638 lightcurves in our sample. We selected the 2% of lightcurves with the largest von Neumann ratios and matched sources with both g-band and r-band lightcurves at the same sky location. 28 objects appeared to have amplification in the lightcurves of both filters which was achromatic to within approximately 0.5 magnitudes. However 25 of the objects had amplification that was quasi-periodic or slowly rising in what appeared by

eye unlikely to be microlensing. Those lightcurves with a characteristic microlensing shape were fit by microlensing models, resulting in one microlensing detection.

The purpose of this search and analysis was to verify that current ZTF cadence and filter coverage is capable of observing a measurable microlensing event. We emphasize that this process was meant to serve as neither a complete search nor a scalable model for microlensing discovery. Price-Whelan et al. (2014) outline a sophisticated statistical approach for determining cuts on statistical parameters, such as the von Neumann ratio, that are tailored to finding microlensing events. Our efforts were not to replicate this procedure but to instead scan the DR1 dataset using one of these statistics until a microlensing event was found. Our focus was on finding an example microlensing event to demonstrate how to use PopSycle to improve microlensing modeling, not to demonstrate a method for microlensing discovery. An improved search strategy could follow the detection algorithm of Price-Whelan et al. (2014) and include (1) removing lightcurves not simply by the number of observations but on the quality of those observations, (2) cutting lightcurves on a von Neumann ratio threshold determined from injecting artificial microlensing events into lightcurves to determine a false positive rate, (3) recalculating the von Neumann ratio after subtracting off a microlensing model, and more. Our search included none of these steps and we are therefore not surprised to find such a small completeness. Future work will include implementing a robust microlensing discovery algorithm resulting in measurements of the microlensing optical depth and event rate across the ZTF footprint.

2.5.2 Event Analysis

Figure 2.9 contains the lightcurves of our example microlensing event which was detected by the ZTF difference imaging alert stream and labelled ZTF18abhxjmj. Mróz et al. (2020b) includes this lightcurve in their list of microlensing events detected in the first year of ZTF’s Galactic Plane Survey; however we discovered this event independently by our event selection process. ZTF18abhxjmj is located at $(\alpha, \delta) = (284.02920^\circ, 13.15229^\circ)$ or $(\ell, b) = (45.19263^\circ, 4.93715^\circ)$ and began to rise at the start of the ZTF DR1 dataset in March 2018. Pan-STARSS1 (PS1) (Chambers et al. 2016) epochal data shows no previous variability in the years leading up to this event. Measurements in the months after ZTF18abhxjmj also show no variability, although more data at later times would help to better measure the baseline magnitude of the event.

We model ZTF18abhxjmj as a point-source, point-lens event allowing for blending and parallax effects. We transformed PS1 g-band and r-band data into the ZTF filter system to include the data in our fit (Medford et al. 2020c), helping to measure the long-duration baseline outside of the event. Bayesian fitting was performed with nested sampling (Skilling 2006) performed by PyMultiNest (Buchner et al. 2014), built on top of MultiNest (Feroz et al. 2009). Our fitter calculates magnifications in a heliocentric reference frame, avoiding the necessity to calculate a parameter reference time $t_{0,par}$. Priors for the Einstein crossing time and microlensing parallax components were taken from one dimensional marginalizations of the microlensing parameters extracted from PopSyCLE simulations pointed at the




Figure 2.9: Microlensing photometric lightcurve for ZTF18abhxjnj with ZTF (circles) and PS1 data (triangles), where the PS1 data has been transformed onto the ZTF filter system, in the g-band (top) and the r-band (bottom). 500 draws from the posterior distribution are in light gray for both filters. Note the break in the middle of the plot, as the PS1 data is from 2009 to 2012. The model captures the asymmetry in the rise and fall time due to parallax, but fails to appropriately match the baseline outside of the event with the PS1 r-band data.

location of the event, with observational cuts applied to the microlensing populations as described in Section 2.3. Following the example of previous work such as Batista et al. (2011), we apply not generic Galactic priors but priors specific to the mass density, galactic rotation, extinction and consequently the microlensing event rate towards this specific line of sight

in the Galaxy. Modelling microlensing events with Bayesian priors derived from PopSyCLE simulations allows for tighter constraints on posteriors than generic priors could otherwise produce.

Figure 2.9 shows 500 draws from our posterior distributions on top of our ZTF and transformed PS1 data. The model correctly captures the parallax effects near the peak of the event that appear as an asymmetry in the rise and fall time of the lightcurve. The model does not agree with the observed PS1 r-band flux, overestimating this contribution in order to fit the ZTF r-band baseline flux from after the event. The point source estimates from the two u_0 solutions in our fit posteriors of ZTF18abhxjmj can be found in Tables 2.3 and 2.4. The event’s Einstein crossing time (t_E) of 76 days is near the peak of the microlensing distribution for the outer Galaxy line of sight as seen in Figure 2.5. We note here that our Einstein crossing time (t_E), r-band baseline magnitude, r-band blend fraction and parallax components for ZTF18abhxjmj are all in agreement with the parameters found by Mróz et al. (2020b) for the same event in their parallax model. Our fit results in different values for t_0 and π_E which can occur due to the correlation between these variables in the heliocentric reference frame.

Transforming from aperture apparent magnitudes to source and lens apparent magnitudes requires using the source flux fraction, which can often be complicated by the presence of neighbor flux. As discussed in Section 2.4, very few microlensing events in the outer Galaxy have significant contributions to their flux from neighboring stars when observed with a relatively smaller photometric aperture of $\theta_{\text{blend}} = 1.0''$. We will assume these optimistic observing conditions because (1) this analysis takes place in the outer Galaxy where there is less confusion due to crowding and (2) the typical seeing on ZTF is around $1.5''$, and therefore an extraction method tuned to these conditions should be able to achieve such a blend radius. Assuming that the presence of neighbor flux is minimal has the convenient consequence of making the measurement of the source flux fraction approximately a measurement of the ratio of source flux to the sum of the flux from both the source and the lens. This approximation can be used to derive the the ratio of flux from the lens and source, or the lens-source-flux ratio, from the source flux fraction as follows:

$$\begin{aligned} b_{\text{sff}} &\approx \frac{f_S}{f_S + f_L + 0} \\ \frac{f_L}{f_S} &\approx \frac{1 - b_{\text{sff}}}{b_{\text{sff}}}. \end{aligned} \tag{2.5}$$

Figure 2.10 reveals that this approximation is valid in the outer Galaxy across 12 decades of $b_{\text{sff},r}$ values. It is in the Galactic bulge where the abundance of neighbors in the observable aperture makes the source flux fraction approximation an overestimation of the lens-source-flux ratio. Given that our fitter solves for the apparent magnitude of the source, we implement this approximation to calculate the apparent magnitude of the lens in each filter as:

$$m_{L,f} = m_{S,f} - 2.5 \log \left(\frac{1 - b_{\text{sff},f}}{b_{\text{sff},f}} \right), \quad f = \{g, r\}. \tag{2.6}$$




Figure 2.10: Comparison of the lens-source-flux ratio to an approximation derived from the source flux fraction (see Equation 2.5) in the ZTF r-band across 12 decades for both a larger photometric aperture ($\theta_{\text{blend}} = 2.25''$) and a smaller aperture ($\theta_{\text{blend}} = 1.00''$). Events in the outer Galaxy have relatively small contributions to their observable flux from neighboring stars when assuming a smaller aperture, making the source flux fraction approximation valid for almost all events. The presence of neighbor stars is the dominant cause of the spread and offset in the source flux fraction approximation in the remaining observations, making such an approximation invalid in the bulge fields and only partially correct in the larger aperture outer Galaxy field. This approximation enables the conversion from the apparent magnitude of the source to the apparent magnitude of the lens using the source flux fraction in Equation 2.6.

Figure 2.11 presents an apparent color-magnitude diagram of ZTF18abhxjmj (and surrounding stars) that results from folding this approximation into our fitting procedure. The ZTF and PS1 magnitudes and colors are derived from apparent aperture magnitudes taken outside of the microlensing event, while the model magnitudes are derived from the fit. The source and lens appear to have approximately the same apparent color due to their approximately equal source flux fractions ($b_{\text{sff},r} \approx b_{\text{sff},g} \approx 0.59$). The g-band source flux fraction ($b_{\text{sff},g}$) is approximately 0.59, meaning that the source and the lens contribute about equally to the apparent g-band brightness. The ZTF color is slightly redder than the PS1 color due

to the mismatch in the baseline magnitude in the lightcurve. The model attributes a color to the source and lens between these two values with appropriately larger error bars, reflecting this discrepancy.

Calculating the absolute magnitudes and stellar types of the source and lens from their apparent magnitudes requires knowing their distances and extinctions. As discussed in Section 2.4, PopSyCLE produces distributions of distances and extinctions for microlensing events along a specific line of sight. We generated PopSyCLE simulations at the location of ZTF18abhxjmj and applied observational cuts to the event catalogs that simulated ZTF observing conditions. Samples were drawn from these trimmed catalogs, weighted by the event model’s Bayesian posteriors for baseline magnitude and source flux fraction in g-band and r-band. Figure 2.12 shows the absolute color-magnitude diagram of the samples that resulted from this procedure.

The source of ZTF18abhxjmj has an absolute magnitude in the g-band of $M_{S,g} = 4.6 \pm 0.6$ and an absolute color of $M_{S,g-r} = 0.49 \pm 0.07$, while the lens has an absolute g-band magnitude of $M_{L,g} = 11.1 \pm 2.6$ and an absolute color of $M_{L,g-r} = 1.3 \pm 0.36$ magnitudes. We matched these source and lens absolute magnitudes to absolute magnitudes of stars generated in synthetic clusters with PyPopStar (Jr. et al. 2020), a python package that generates single-age, single-metallicity populations from user specified initial mass functions, stellar evolution models, and stellar atmospheres. The source approximately resembles a 1.04 solar mass G-star in a $10^{9.82}$ year old cluster, and the lens approximately resembles a 0.39 solar mass M-dwarf in a $10^{7.8}$ year old cluster. Systematic errors in the Galactic model implemented in PopSyCLE significantly contribute to the uncertainty in these conclusions but are not captured by our stated errors.

We have included in Figure 2.12 the source and lens absolute magnitudes that would have been calculated if a simulated catalog from the outer bulge was used instead of one produced along the target’s line of sight. Microlensing source and lenses towards the bulge are, on average, at closer distances and are behind more magnitudes of extinction. These two facts have opposite effects on the estimate of the source’s absolute magnitude. The additional extinction pushes the source star’s probability to a smaller absolute magnitude in order for the source or lens to appear at the apparent magnitude determined by the Bayesian fit, with the closer distance having the opposite effect. The results of these two competing effects can be resolved with PopSyCLE simulations at the location of each microlensing event that ZTF discovers modelled after this fitting procedure.

This example analysis demonstrates how data from ZTF and simulations from PopSyCLE can be combined to fit microlensing models and estimate stellar types of microlensing sources and lenses. The results of this particular analysis are not exceptional as M-dwarfs are extremely common throughout the Galaxy and are often found to be lenses of microlensing events, although this method could be used to find more exotic lenses such as free-floating planets and black holes. We have outlined the steps of this analysis to illustrate how probabilistic priors for a specific event can be quickly generated through modelling microlensing populations toward a particular line of sight.

2.6 Discussion

The Zwicky Transient Facility and its surveys are an excellent opportunity to discover microlensing events. We find that ZTF will observe ~ 1100 events in three years of observing, with ~ 500 events occurring outside of the Galactic bulge in the outer Galaxy ($\ell \geq 10^\circ$). This total can be increased to ~ 2400 events (~ 1300 events in the outer Galaxy) by extending ZTF operations to five years and executing a post-processing image co-addition pipeline. The event rate of microlensing is proportional to the number of observed luminous sources. While ZTF's single image limiting magnitude is not as deep as other optical surveys, its massive 49 deg^2 camera is able to cover the entire northern sky every three nights in multiple filters. The decrease in microlensing event rate outside of the Galactic bulge that discourages other microlensing surveys is compensated for by the billions of stars observed within this large footprint. Observing in the outer Galaxy almost doubles the total number of microlensing events that ZTF will observe.

Microlensing events can be discovered in ZTF by searching through the epochal photometric catalogs present in the public data releases described in Section 2.2. These catalogs contain observations in multiple filters that allow for confirming a potential microlensing event through its achromaticity. ZTF also generates subtraction images for all of its exposures and serves a real-time alert stream of transient detections found on these difference images. Filters could be developed that search for microlensing events on a nightly basis (Price-Whelan et al. 2014; Godines et al. 2019), generating a list of candidates that could trigger photometric or astrometric followup. This would be particularly helpful in attempting to detect exoplanets through microlensing, which requires triggering higher cadence followup near the photometric peak of the event, as well as discovering black holes lenses which requires astrometric follow-up.

Microlensing detections made outside of the Galactic plane will be extremely rare due to the decrease in luminous stellar sources. Galaxies begin to be the dominant luminous sources in these fields and the distance ratio of luminous sources and massive lenses does not result in observable microlensing events. Galaxies are far away and microlensing is maximized when the lens is halfway between the source and the observer, so we therefore cannot hope to observe any microlensing events where galaxies are the luminous source. However this challenge can be inverted to provide an interesting opportunity. There is a possibility that primordial black holes (PBHs) significantly contribute to dark matter and could be observed through microlensing. Previous work suggest that the dark matter mass fraction contributed to by PBHs could be constrained through an effect on the shape of the Einstein crossing time distribution (Green 2016, 2017; Niikura et al. 2019; Lu et al. 2019). Given the lack of observable microlensing events outside of the Galactic plane, and the isotropic distribution of dark matter, any microlensing detections made outside the plane could place constraints on the PBH dark matter fraction. The likelihood that a microlensing event is caused by a PBH lens relative to a stellar lens increases when observing outside the Galactic plane. A ZTF microlensing survey would be one of the only microlensing surveys conducted that includes observing in these fields, making it one of the few surveys that

could make this measurement. There may also be advantages in looking for black holes as microlensing lenses in the outer Galaxy as compared to the Galactic bulge. Detecting a black hole through microlensing requires weighing the mass of the lens despite the lens mass' degeneracy with microlensing parallax when using photometric data. This degeneracy can be avoided by astrometric measurement which can determine the mass of the lens directly. As discussed in Section 2.4.2, outer Galaxy microlensing events have larger Einstein radii and therefore have an astrometric signature that is easier to detect. (Lam et al. 2020) outline how black hole lenses have significantly larger maximum astrometric shifts, longer Einstein crossing times and less microlensing parallax than star, white dwarfs or neutron stars lenses with PopSyCLE simulations. Figure 2.13 replicates Figures 12 and 13 from Lam et al. (2020) in our outer bulge and outer Galaxy fields. All events in the outer Galaxy sample occur at longer Einstein crossing times and with larger microlensing parallaxes, making them easier to measure and therefore distinguish black hole lenses. The maximum astrometric shift is significantly larger, averaging almost an order of magnitude above the 0.2 milliarcseconds that the Keck laser guide star adaptive optics system is capable of measuring (Lu et al. 2016) and maxing out at over 5 milliarcseconds. Decreased stellar densities in the outer Galaxy will present a challenge to making this measurement and requires the Hubble Space Telescope or wide field adaptive optics such as an upgraded Gemini North adaptive optics system if measured from the ground. Future space instruments such as the James Webb Space Telescope (Gardner et al. 2006) or the Nancy Grace Roman Space Telescope (Spergel et al. 2013) will be more than capable of detecting black holes using this technique.

There are also challenges that arise when attempting to use ZTF to make microlensing measurements. ZTF's photometric precision of ~ 0.1 magnitudes at a limiting magnitude of $m_{\text{lim}} < 21$ (Masci et al. 2018) can make it difficult to detect events with a large impact parameter or small maximum amplification. These events will be difficult to distinguish from background noise or variability of faint stars. ZTF is located in the Northern hemisphere, limiting exposure to the Galactic plane to select summer months of the year, reducing the total number of observable short duration events. ZTF is also a collaboration with many priorities both Galactic and extra-Galactic resulting in decisions on survey design, cadence and scientific goals that are not necessarily optimized for microlensing.

Future synoptic surveys such as the Rubin Observatory Legacy Survey of Space and Time (LSST) could continually monitor billions of stars across the Milky Way for many years, providing opportunities to learn about galactic structure, stellar populations and possibly even dark matter through photometric microlensing. The massive footprints of surveys such as ZTF and LSST unlock the potential to observe thousands of microlensing events across the entire Galactic plane and possibly even off the plane, expanding beyond the scope of microlensing surveys to date that have been pointed at the Galactic bulge and other nearby galaxies. Combining these datasets with sophisticated microlensing modelling software can result in improvements to stellar categorization and population statistics that would otherwise be out of reach for these photometric surveys.

Table 2.3: Microlensing Parameters of ZTF18abhxjmj

ℓ deg.	b deg.	t_0 MJD	t_E days	u_0 -	$m_{S,g}$ mag	$b_{\text{sf},g}$ -	$m_{S,r}$ mag	$b_{\text{sf},r}$ -	$\pi_{E,N}$ -	$\pi_{E,E}$ -
284.02916	13.15228	58229.9	76.7	0.14	21.84	0.592	20.33	0.593	0.187	0.257
		+4.0	+8.7	+0.04	+0.17	+0.077	+0.17	+0.078	+0.054	+0.050
		-4.1	-8.8	-0.03	-0.15	-0.102	-0.15	-0.097	-0.040	-0.036
		58227.1	75.8	-0.05	21.85	0.591	20.33	0.589	0.198	0.241
		+3.7	+8.0	+0.05	+0.16	+0.065	+0.16	+0.056	+0.047	+0.051
		-4.0	-6.6	-0.05	-0.12	-0.096	-0.11	-0.094	-0.036	-0.036

The microlensing parameters of the median best-fit point-source point-lens microlensing model of ZTF18abhxjmj, including the time of maximum heliocentric amplification (t_0), Einstein crossing time (t_E), minimum source-lens separation in units of the Einstein radius (u_0), baseline magnitudes for the source in g-band and r-band ($m_{S,g}, m_{S,r}$), source-flux-fractions in g-band and r-band ($b_{\text{sf},g}, b_{\text{sf},r}$) and the two components of the microlensing parallax ($\pi_{E,N}, \pi_{E,E}$). We find an Einstein crossing time of 76 days in our two u_0 solutions and a blend fraction in both g-band and r-band around 0.59. These values indicate that the flux in the aperture is about equally split between the source and the lens in both filters. The visible parallax in the lightcurve appears in the fit, confirmed by significant components of π_E .

Table 2.4: Model Magnitudes of ZTF18abhxjmj

	M_g	M_r	M_{g-r}
Lens	11.12 ± 2.64	9.84 ± 2.28	1.28 ± 0.36
Source	4.58 ± 0.59	4.08 ± 0.57	0.49 ± 0.07

The absolute magnitudes (M_g, M_r), and absolute color (M_{g-r}) of the point-source point-lens microlensing model of ZTF18abhxjmj. The absolute magnitudes are calculated by drawing samples from the PopSyCLE simulations generated at the location of the event weighted by the posteriors of our Bayesian fit. The errors on these measurements do not include systematics from PopSyCLE’s Galactic model.

Acknowledgements

This work was done with support from the University of California Office of the President for the UC Laboratory Fees Research Program In-Residence Graduate Fellowship (Grant ID: LGF-19-600357). Part of this work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and was supported by the LLNL-LDRD Program under Project Number 17-ERD-120. J.R.L. and C.L. acknowledge funding from the National Science Foundation (AST-1909641). Authors would also like to thank Danny Goldstein and Matt Hosek for their support both technical and otherwise.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes. Lastly we thank the lawyers of Lawrence Livermore National Laboratory for their contributions to this acknowledgement section.




Figure 2.11: Apparent color-magnitude diagram for 88,914 stars cross-matched in the ZTF g-band and r-band in the 0.77 square degrees surrounding ZTF18abhxjmj. Highlighted are the apparent magnitudes of the event as calculated by the ZTF observations outside of the event (gray), the PS1 observations placed onto the ZTF filter system (blue) and the baseline apparent magnitudes as calculated by the point source point lens model for the source (yellow) and the lens (green). The ZTF measurement is slightly redder than the PS1 measurement, consistent with the mismatched out-of-event flux shown in Figure 2.9, but still within the error of the measurement. The g-band source flux fraction of $b_{\text{sf},g} = 0.59$ places the source and the lens at nearly the same observable g-band magnitude, while the similar source flux fractions in both filters ($b_{\text{sf},g} \approx b_{\text{sf},r}$) places the source and the lens at nearly the same observable color.

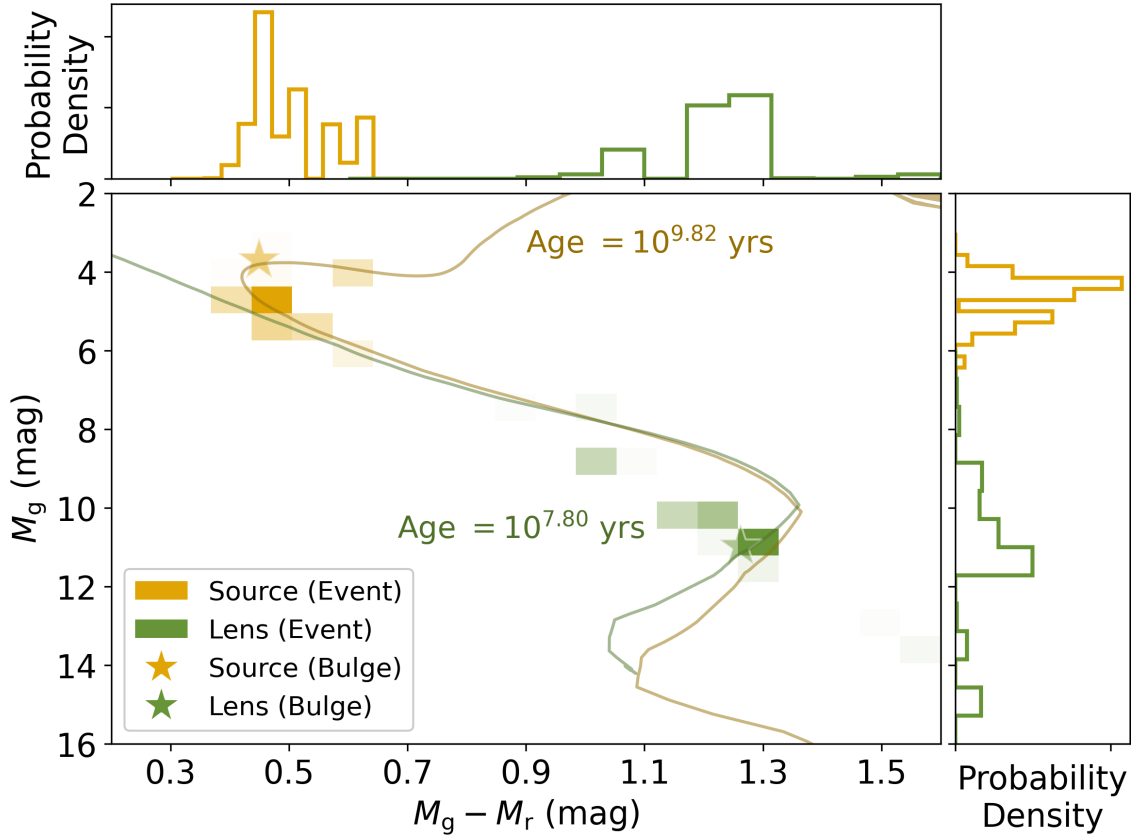


Figure 2.12: Absolute color-magnitude diagram for the source (yellow) and lens (green) derived from combining Bayesian modeling and PopSyCLE simulations, with histograms on both axes showing the marginalized distributions of the parameters. Isochrones generated by PyPopStar have been drawn to approximate the source and lens ages to be $10^{9.82}$ years and $10^{7.8}$ years respectively. Point estimates for the source and lens calculated using PopSyCLE catalogs generated in the outer bulge (stars) find a slightly brighter source due to the additional extinction in the Galactic bulge. These estimations are highly sensitive to the systematic errors discussed throughout Section 2.5.

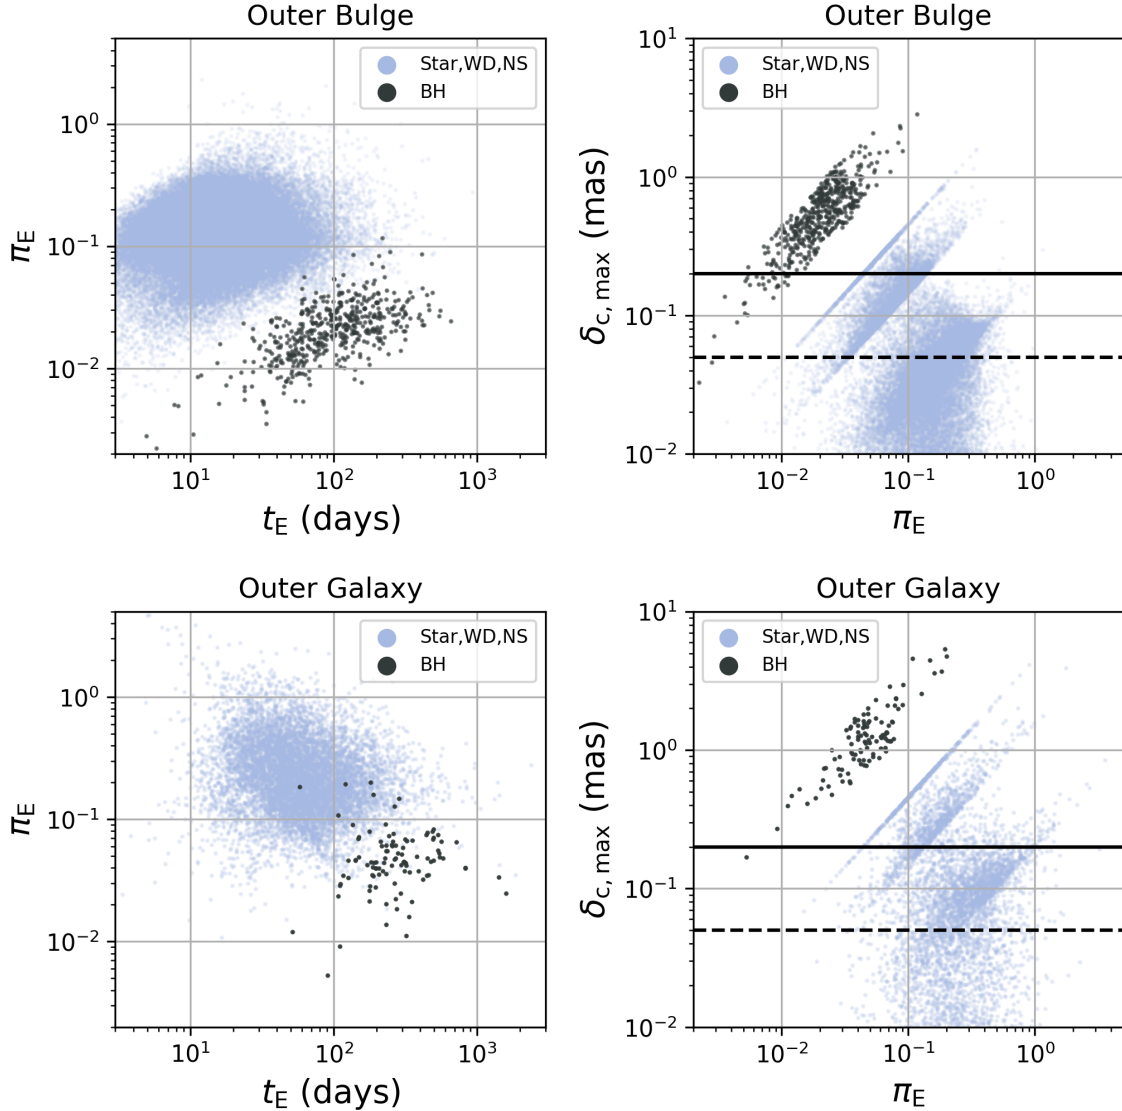


Figure 2.13: Microlensing parallax π_E vs. Einstein crossing time (top) and maximum astrometric shift $\delta_{c,max}$ vs. microlensing parallax π_E broken out by astrophysical type of lens for the outer bulge (top) and outer Galaxy (bottom) fields. PopSyCLE simulations reveal that black hole microlensing lenses are distinct from stars, white dwarfs and neutron stars in these two spaces (Lam et al. 2020). Both the Einstein crossing times and microlensing parallaxes increase when measured in the outer Galaxy (bottom) as compared to the outer bulge (top), making it easier to constrain black holes in this plane. The maximum astrometric shift for black holes increases to a decade above the detection limit of the Keck laser guide star adaptive optics system (solid) and almost two decades above the anticipated limits of the Nancy Grace Roman Space Telescope or the Thirty Meter Telescope (dashed). Introducing observational cuts not present in these figures reduces the total number of events but maintains the same trends.

Chapter 3

Removing Atmospheric Fringes from Zwicky Transient Facility i-Band Images using Principal Component Analysis

This is a replication of the article published as [Medford, Nugent, Goldstein, Masci, Andreoni, Beck, Coughlin, Duev, Mahabal, & Riddle 2021](#), PASP, 133, 1024.

Abstract

The Zwicky Transient Facility is a time-domain optical survey that has substantially increased our ability to observe and construct massive catalogs of astronomical objects by use of its 47 square degree camera that can observe in multiple filters. However the telescope's i-band filter suffers from significant atmospheric fringes that reduce photometric precision, especially for faint sources and in multi-epoch co-additions. Here we present a method for constructing models of these atmospheric fringes using Principal Component Analysis that can be used to identify and remove these artifacts from contaminated images. In addition, we present the Uniform Background Indicator as a quantitative measurement of the reduced correlated background noise and photometric error present after removing fringes. We conclude by evaluating the effect of our method on measuring faint sources through the injection and recovery of artificial stars in both single-image epochs and co-additions. Our method for constructing atmospheric fringe models and applying those models to produce cleaned images is available for public download in the open source python package [fringeZ](#).

3.1 Introduction

Large scale synoptic surveys have produced an abundance of optical images at scales previously unseen. Surveys such as the Palomar Transient Factory (Law et al. 2009), Catalina Real-Time Transient Survey (Drake et al. 2009), Zwicky Transient Facility (Bellm et al. 2019c; Graham et al. 2019; Masci et al. 2018) and others have revolutionized our understanding of the universe by generating massive transient datasets. These datasets require advances in computational processing techniques. The Vera C. Rubin Observatory will generate 20 terabytes of data per night, totalling over 500 petabytes of imaging data over the 10 years of the survey (Ivezić et al. 2019). Such data flows require that the reduction of raw images into calibrated science images must be systematic and require minimal human intervention.

One significant source of noise in long wavelength optical imaging is atmospheric emission lines. These emission lines are produced by highly non-thermal atomic and molecular transitions (primarily O II and OH) and are influenced by the temperature and density of the upper and lower atmosphere as well as the current solar activity. Thus the strength of these lines can vary throughout the night and is proportional to the airmass through which the telescope is pointed. Fringe patterns appear when photons at these wavelengths fall onto thin charge-coupled devices (CCDs) due to the self-interference caused from light reflecting off of the back of the imaging instrument before it is absorbed by the CCD itself (Bernstein et al. 2017). Thick CCDs rarely see this effect save at the longest wavelengths. For thin CCDs, these fringe patterns can introduce significant noise into i-band and z-band images, rendering photometry and image subtraction ineffective at faint magnitudes without a calibration correction that successfully removes them. These fringes also appear in interference image spectrometers where attempts have been made to remove them using wavelet transformations (Ren et al. 2017).

Principal component analysis (PCA) is a statistical method for reducing data to a set of orthogonal components by finding vectors of minimal variance within the data (Jolliffe & Cadima 2016). PCA has been shown to be an effective method for modeling and removing atmospheric fringes from red optical and near-IR images. These methods work by building a set of orthogonal component images, constructed from a large representative sample, which can approximate the fringes of a single image through linear combination. The dot product of individual images against these orthogonal components results in eigenvalues that, when used as weights to these orthogonal images, generate bias images of the atmospheric fringes. Subtracting these bias images removes the fringe noise and improves the photometric precision of astrophysical source measurements.

Previous PCA methods for atmospheric fringe subtraction constructed orthogonal images by down-sampling fringed images into a lower resolution before re-parameterization of the data. Bernstein et al. (2017) compress an image into a sparse set of features before attempting reduction into orthogonal components. This method has strong results but lacks the ability to resolve contributions from individual atmospheric lines due to its compression. PCA performed on a per-pixel basis is more computationally expensive but has the power to

capture these individual fringe effects, resulting in more accurate photometry.

The computational resources dedicated to the reduction and calibration of astronomical images has grown over the decades to keep up with increasing data flows. The development of highly optimized PCA algorithms, along with additional computational resources, have now made it possible to model fringes via per-pixel PCA analysis. Here we present the implementation of such a method on full resolution Zwicky Transient Facility i-band data. In section 5.2, we outline the Zwicky Transient Facility instrument and dataset. In Section 5.4, we present our method for implementing per-pixel PCA atmospheric fringe modeling and removal, as well as the Uniform Background Indicator as a quantitative measurement of correlated background noise. In Section 5.5, we analyze the results of applying our method, including increased photometric precision on faint sources and the ability to detect otherwise undetectable sources in multi-epoch co-additions. We discuss and conclude in Section 3.5.

3.2 The Zwicky Transient Facility Instrument

The Zwicky Transient Facility (ZTF) is an optical time-domain survey that has been operating on the 48-inch Samuel Oschin Telescope at Palomar Observatory since March 2018 (Bellm et al. 2019c; Graham et al. 2019). ZTF’s camera covers 47 square degrees in a single exposure, enabling coverage of the entire visible Northern sky every few nights in ZTF g-band, r-band and i-band filters with an average point spread function (PSF) full width at half maximum of $2.0''$ on a plate scale of $1.01'' \text{ pixel}^{-1}$. The ZTF camera is divided into 16 CCDs, each covered with an anti-reflective (AR) coating, with each CCD split into four separate readout channels for a total of 64 readout channels (Dekany et al. 2020). Surveys with the telescope over its first several years of operations (Bellm et al. 2019b) take standard 30 second exposures that achieve median five-sigma limiting magnitudes of $r \approx 21.0$, $g \approx 21.25$, and $i \approx 20.0$ (Bellm et al. 2019c). The Infrared Processing and Analysis Center (IPAC) reduces the data for the ZTF survey, including producing a real-time alert stream triggered by transient event detections on difference images (Patterson et al. 2018). In addition to these alerts, the ZTF collaboration routinely produces public data releases that contain, among other data products, lightcurves assembled from single image PSF photometry for every star in the northern sky that appears in a deep co-added reference image (Masci et al. 2018). Reference images are ideally constructed from 40 individual exposures, although weather and visibility produces variable results for different areas of the sky. ZTF’s observing time is split between public observations (40%) (funded by the National Science Foundation’s Mid-Scale Innovations Program or MSIP), partnership observations (40%) (which are held in a proprietary period for collaboration members of the survey), and Caltech observations (20%). Starting in 2020, ZTF began operating its second phase as ZTF-II with 50% of observation time funded by a MSIP survey that images the entire visible sky every other night. The ZTF i-band filter is used exclusively for partnership observations, such as the ZTF Ultra Deep Survey (ZUDS), before being released to the public after the partnership’s proprietary period. Due to the decreased limiting magnitude of the i-band relative to the g-

and r-band, standard exposures in the i-band are set to 90 seconds and increase the filter’s median limiting magnitude to a comparable value. Atmospheric fringes appear exclusively in the i-band and therefore this work uses partnership data not yet released to the public.

3.3 Method for Removing Atmospheric Fringes

Our method for removing atmospheric fringes was a two step process. First the PCA eigen-vectors for a single readout-channel were extracted from a large set of images with significant fringing. Second each image was processed through this model to generate a unique bias image that, when subtracted from the original image, removed the atmospheric fringes. Figure 3.1 outlines a visual representation of the steps in this method. An explanation of the data products and corresponding symbols used throughout this section can be found in Table 3.1.

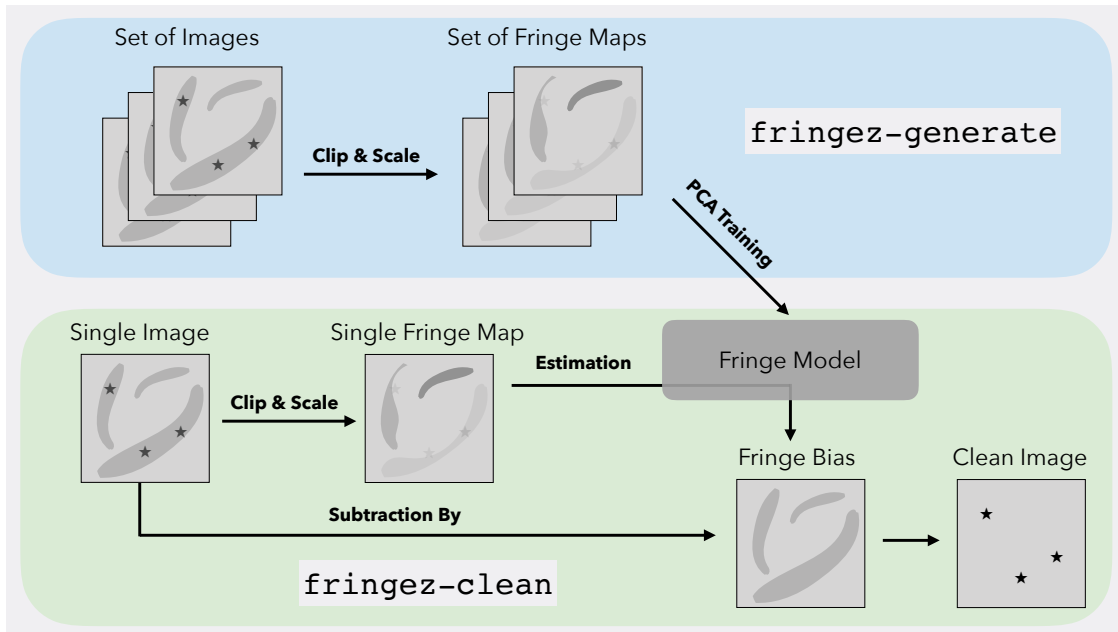


Figure 3.1: The process for removing atmospheric fringes was done in two steps. A fringe model for each readout channel was constructed by running fringe maps through PCA training, where fringe maps are a clipped and scaled version of single-epoch images. This step was only performed once per readout channel. Cleaning every single-epoch image thereafter was performed by using the eigen-vectors in the fringe model to generate a fringe bias image for each fringed image. This fringe bias was subtracted from the fringed image to create a clean image. Functions for generating fringe models and cleaning fringed images can be found in the open source [fringe](#) package under the executables `fringe-generate` and `fringe-clean` respectively.

Table 3.1: `fringe` Math Symbols

Description	Symbol	Dimensions
Fringed Image	$\mathbf{I}_{\text{fringe}}$	$[N_{\text{rows}} \times N_{\text{cols}}]$
Flattened Fringed Image	\vec{I}_{fringe}	$[1 \times N_{\text{pixels}}]$
Fringe Map	\vec{F}_{map}	$[1 \times N_{\text{pixels}}]$
Fringe Map Eigen-values	$\vec{\lambda}_{\text{map}}$	$[1 \times N_{\text{comp}}]$
Fringe Bias	\vec{F}_{bias}	$[1 \times N_{\text{pixels}}]$
Average of Training Set	$\vec{\mu}_{\text{model}}$	$[1 \times N_{\text{pixels}}]$
Model Eigen-images	$\mathbf{v}_{\text{model}}$	$[N_{\text{comp}} \times N_{\text{pixels}}]$
Model Explained Variance	\vec{R}_{model}	$[1 \times N_{\text{pixels}}]$

3.3.1 Generating PCA Fringe Models

ZTF’s reduction pipeline was designed to process each of the camera’s 64 readout-channels separately. We therefore began by gathering images of a common readout-channel together for model generation, with the goal of generating 64 fringe models. Training images were selected by gathering all i-band images between 2019-04-01 and 2020-04-01 and removing images with a limiting magnitude less than 19. This cut on limiting magnitude removed cloudy images from our sample that would fail to have a photometrically accurate measurement of the night sky and thus its atmospheric lines. By including all images within these dates we ensured a representative sample of atmospheric conditions and airmasses which are correlated with the strength of atmospheric fringes. There were 550,365 i-band images included in total, with each fringe model trained on between 8,062 and 8,898 images.

We began by selecting all fringed images ($\mathbf{I}_{\text{fringe}}$) for a single readout channel. The location and strength of atmospheric fringes are independent of the astrophysical sources within an image. Therefore we needed to remove sources from each image prior to training in order to reduce the confusion in the PCA variance reduction process. Our preferred method for identifying pixels containing astrophysical sources was to use the quality masks produced by IPAC that labels pixels containing a source in each image’s source catalog. In cases where this pixel mask could not be obtained, pixels containing sources were identified by calculating the median absolute deviation for the image and flagging pixels that were 5 standard deviations above or below the image’s median value. This process did not clip the atmospheric fringes themselves, which were at the level of the background noise and therefore significantly below the level of these thresholds. All pixels flagged as containing sources were replaced with the value of the image’s global median, removing stars from our images without changing the value of the vast majority of pixels. A possible improvement to our method could be to replace the pixel instead with an estimate of the local background. We choose the global median because the local background can be difficult to estimate due to the fringes themselves and in particular in regions where the transverse length scale of the fringes approaches the pixel scale of the image. Each image was then scaled so that

they could be used for training together with other images taken at different airmasses and limiting magnitudes. We named each training image a fringe map (\vec{F}_{map}), as it traces the relative location and strength of only the atmospheric fringes. Fringe maps were created by subtracting the median from each image, dividing the result by the image’s median absolute deviation, and flattening the image into a $[1 \times N_{\text{pixels}}]$ array:

$$\vec{F}_{\text{map}} = \frac{\mathbf{I}_{\text{fringe}} - \tilde{\mathbf{I}}_{\text{fringe}}}{\text{MAD}(\mathbf{I}_{\text{fringe}})} \Big|_{\text{flattened}} . \quad (3.1)$$

(Because PCA is a linear method that calculates coefficients irrespective of the basis in which the data is represented, performing our analysis in one dimension has no effect on the results.) Each pixel in the fringe map was treated as a separate variable in our PCA analysis, resulting in 9,461,760 variables (3080 rows by 3072 columns) for each model. The eigenvalues of the feature vector was determined using the randomized Singular Value Decomposition method (Martinson et al. 2011; Halko et al. 2011) in the `scikit-learn` python package (Pedregosa et al. 2011). Each 1-dimensional eigen-vector was reconstructed into the original image shape to create an eigen-image. The final fringe model was a set of eigen-images that captured orthogonal contributions to the training set variance in each pixel. We also recorded each component’s explained variance (\vec{R}_{model}), or the amount of training set variance pointed in each eigen-vector’s direction, for later use in the processing pipeline. We then repeated this process to construct a unique fringe model for each of the camera’s 64 readout channels.

3.3.2 Removing Fringes from Single-Epoch Images

Fringes were removed from images by applying the fringe model in these steps:

$$\vec{\lambda}_{\text{map}} = \frac{(\vec{F}_{\text{map}} - \vec{\mu}_{\text{model}}) \cdot \mathbf{v}_{\text{model}}^T}{\sqrt{\vec{R}_{\text{model}}}} , \quad (3.2)$$

$$\vec{F}_{\text{bias}} = \left[\left[\vec{\lambda}_{\text{map}} \cdot \left(\sqrt{\vec{R}_{\text{model}}} \mathbf{v}_{\text{model}} \right) \right] + \vec{\mu}_{\text{model}} \right] \cdot \text{MAD}(\vec{I}_{\text{fringe}}) , \quad (3.3)$$

$$\vec{I}_{\text{clean}} = \vec{I}_{\text{fringe}} - \vec{F}_{\text{bias}} . \quad (3.4)$$

First, a fringe map was regenerated for each fringed image in our sample. The dot product of this fringe map (less the average value of each pixel in the training set $\vec{\mu}_{\text{model}}$) and the model’s eigen-images was calculated, and then divided by the square root of that component’s explained variance (Equation 3.2). This produced the eigen-values ($\vec{\lambda}_{\text{map}}$) for each fringe map. These eigen-values were used as coefficients to linearly combine re-scaled eigen-images and the average values of the training sets were added back in (Equation 3.3). This image, after being re-scaled by the median absolute deviation of the original fringed

image, was named a fringe bias (\vec{F}_{bias}). When this fringe bias was subtracted from the original fringed image it produced a clean image (\vec{I}_{clean}) with significantly less fringing (Equation 3.4).

An example of each of the images in this process are shown in Figure 3.2. This figure visually demonstrates the results of our method. The original image is a 90 second i-band exposure that contains a representative amount of atmospheric fringes. The fringe map is almost entirely devoid of individual sources, although the source pixel identification method does struggle around particularly bright stars. The fringe bias generated from processing the fringe map through the fringe model successfully identifies the location and strength of each atmospheric fringe. The clean image has nearly all of the atmospheric fringing pattern removed while retaining nearly all of the astrophysical sources.

The arithmetic average of the fringe maps used to train the fringe model for each readout-channel is shown in Figure 3.3, with each readout-channel placed on a common gray-scale. The pattern that clearly emerges is coherent on the level not just of the readout channel, but also over the full CCD. The etching process for creating a thinned CCD uses a circular buffer that removes layers of the chip to create a uniformly thick device. The average of the training fringe maps identifies the residual thickness variations for each CCD. Work is ongoing to use these psuedo-measurements of the thickness to improve the ZTF data quality pipeline (Richard Dekany, private communication). The inner 32 readout channels ($16 \leq \text{rcid} < 48$) and the outer 32 readout channels ($0 \leq \text{rcid} < 16$, $48 \leq \text{rcid} < 64$) have distinctly different amounts of atmospheric fringing present in their images. The inner readout channels have two layers of AR coating while the top and bottom rows of CCDs have a single layer coating. This causes the outer readout channels to have a higher reflectivity at the longer wavelengths where fringing occurs and thus a larger contrast of the fringing pattern.

The eigen-images for readout channel 13 are shown in Figure 3.4 as a representative example of a fringe model. Each pixel in the image was trained as a separate variable and yet the fringe patterns across pixels remains coherent after reconstruction into the original image dimensions. This confirms that the components contain correlated eigen-vectors for the different features. In addition to the atmospheric fringes, the PCA training captured large scale variations in the background that remain after flat fielding. These backgrounds dominate the higher order components. This indicates that some global scale variations in flux remains after the execution of the ZTF flat-fielding pipeline that could possibly be improved by implementing a PCA method (Bernstein et al. 2017). We note that PCA requires setting the number of components, or number of eigen-vectors, for the reduction algorithm. The fractional explained variance for all of the components are also shown in Figure 3.4. On average across the 64 readout channels, the first component captured 64.6% of the pixel variance while the sixth component captured only 5.0% of the pixel variance. In total six PCA components reconstructed 95.0% of the variance seen in our training sample. Adding additional components failed to significantly increase the amount of fractional explained variance as the previous component. We therefore determined that six components was sufficient to




Figure 3.2: An example of an i-band image in the various steps of our fringe removal method. Fringed images were clipped and scaled to produce a fringe map. The fringe map was processed through a fringe model to generate a fringe bias. This bias image was subtracted from the original fringed image to create a cleaned image. While a fringe model only needs to be constructed once, each single-epoch image has a unique fringe bias determined by the linear combination of eigen-images in the fringe model that best reconstructs the fringe map. This ensures that the fringe model can successfully remove fringes arising from a variety of airmasses and seeing conditions.




Figure 3.3: The average of the training fringe maps for each readout channel’s fringe model on the ZTF camera, all placed onto a common gray-scale. These circular patterns roughly trace the thickness variations in the CCD and show the circular pattern resulting from flattening the device. The inner 32 readout channels show significantly less atmospheric fringing due to an additional layer of anti-reflective coating, as compared to the outer 32 readout channels.

capture the variation due to atmospheric fringes.

Successfully training the fringe models required simultaneously holding large numbers of training images in computer memory in order to execute singular value decomposition. The large number of images that trained each of our 64 fringe models ranged in memory from 285 gigabytes to 314 gigabytes. We trained our models on the Cori Haswell nodes at the NERSC

Supercomputer located at the Lawrence Berkeley National Laboratory. The Haswell nodes each have 128 gigabytes of RAM, making it impossible to include such a large number of images in our training sets without modifying our method. To adapt to cases where access to computational resources are limited such as this one, our training method was modified by down-sampling the training set. Training images were sorted by limiting magnitude and split into 300 sub-stacks of approximately equal depth. In order to combine each sub-stack into a single training image we calculated the sub-stack’s median. While taking the median can distort the correlations between pixels that PCA is attempting to solve for, the median is also more robust than the mean to removing outliers that have contaminated the signal from a different distribution than the one attempting to be measured. Smooth background gradients remain on some ZTF images that can confuse the PCA reduction algorithm if not removed as outliers. We therefore settled on taking the median instead of the mean to combine the sub-stack of fringe maps into a single training image. An alternative approach could be to entirely remove the images with smooth gradients through a different method and then combine the remaining images in the sub-stacks using the mean. These 300 training images totalled approximately 11 GB and were therefore able to be fed simultaneously into our training process. Comparing the fringe models that result from training on down-sampled fringe maps produced comparable results to training on individual fringe maps for smaller sample sizes.

The code used for generating fringe models from ZTF i-band images and cleaning fringed images has been released in the open-source package `fringe` (Medford & Nugent 2019). The `fringe` package includes command line executables and python functions for downloading the fringe models computed as a result of this work, producing fringe bias images, and cleaning fringed images. This package is available for download and installation at <https://github.com/MichaelMedford/fringe>, and all fringe models are available for download at <https://portal.nersc.gov/project/ptf/iband>. In October 2019, `fringe` was implemented into the ZTF IPAC data reduction pipeline with fringe models built on 500 training images per readout channel. All i-band images taken previous to this date were also re-processed by IPAC to remove atmospheric fringes. In November 2020, fringe models were updated to include the 550,365 i-band images described and investigated in this paper. This implementation of `fringe` and the current version of fringe models will continue to be a part of the IPAC data reduction pipeline in ZTF-II. The current implementation of `fringe` can also be generalized to other instruments by simply altering the FITS extensions and header keywords that are currently chosen for ZTF images.

3.4 Measuring Improved Photometric Precision

3.4.1 The Uniform Background Indicator

In order to assess the effectiveness of our de-fringing method, we created a procedure for measuring correlated background noise. In this method, aperture photometry is taken

at random locations on the background of an image. At the location of each aperture the flux (F_{bkg}) and its associated measurement error (σ_{bkg}) are measured, resulting in a set of aperture fluxes ($\{F_{bkg}\}$) and a set of errors ($\{\sigma_{bkg}\}$). We then compare the standard deviation of background fluxes ($\text{std}(\{F_{bkg}\})$) to the average error on the background flux measurements ($\text{median}(\{\sigma_{bkg}\})$). We call the ratio of these two values the Uniform Background Indicator:

$$\Psi = \frac{\text{std}(\{F_{bkg}\})}{\text{median}(\{\sigma_{bkg}\})}. \quad (3.5)$$

For an image of uniform Gaussian noise, $\Psi \approx 1$ as the distribution in background fluxes will be equal to the average flux error across all measurements. Here there exists no global variance which is not captured by the local error term. For an image with correlated background noise, $\Psi > 1$ because the different background values sampled by the apertures will introduce additional variance into the numerator of Equation 3.5 not captured by the local error terms in the denominator. The flux errors in Equation 3.5 must be calculated locally and be in the same units as the flux measurement with respect to aperture area. The aperture size should be chosen to be approximately the PSF scale of the instrument so that UBI measurements will indicate what fringe signal (or any other source of correlated background noise) a particular instrument would observe.

We validated using the UBI as a measurement of correlated background noise through a controlled experiment. We generated images containing only Gaussian noise that were the same size as ZTF images. The error on each pixel value was set as the root-mean-square of the median-subtracted pixel values in a 10 pixel by 10 pixel box around that pixel. The size of this square had to be smaller than the spatial scale of the correlated background noise so that it contained minimal variation in the background value. We then laid 50,000 circular apertures onto the image and calculated the aperture flux and aperture flux error of each measurement. The UBI as outlined in Equation 3.5 was evaluated five times on each image producing an average UBI as well as an error on that measurement.

In Figure 3.5, we show the results of this experiment for varying aperture sizes. For all Gaussian images and aperture sizes we found $\Psi \approx 1$, confirming our interpretation that this value indicates no correlated background noise on an image. The location and scale of the Gaussian noise was also found to have no effect on this measurement. Figure 3.5 shows three example images with correlated background noise and the UBI values that they produce, covering the range of Ψ that we detected throughout this analysis. These example images show the correlation between a quantitatively larger Ψ and a qualitative increase in the appearance of atmospheric fringes.

Having verified that the UBI is a valid indicator of correlated background noise, we next measured the effect of removing atmospheric fringes on the UBI. A sample of g-band, r-band and i-band images was created by downloading one random image for each filter, readout-channel and field observed in the ZUDS survey from the week of 2020-02-01 to 2020-02-08 with a limiting magnitude greater than 19. This sampling method ensured a representative sample of airmasses and limiting magnitude for ZTF observations, while removing images

of extremely low quality from the sample. Our final sample contained 1408 g-band, 1407 r-band, and 1472 i-band images across a range of stellar densities. Each of the i-band images was cleaned using our method described in Section 3.3.2. The UBI was then calculated for each of the fringed images, as well as the cleaned images.

The distribution of the UBI for all of the images in our sample is shown in Figure 3.6, with $\Psi \geq 1$ as expected. The g-band and r-band images each have a nearly normal UBI distribution with medians of $\Psi = 1.13$ and $\Psi = 1.16$, indicating a small but measurable amount of correlated background noise remaining after flat-fielding. The i-band images contain far more correlated background noise and are notably bimodal, with one population averaging $\Psi = 1.33$ and a second wider distribution averaging a larger $\Psi = 1.91$. This split is caused by the location of the readout channel on the image plane. The inner 32 readout channels ($16 \leq \text{rcid} < 48$) with their additional layer of AR coating are less susceptible to atmospheric fringing. The outer 32 readout channels ($0 \leq \text{rcid} < 16$, $48 \leq \text{rcid} < 64$) experience significantly more fringing due to a lack of an additional layer of AR coating. Less than 20% of the images have a $\Psi < 1.72$ and the 20% most fringed images have a $\Psi \geq 2.20$. The cleaned i-band images show indistinguishable behavior between the inner and outer readout channels forming a single distribution averaging $\Psi = 1.15$. This population appears similar to the g-band and r-band populations, indicating successful removal of atmospheric fringes. However the g-band and r-band images have significantly longer tails with an 80th percentile of $\Psi = 1.28$ and $\Psi = 1.32$ respectively, compared to an 80th percentile for cleaned i-band images of $\Psi = 1.20$. This indicates that there is correlated background noise occurring in g-band and r-band images that PCA analysis could potentially model and remove. It is clear that the process of removing atmospheric fringes significantly reduces correlated background noise from i-band images.

It is reasonable to predict that the UBI will increase with airmass, as the presence of atmospheric fringes is caused from the column of atmosphere through which images are taken. This is found to be the case in Figure 3.7, where collecting the images into airmass bins and calculating the median UBI shows a trend toward larger UBI for larger airmass in i-band images. There is a slight increase in the g-band and r-band images as well, although the effect is relatively weak. The cleaned i-band images have comparable UBI values to the g-band and r-band images, again confirming the effectiveness of our method.

3.4.2 Photometric Error Due to Fringes

While the UBI is a useful measurement of the presence of correlated background noise, we sought to quantitatively measure the improvement in photometric precision that resulted from removing atmospheric fringes. First we measured the photometric error caused by atmospheric fringes relative to a reference catalog. Next we measured how de-fringing removes this error by injecting faint artificial sources into single images. Last we measured how de-fringing recovers lost signals by injecting extremely faint sources into images before combining them through co-addition. For each of these experiments we will demonstrate how

our method significantly improves the photometric precision of faint source measurements affected by atmospheric fringes.

First we measured the photometric error caused by atmospheric fringes. For each of the i-band images in our sample, we used SExtractor (Bertin & Arnouts 1996) to generate an instrumental photometric aperture catalog of astrophysical sources. We then cross-matched the original and cleaned images with Pan-STARRS1 (PS1) (Chambers et al. 2016) i-band catalogs downloaded from the VizieR database using astroquery (Ginsburg et al. 2019) to create cross-matched catalogs. A zeropoint was calculated for each image using matching sources with PS1 i-band magnitudes less than 17. This zeropoint was used to transform the instrumental photometric catalogs into ZTF i-band magnitudes comparable with PS1 i-band magnitudes, as well as to calculate a 5-sigma limiting magnitude for each image. After removing images with a limiting magnitude less than 21 to ensure reasonable measurement of faint sources, our final sample size was 738 images. For each pair of fringed and cleaned images, the cross-matched catalogs were used to find the variance in the difference between ZTF and PS1 magnitudes for stars within a magnitude bin. The photometric error caused by fringes was then calculated as:

$$\sigma_{mag} = \text{std}(m_{PS1} - m_{ZTF}) \quad (3.6)$$

$$\sigma_{fringe} = \sqrt{\sigma_{mag,fringed}^2 - \sigma_{mag,cleaned}^2} \quad (3.7)$$

The photometric error we measured due to fringing is shown in Figure 3.8, plotted separately for different faint magnitude bins against the UBI of the fringed images before cleaning. Images with larger UBI values have larger amounts of additional magnitude scatter relative to the PanSTARRS catalog than those images with smaller UBI values. Those images with $\Psi \leq 1.3$ appear to have marginal improvement to their photometric precision due to removing fringes. However, images with $\Psi > 1.3$ show significant photometric errors due to fringing that our method removes. Images with $\Psi = 1.75$ have a 0.21 magnitude error on 19.5 magnitude stars, 0.35 magnitude error on 20.5 magnitude stars and a 0.43 magnitude error on 21.5 magnitude stars. Images with $\Psi = 2.0$ have a 0.28 magnitude error on 19.5 magnitude stars, 0.46 magnitude error on 20.5 magnitude stars and a 0.57 magnitude error on 21.5 magnitude stars. The worst 10% of i-band images in the outer readout channels have $\Psi \geq 2.4$, resulting in errors as large as 0.21 magnitudes for 18.5 magnitude stars and up to 0.39 magnitudes for 19.5 magnitude stars. Images with $\Psi \geq 2.3$ have very few sources fainter than 19.5 magnitude in the figure, indicating that these images fail to pass the limiting magnitude cut. Our method enables the recovery of the faint end of the luminosity function that would otherwise be unobservable in ZTF i-band images.

3.4.3 Effects on Fake Sources: Single Epoch

Another way to measure the effect of atmospheric fringes on photometric precision of faint sources is through the injection and recovery of fake sources. We selected an i-band image with $\Psi = 1.78$ as a representative image. The image's zeropoint and 5-sigma limiting

magnitude were calculated using a cross-match to PanSTARRS1 sources as described above. A PSF for the image was derived using `psfex` (Bertin 2011) and executed using wrappers written in the `galsim` (Rowe et al. 2015) python package. 100 sources with magnitudes equal to the 5-sigma limiting magnitude were injected into the original image using this PSF model at random locations, including Poisson noise. The image was then cleaned using the `fringe` package. Photometric catalogs using both aperture and PSF photometry were calculated on the original and cleaned images using `SExtractor` with a 3-pixel 1-sigma detection threshold. Measurements in the catalogs at the location of the injected fake sources were then recovered for comparison to the true modeled flux. We also calculated the ideal aperture corrected flux of the injected fake sources by calculating the median ratio of aperture to PSF fluxes of high signal-to-noise astrophysical sources from the catalogs, for both fringed and cleaned images. We then repeated this process 50 times for a total of 5000 fake sources injected to form our sample. We also duplicated this experiment with the `SExtractor BACKGROUND` parameter set to `LOCAL` and `GLOBAL` to test the affect of forcing the source identification algorithm to attempt to characterize local variations in the background noise. Figure 3.9 shows the results of this experiment. The distributions show the fractional offset of the measured flux to the injected flux scaled by the theoretical error on a 5-sigma source. Plots are drawn in log-scale to highlight the long tail of overestimated measurements for images that have not been cleaned, with a Gaussian distribution drawn in gray as a reference.

We first note some observations about the quality of recovered fake sources on the images which have not been cleaned. Setting the `SExtractor BACKGROUND` parameter to `LOCAL` improves photometric measurements on these fringed images. `GLOBAL` computes the background flux across the entire image and underestimates the amount of background underneath a source sitting on a bright fringe. These sources will have overestimated brightness and form the long tail of measurements with larger than expected fluxes. `LOCAL` computes the background flux with a rectangular annulus around the source that prevents attributing the additional brightness to the source but instead to the background. Switching to a `LOCAL` background therefore reduces this tail. Aperture photometry performs better than the PSF photometry on these fringed images with a reduced tail for both `SExtractor` settings. This prominent tail may be due to a bias in the PSF model caused by including stars that fall on fringes, inflating the wings of the PSF model. When that model is applied to stars that fall on particularly bright fringes the increased background may be included in the measured flux, resulting in an overestimate of the flux. Aperture photometry does not suffer from this potential model bias and therefore has a less prominent tail of overestimated fluxes.

The sources from the cleaned catalogs are more accurately and precisely measured under all conditions. There is a slight tail of overestimated fluxes when performing PSF photometry but it is far closer to a Gaussian distribution than without cleaning. For both `LOCAL` and `GLOBAL` backgrounds, the distribution of recovered sources in the cleaned catalogs very nearly resembles a Gaussian distribution. There exists only a few sources with measured fluxes exceeding the injected flux causing a deficit of fainter sources. There still remains a noticeable deficit of sources at the faint end of the brightness distribution, indicating that our method is not removing all additional flux in the image background. Aperture photometry remains

the best way to evaluate sources even where atmospheric fringes have been removed. Failure to clean images containing atmospheric fringes results in a systematic overestimation of the flux of faint sources. Applying our method enables near-ideal recovery of 5-sigma sources for a variety of measurement methods.

3.4.4 Effects on Fake Sources: Multi Epoch

It is particularly difficult to overcome the effects of fringes when combining multiple images of the same field to recover sources fainter than a single image’s limiting magnitude. Changing atmospheric conditions and various observational airmasses will alter the strength of the fringes, while dithering and sky motion will place astrophysical sources onto slightly different pixels for each exposures. Failure to remove fringes will contribute significant excess flux to a co-addition of multiple images. We demonstrate here this effect quantitatively by injecting extremely faint sources into individual images and attempting to measure them in a co-added image.

100 i-band images of the same field were zero-pointed using the previously outlined method. Sources in each image were measured and the signal-to-noise versus magnitude was fit to extrapolate the 0.5 sigma magnitude of all images. Each image was injected with 100 sources at this 0.5 sigma magnitude at a common list of coordinates in right ascension and declination. These sources would be expected to appear as 5-sigma sources after combining the 100 i-band images because signal-to-noise increases as the square root of the total exposure time. All images were then cleaned with the `fringe` package. `SCAMP` (Bertin 2006) was run to find astrometric projection parameters for each of the images such that they could be transformed onto a common reference frame. All fringed and cleaned images were then combined into two separate co-additions using `SWarp` (Bertin et al. 2002), produced with a median combination filter and background subtraction. Aperture photometry catalogs were then generated on the final co-additions using `SExtractor` with the `LOCAL` background setting. Sources at the locations of the injected signals were recovered and their signal-to-noise measured as the ratio of their aperture flux to their locally determined aperture flux error. We repeated this process 50 times for a total of 5000 fake sources distributed over 50 fringed and 50 cleaned co-additions.

Figure 3.10 shows the quantitative and qualitative photometric improvements to these recovered sources. Sources observed on fringed co-additions peak at a signal to noise of 3, pushing them below the typical 5-sigma observable threshold. These sources also have a long tail of excessive flux stretching as high as 10 sigma due to the additional flux caused by atmospheric fringes. Sources observed on cleaned co-additions are much closer to their theoretical distribution. These sources peak at exactly 5-sigma and the large majority of sources are within the normal signal-to-noise range of 4 to 6. There still exists a tail of excessive fluxes, demonstrating that our method is failing to remove all excessive correlated background from the individual exposures. However our method clearly produces an observed flux distribution that is much closer to a Gaussian distribution.

The largest difference between the two populations appears in their yields. Of the 5000 sources injected into each of the fringed and cleaned images, 1161 sources were recovered from cleaned co-additions and only 156 sources were recovered from fringed co-additions. Over 95% of the 0.5 sources injected into the fringed images failed to be recovered after co-addition. The order of magnitude increase in the number of sources recovered by our method enables ZTF i-band surveys to recover faint sources that would otherwise have been extremely unlikely to observe. We note that forced photometry at the known locations of the fake sources may have increased the recovered yields for both populations but would not be an accurate representation of the observation process undertaken for unknown sources.

Our analysis demonstrates the power of our method to remove photometric error due to atmospheric fringes and enable the recovery of faint sources in both single images and co-additions that would otherwise have been undetectable. Failure to implement a method for removing atmospheric fringes greatly reduces the effectiveness of the i-band filter for observing any sources fainter than 18th magnitude. Our method increases the photometric precision of the i-band to that of g-band and r-band images that do not suffer from atmospheric fringes.

3.5 Discussion

Our method has several benefits beyond improving the photometric quality of individual i-band images. Image references are constructed through the combination of as many as 40 epochs taken at a single field. As shown in Section 3.4.4, removing fringes significantly improves the photometric precision on a multi-epoch co-addition. Application of our method also greatly improves the quality of the ZTF alert stream. The alert stream packets are generated on difference images created from subtracting individual epochs from a multi-epoch co-addition reference (after appropriate scaling). Correctly removing fringes makes visible the faint end of the luminosity function that would otherwise not be observable in the i-band alert stream.

Future work on removing atmospheric fringes using PCA could improve upon our method in several ways. Training on CCDs instead of readout channels could produce better eigen-images by including more correlated pixels in the PCA feature identification. Also, different methods of calculating the local flux error such as calculating the root-mean-square error on a smaller (or larger) square around each pixel could produce a less noisy UBI measurement. Lastly, investigating the eigenvalues for fringe bias images generated on different readout channels from the same exposure could reveal correlations that could be used to perturb the fringe bias into a better fit for each readout channel. Treating each readout channel as entirely independent, while convenient and a natural fit for the IPAC processing pipeline, may leave out valuable information that could improve our method.

The eigen-images shown in Figure 3.4, as well as the eigen-images of many of the other fringe models, show significant smooth variations that are not being removed by the current flat fielding pipeline. This indicates that the fringe bias images include not only atmospheric

fringes, but residual global gradients. Future work could be done on exploring the application of this PCA method to supplement or even replace the current flat fielding pipeline on not only the i-band images, but g-band and r-band images as well. The Uniform Background Indicator can be used as a quantitative measurement to compare how well a PCA method, as compared to more classical flat fielding methods, generates astronomical images with normal backgrounds.

Acknowledgments

This work is based on observations obtained with the Samuel Oschin Telescope 48-inch and the 60-inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. ZTF is supported by the National Science Foundation under Grant No. AST-1440341 and a collaboration including Caltech, IPAC, the Weizmann Institute for Science, the Oskar Klein Center at Stockholm University, the University of Maryland, the University of Washington, Deutsches Elektronen-Synchrotron and Humboldt University, Los Alamos National Laboratories, the TANGO Consortium of Taiwan, the University of Wisconsin at Milwaukee, and Lawrence Berkeley National Laboratories. Operations are conducted by COO, IPAC, and UW.

We acknowledge support from the University of California Office of the President for the UC Laboratory Fees Research Program In-Residence Graduate Fellowship (Grant ID: LGF-19-600357). This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. We acknowledge support from the DOE under grant DE-AC02-05CH11231, Analytical Modeling for Extreme-Scale Computing Environments. M. W. Coughlin acknowledges support from the National Science Foundation with grant number PHY-2010970.

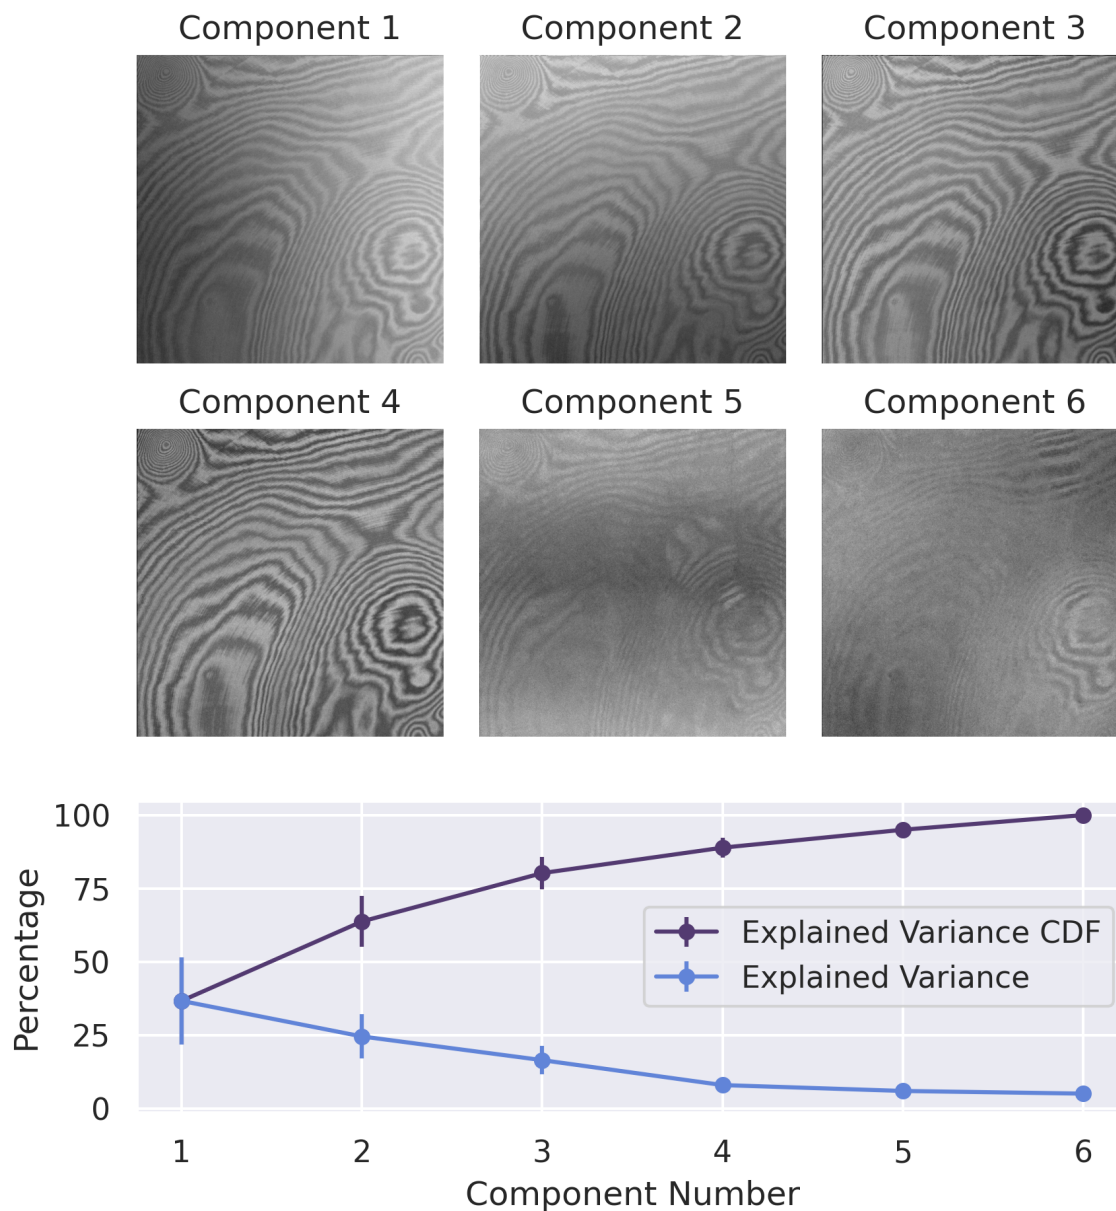


Figure 3.4: Each readout channel has a distinct fringe model with six eigen-images constructed from the reduction of thousands of training fringe maps. Here is an example of the eigen-images from readout channel 13 (top). While the fringe models were trained on a 1-dimensional array of pixels, the 2-dimensional fringes remain intact. The fringe pattern of the readout channel is clearly evident with slight variations in position and strength amongst the first four eigen-images. However the last two components (and to some extent the first four) contain smooth global gradients that would ideally be removed by flat fielding. The fractional explained variance for each component across all 64 fringe models (bottom) show that the fifth and sixth components captured far less variance in the training sample. The first component captured 64.6% of the pixel variance while the sixth component captured only 5.0% of the pixel variance. In total six PCA components reconstructs 95.0% of the variance seen in our training sample. We therefore chose to have six components in our models.




Figure 3.5: The Uniform Background Indicator is $\Psi \approx 1$ for a range of aperture sizes on images of Gaussian noise (top). This validates interpreting $\Psi \approx 1$ as measuring no correlated background noise in an image. Three example i-band images and their Ψ values are shown (bottom). More prominent atmospheric fringes correlates with an increase in the value of Ψ .




Figure 3.6: Our sample of g-band, r-band and i-band images show distinctively different distributions in their Uniform Background Indicator (Ψ) as shown in both a histogram (top) and cumulative distribution function (bottom). The g-band (green) and r-band (red) average $\Psi \approx 1.15$ indicating a small but measurable amount of correlated background noise after flat fielding. The i-band is split between two populations. The images taken on the inner 32 readout channels (light blue) are only moderately affected by atmospheric fringing, averaging $\Psi = 1.33$. However the outer 32 readout channels (dark blue) are significantly affected by these fringes, with an average of $\Psi = 1.91$ and less than 20% of the images with $\Psi < 1.72$. The population of cleaned i-band images (yellow) is monomodal and has a median value similar to the g-band and r-band of $\Psi = 1.15$. ZTF i-band images processed with our method show similar amounts of correlated background noise as is present in g-band and r-band images.




Figure 3.7: The Uniform Background Indicator (Ψ) correlates with airmass for i-band images with significant amounts of atmospheric fringes. Observing sources through the additional column of atmosphere increases the exposure to stimulated emission of atmospheric lines that causes the emergence of fringes. Cleaning the i-band images removes this correlation and produces a relationship with airmass indistinguishable from the g-band and r-band populations.




Figure 3.8: Measurement of the photometric error on faint sources due to fringing on 738 high quality i-band images as shown against the Uniform Background Indicator (Ψ) of the fringed image before cleaning. This error is calculated by comparing the variance in the difference between ZTF and PS1 magnitudes before and after removing fringes. Larger UBI values correlate with larger amounts of photometric error, getting as large as 0.46 magnitudes for 20.5 magnitude sources and 0.57 magnitudes for 21.5 magnitude sources at $\Psi = 2.0$. Atmospheric fringes add a significant systematic error to the photometry of faint sources that our method is able to remove.

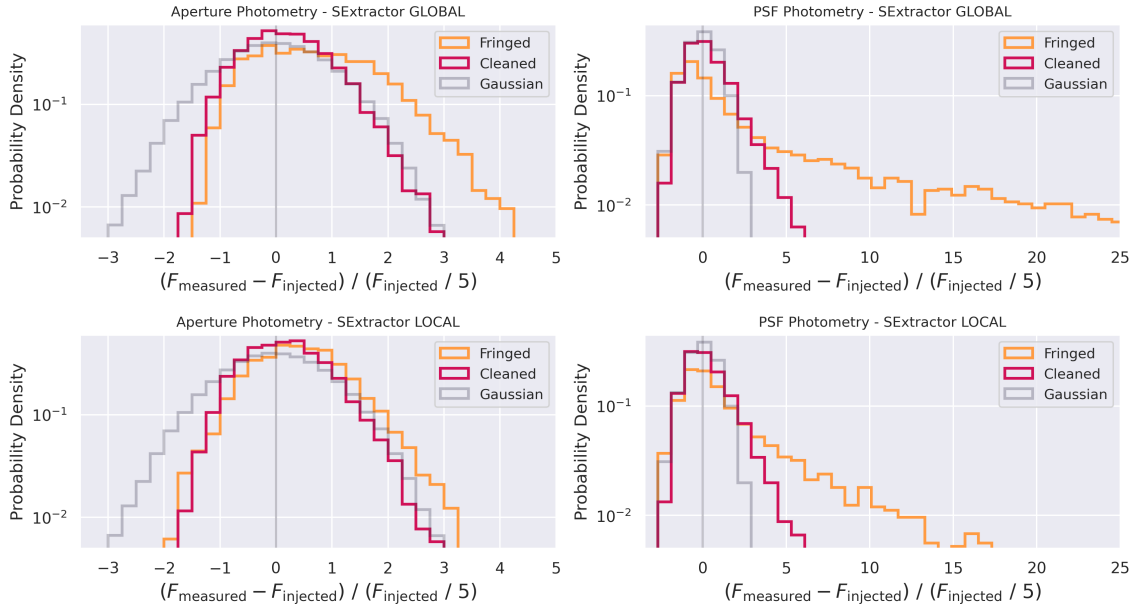


Figure 3.9: The measured flux of 5000 fake 5-sigma sources injected into an i-band image on the fringed images (orange) and after cleaning (red) using the `SEExtractor GLOBAL` background setting (top), `LOCAL` background setting (bottom), aperture photometry (left) and PSF photometry (right). On each sub-figure a Gaussian of 5000 sources is drawn (gray) as a visual guide. In all cases the cleaning method significantly increases the accuracy and precision of the recovered flux, particularly for PSF photometry where attempting photometry on images with fringes can often result in overestimating the brightness of the source. Aperture photometry also overestimates the flux resulting in a deficit of lower flux detections than would be statistically expected, although to a lesser degree. If a method for removing fringes cannot be applied, it is best to use the `SEExtractor LOCAL` background setting and a catalog of aperture photometry to most accurately measure the true magnitude of faint sources.

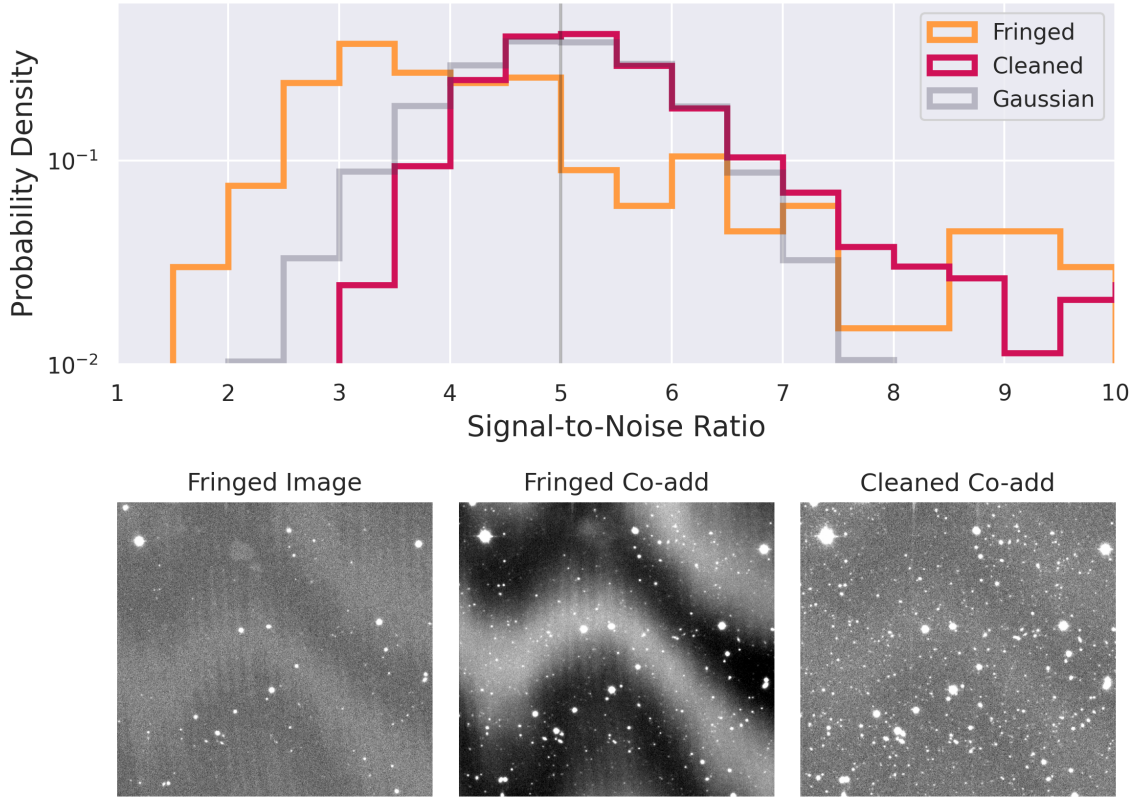


Figure 3.10: The signal-to-noise distribution of 5000 0.5 sigma sources injected into 100 images after a median co-addition (top) for images with atmospheric fringes (orange) and after cleaning (red). The photometric catalogs were generated with a `SExtractor LOCAL` background and aperture photometry. 1161 sources were recovered in the cleaned images and only 156 sources were recovered in the fringed images, demonstrating the necessity to clean i-band images in any attempt to find faint sources after multi-epoch co-addition. Those sources that were recovered in the fringed images are most likely to be detected at an artificially low signal-to-noise, with a long tail of higher signal-to-noise due to falling on positive fringes. The cleaned images have a distribution much closer to Gaussian, with a small deficit at the low signal-to-noise end that also appears as a tail at the higher end. The photometric improvement in multi-epoch co-additions can be clearly seen in as a smoothly varying background and the presence of faint sources after cleaning has been applied (bottom).

Chapter 4

Transformations from Pan-STARRS1 and UBV Filters into ZTF Filters

This is a replication of the article published as [Medford, Lu, & Schlafly 2020c](#), RNASS, 4, 38.

4.1 Introduction

The Zwicky Transient Facility (ZTF) is an optical time-domain survey that has been operating on the 48-inch Samuel Oschin Telescope at Palomar Observatory since March 2018 ([Bellm et al. 2018](#)). ZTF surveys the visible Northern sky every few nights in ZTF g-band and r-band, as well as an i-band filter used only for partnership observations and thus absent from this analysis. Transformations between the photometric systems of ZTF, Pan-STARRS1 (PS1) ([Chambers et al. 2016](#)), and Johnson-Morgan-Cousins (UBV) ([Johnson et al. 1966](#)) are essential for extending both the time-baseline and wavelength coverage of ZTF catalogs with information from other catalogs. We use cross-matching catalogs and simple stellar population (SSP) synthesis models to derive photometric transformations from PS1 and UBV to ZTF.

4.2 Pan-STARRS1 to ZTF Photometric Conversions

We selected g-band and r-band photometric catalogs from ZTF observations at random locations throughout the sky. These stars were cross-matched to PS1 mean magnitudes in the DR1 catalog. Saturated stars and stars with significant noise were removed from the sample by applying a cut of $12 \text{ mag} < m_{\text{PS}} < 19 \text{ mag}$ for each band. We also removed any sources with bad PS1 photometry in any filter. We calculated a linear fit to the difference of each ZTF filter to its corresponding PS1 filter as compared to the star’s PS1 g minus r color. We find

$$\begin{aligned} m_{\text{ZTF,g}} - m_{\text{PS,g}} &= 0.055(m_{\text{PS,g}} - m_{\text{PS,r}}) - 0.012 \\ m_{\text{ZTF,r}} - m_{\text{PS,r}} &= -0.087(m_{\text{PS,g}} - m_{\text{PS,r}}) - 0.0035. \end{aligned}$$

To estimate the error of our method, we re-sampled the sky and calculated the difference between the ZTF magnitude calculated from our transformation and the star’s magnitude as recorded in the ZTF catalog. We measure the difference between ZTF magnitudes transformed from PS1 and true ZTF magnitudes for stars with $12 \text{ mag} < m_{\text{PS}} < 19 \text{ mag}$ to be

$$\begin{aligned} \sigma_{\text{PS-TO-ZTF, g}} &= 0.114 \text{ mag} \\ \sigma_{\text{PS-TO-ZTF, r}} &= 0.063 \text{ mag}. \end{aligned}$$

Because we are able to use a large number of stars, the formal uncertainties on the fit parameters are negligible, though systematic effects in the surveys’ photometric calibrations or through-puts may lead to slight variations over the sky. The left subplot in Figure 4.1 shows our sample and the linear fit that well captures the trend in the data. Note that this is a transformation from PS1 to ZTF apparent magnitudes.

4.3 UBV to ZTF Photometric Conversions

PyPopStar (Jr. et al. 2020) is a python package that generates single-age, single-metallicity populations from user specified initial mass functions, stellar evolution models, and stellar atmospheres. PyPopStar contains transmission curves for all UBV filters, as well as ZTF g-band and r-band filters as a result of this work, and produces photometric observations of stars in a synthetic cluster. We produced clusters with

- Initial Mass Function: (Kroupa 2001) ($0.08M_{\odot} < M_{*} < 120M_{\odot}$)
- Cluster Mass: $10^5 M_{\odot}$
- Age: 8×10^9 years
- Metallicity: $[M/H] = 0$

- Evolution: MIST v1.2 (Choi et al. 2016)
- Atmospheres: ATLAS9 (Castelli & Kurucz 2003), PHOENIXv16 (Husser et al. 2013), BARAFFE (Baraffe et al. 2015)
- Reddening Law: Damini et al. 2016
- Extinction: $0.0 \text{ mag} < A_{K_s} < 1.0 \text{ mag}$

We chose a grid of extinctions from $0.0 \text{ mag} < A_{K_s} < 1.0 \text{ mag}$, which corresponds to a $A_V \lesssim 10 \text{ mag}$, in order to fully explore high extinctions found throughout the Galactic plane and towards the Galactic bulge. After executing these simulations, we extracted the UBVR (B, V, R) and ZTF (g, r) absolute magnitudes for stars in the resulting clusters. This photometry was placed onto a 2-D plane of $m_{UBV,B} - m_{UBV,V}$ versus $m_{UBV,V} - m_{UBV,R}$ resulting in stars separating without significant overlap. We calculated the difference between the ZTF magnitude and a selected UBVR magnitude for each star: $m_{ZTF,r} - m_{UBV,R}$ and $m_{ZTF,g} - m_{UBV,V}$. The choice to calculate the offset of $m_{ZTF,g} - m_{UBV,V}$ instead of the perhaps more intuitive $m_{ZTF,g} - m_{UBV,B}$, avoiding interpolation artifacts. A linear interpolation was fit to the UBVR photometry in this 2-D plane, calculated separately for each of the ZTF filters. We applied a nearest extrapolation to points that fell outside the convex hull of our data but within the range of $0 < m_{UBV,V} - m_{UBV,R} < 6$ and $0 < m_{UBV,B} - m_{UBV,V} < 6$. We believe the region outside of the convex hull of our data to be largely un-physical but chose to extrapolate into this region to provide approximate solutions.

We estimate the error of our method by randomly sampling the transformation via bootstrap sampling. We generated new synthetic clusters with both UBVR and ZTF photometry at different extinctions and compared the ZTF magnitude predicted by our method against artificial photometry measured by PyPopStar. The resulting spread in our magnitude differences came out to be

$$\begin{aligned}\sigma_{UBV\text{-TO-ZTF, g}} &= 0.023 \text{ mag} \\ \sigma_{UBV\text{-TO-ZTF, r}} &= 0.020 \text{ mag}.\end{aligned}$$

Any star with a measured $m_{UBV,B}$, $m_{UBV,V}$, $m_{UBV,R}$ magnitude can be placed on these planes, producing a ZTF g-band and r-band magnitude with relatively small errors. These planes, along with a set of python functions that executes the UBVR to ZTF transformations, can be found at https://portal.nersc.gov/project/uLens/ZTF_Photometric_Transformations. Note that this is a transformation from UBVR to ZTF absolute magnitudes.

Acknowledgements: This work was done with support from the University of California Office of the President for the UC Laboratory Fees Research Program In-Residence Graduate Fellowship (Grant ID: LGF-19-600357). Part of this work was performed under the auspices

of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and was supported by the LLNL-LDRD Program under Project Number 17-ERD-120.

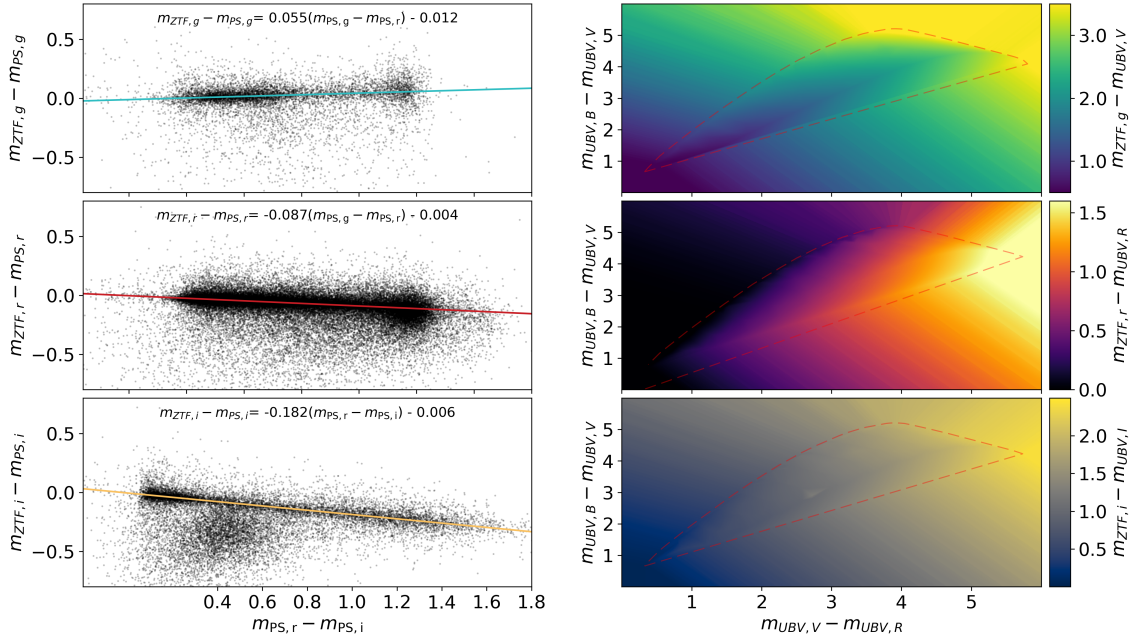


Figure 4.1: Results for transforming from the PS1 and UBV photometric systems to ZTF filters. The PS1 to ZTF transformation (left) is derived from cross-matching stars observed on ZTF, limiting the sample to only stars with $12 < m_{PS,g/r} < 19$, and correlating the difference in ZTF and PS1 magnitude with the PS1 color. ZTF observations for these fits were drawn from random dates and sky locations. The UBV to ZTF transformation (right) shows a 2-D color relationship derived from simple stellar population synthesis models. Applying a linear interpolation to photometry of synthetic clusters with varying amounts of extinction produce these relationships. Points outside of the convex hull (dashed red) of our synthetic data are extrapolated and occupy largely un-physical regions. Functions for executing a photometric transformation with these 2-D planes can be found at https://portal.nersc.gov/project/uLens/ZTF_Photometric_Transformations.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors

expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Chapter 5

128 Microlensing Events from the Three Years of Zwicky Transient Facility Phase One

Abstract

Microlensing events have historically been discovered throughout the Galactic bulge and plane by surveys designed solely for that purpose. We conduct the first multi-year search for microlensing events on the Zwicky Transient Facility (ZTF), an all-sky optical synoptic survey that observes the entire visible Northern sky every few nights. We discover 128 high quality microlensing events in the three years of ZTF-I using the bulk lightcurves in the ZTF Public Data Release 5. We find an excess of long-duration Galactic plane events when comparing our results to both previous surveys and simulations. 35 of our events are found outside of the Galactic plane ($|b| \geq 15^\circ$), nearly tripling the number of previously discovered events in the stellar halo from surveys pointed toward the Magellanic Clouds and the Andromeda Galaxy. We also record 1690 ongoing candidate events as potential microlensing that can continue to be observed by ZTF-II for identification. The scalable and computationally efficient methods developed in this work can be applied to future synoptic surveys such as the Rubin Observatory's Legacy Survey of Space and Time and the Nancy Grace Roman Space Telescope as they attempt to find microlensing events in even larger and deeper datasets.

5.1 Introduction

Einstein (1936) first derived that objects with mass could bend the light of a luminous star on its way to an observer and produce multiple images of the background star. This phenomenon is called gravitational microlensing because the individual images are separated by microarcseconds and are therefore unresolvable to any instrument. The observer can only measure a lightcurve that includes both the constant luminosity of the lens and the apparent increase in the brightness of the background star (Refsdal & Bondi 1964). Microlensing events will last as long as there is an apparent alignment, occurring on timescales ranging from days to years for one star magnifying another within the Milky Way galaxy (B. Paczyński 1986).

The amplification of the luminous source is maximized at the time (t_0) when the impact parameter (u_0) between the location of the source and lens is minimized. The duration of this amplification, named the Einstein crossing time (t_E), is approximately the time during which the length-scale of the lens' multiple images, named the Einstein radius (θ_E), is traversed by the relative proper motion between the lens and source (μ_{rel}),

$$t_E = \theta_E / \mu_{\text{rel}} .$$

For a more thorough derivation of these parameters we refer the reader to Gould (1992).

Microlensing can be distinguished from other astrophysical transients due to several unique characteristics. The amplification of a source is ideally followed by a decrease in the source's brightness that is symmetric across the point of maximum amplification when a constant velocity for the lens and source is assumed. However this symmetry is complicated by the non-uniform motion of the Earth as it orbits around the sun (Gould 1992). This produces a sinusoidal variation to the source amplification with an amplitude determined by the relative distances between the observer to the source (d_S) and the observer to the lens (d_L) with a period of one year called microlensing parallax (π_E),

$$\begin{aligned} \pi_E &= \pi_{\text{rel}} / \theta_E , \\ \pi_{\text{rel}} &= 1/d_L - 1/d_S . \end{aligned}$$

Either the symmetry of the observable lightcurve when there is low parallax or the presence of a one-year sinusoidal variation to the amplification where there is high parallax is evidence of microlensing.

The amplification of the background star is achromatic as the geometric effects of bending space-time occur equally for all wavelengths of light. However the constant presence of the lens light in the lightcurve will produce an chromatic effect called blending (Stefano & Esin 1995). Imagine a source star that is entirely one flux unit of wavelength of red and a lens star of equal flux that is entirely one wavelength of blue. Before (and after) the microlensing event the brightness measured in corresponding red and blue filters would be equal to each other, but during the event the red filter would increase in brightness while the blue filter measured a constant flux. This is further complicated by the presence of multiple neighbor stars not close enough to be involved in the microlensing phenomenon but within the point-spread

function of the instrument. These stars contaminate the lightcurve and also contribute additional unamplified light that causes chromatic blending. Together these effects produce a chromatic signal that can be modelled as a combination of unamplified and amplified light. The fraction of light that originates from the source (as compared to the light from the unamplified lens and neighbors) is given the name source-flux-fraction of b_{sf} .

The probability of two stars crossing the same line of sight is proportional to the apparent stellar density. This has motivated most microlensing surveys to look for events toward the Galactic bulge (Sumi et al. 2013; Udalski et al. 2015a; Navarro et al. 2017; Kim et al. 2018a; Mróz et al. 2019; Golovich et al. 2020), as well as the Magellanic clouds (Alcock et al. 2000; Tisserand et al. 2007; Wyrzykowski et al. 2011b,c) and M31 (Novati et al. 2014; Niikura et al. 2019). These surveys discover hundreds to thousands of microlensing events per year. There have been fewer microlensing surveys dedicated to looking for these events throughout the Galactic plane. The Expérience pour la Recherche d’Objets Sombres (EROS) spiral arm surveys (Rahal et al. 2009) observed 12.9 million stars across four directions in the Galactic plane and discovered 27 microlensing events over seven years between 1996 and 2002. The Optical Gravitational Lensing Experiment (OGLE; Mróz et al. (2020c)) discovered 630 events in 3000 square degrees of Galactic plane fields in the 7 years between 2013 and 2019. They measured a threefold increase in the average Einstein crossing time for Galactic plane events as compared to the Galactic bulge and an asymmetric optical depth that they interpret as evidence of the Galactic warp. Mróz et al. (2020a) found 30 microlensing candidates in the first year of Zwicky Transient Facility (ZTF) operations.

While the number of events discovered thus far are much smaller than in Galactic bulge fields, simulations predict that there are many more events left to discover in large synoptic surveys. Sajadian & Poleski (2019) predict that the Vera Rubin Observatory Legacy Survey of Space and Time (LSST), can discover anywhere between 7900 and 34000 microlensing events over ten years of operation depending on its observing strategy. Chapter 5 estimate that ZTF would discover ~ 1100 detectable Galactic plane microlensing events in its first three years of operation, with ~ 500 of these events occurring outside of the Galactic bulge ($\ell \geq 10^\circ$). Microlensing events throughout the Galactic plane can yield interesting information about Galactic structure and stellar evolution that we cannot learn from only looking toward the Galactic bulge.

Gravitational lenses only need mass but not luminosity in order to create a microlensing event, thereby making microlensing the only method available for detecting a particularly interesting non-luminous lens object: black holes. Inspired by the suggestion by B. Paczyński (1986) that constraints on the amount of dark matter found in MASSive Halo Compact Objects (MACHOs) could be measured by microlensing, there had been ongoing efforts to find these MACHOs, including black holes, as gravitational lenses for many years. The MACHO (Alcock et al. 2000) and EROS (Afonso et al. 2003) projects calculated these upper limits after several years of observations toward the Galactic bulge and Magellanic clouds. While black hole candidates have been proposed from individual microlensing lightcurves (Bennett et al. 2002; Wyrzykowski et al. 2016), there is a degeneracy between the mass of the lens and the distances to the source and lens that cannot be broken through photometry

alone. Essentially photometry cannot distinguish between a massive lens that is relatively distance and a less massive lens that is closer to the observer. [Lu et al. \(2016\)](#) outlined how direct measurement of the apparent shift in the centroid of the unresolved source and lens, or astrometric microlensing, can be used to break this degeneracy and confirm black hole microlensing candidates. This shift is extremely small, on the order of miliarcseconds, and therefore requires high resolution measurements on an adaptive optics system such as Keck AO. Several candidates have been followed up using this technique ([Arredondo et al. 2019](#)) but no significant black hole detections have been confirmed. Only a few can be rigorously follow up astrometrically and the astrometric signal of events found toward the bulge is at the edge of current observational capabilities. Simulations indicate that black hole candidates can be ideally selected from a list of microlensing events by searching for those with larger Einstein crossing time and smaller microlensing parallaxes ([Lam et al. 2020](#)). Furthermore simulations indicate that events in the Galactic plane have larger microlensing parallaxes and astrometric shifts overall that can make finding black hole candidates amongst these surveys easier than surveys pointed toward the Galactic bulge ([Medford et al. 2020a](#)).

In this paper we will conduct a search for microlensing events in the Zwicky Transient Facility dataset that is optimized to find these black hole candidates, increasing the number of candidates that can be astrometrically followed up for lens mass confirmation. In [Section 5.2](#) we discuss the Zwicky Transient Facility telescope, surveys and data formats. In [Section 5.3](#) we outline the software stack we have developed to ingest and process lightcurve data for microlensing detection. In [Section 5.4](#) we step through the detection pipeline constructed to search for microlensing in ZTF’s all sky survey data. In [Section 5.5](#) we analyze the results of our pipeline and share our list of microlensing candidates. We conclude with a discussion in [Section 5.6](#).

5.2 The Zwicky Transient Facility

5.2.1 Surveys and Data Releases

The Zwicky Transient Facility (ZTF) began observing in March 2018 as an optical time-domain survey on the 48-inch Samuel Oschin Telescope at Palomar Observatory ([Bellm et al. 2019c](#); [Graham et al. 2019](#)). In the almost three years of Phase-I operations, ZTF has produced one of the largest astrophysical catalogs in the world. Nightly surveys were carried out on a 47 square degree camera in ZTF g-band, r-band and i-band filters averaging $\sim 2.0''$ FWHM on a plate scale of $1.01'' \text{ pixel}^{-1}$. The surveys during this time were either public observations funded by the National Science Foundation’s Mid-Scale Innovations Program, collaboration observations taken for partnership members and held in a proprietary period before later being released to the public, or programs granted by the Caltech Time Allocation Committee ([Bellm et al. 2019b](#)).

Three surveys of particular interest for microlensing science are the Northern Sky Survey and the Galactic Plane Survey, both public, and the partnership High-Cadence Plane Survey.

The Northern Sky Survey observed all sky north of -31° declination in g-band and r-band with an inter-night cadence of three days, covering an average of 4325 deg^2 per night. The Galactic Plane Survey (Prince & Zwicky Transient Facility Project Team 2018) observed all ZTF fields falling within the Galactic plane ($-7^\circ < b < 7^\circ$) in g-band and r-band every night that the fields were visible, covering an average of 1475 deg^2 per night. These two public surveys have been run continuously since March 2018. The collaboration High-Cadence Plane Survey covered 95 deg^2 of Galactic plane fields per night with 2.5 hour continuous observations in r-band totaling approximately 2100 deg^2 . ZTF Phase-II began in December 2020 with a public two-night cadence survey of g-band and r-band observations of the Northern sky dedicated to 50% of available observing time.

The resulting data from these surveys is reduced and served to the public by the Infrared Processing and Analysis Center (IPAC) in different channels tailored to different science cases (Masci et al. 2018). IPAC produces seasonal data releases (DRs) containing the results of the ZTF processing pipeline taken under the public observing time and a limited amount of partnership data. The products included in these DRs include instrumentally calibrated single-exposure science images, both point-spread-function and aperture source-catalogs from these individual exposures, reference images constructed from a high quality set of exposures at each point in the visible sky, an objects table generated from creating source-catalogs on these reference images, and lightcurves containing epoch photometry for all sources detected in the ZTF footprint. There have been DRs released every three to six months, starting with DR1 on 2019 May 8. As discussed in more detail in Section 5.3.2, this work exclusively uses data publicly available in DR3, DR4 and DR5 released on 2021 March 31. Additional details about the content and structures of products in these data releases can be found at <https://www.ztf.caltech.edu/page/dr5>.

These observations cover a wide range of Galactic structure in multiple filters with almost daily coverage over multiple seasons. ZTF’s capability to execute a wide-fast-deep-cadence opens an opportunity to observe microlensing events throughout the Galactic plane that microlensing surveys only focused in the Galactic bulge cannot observe with equivalent temporal coverage. Our previous work simulating microlensing events observable by ZTF (Medford et al. 2020a) indicated that there exists a large population of microlensing events both within and outside of the Galactic bulge yet to be discovered. We seek to find black hole microlensing candidates amongst these events that could be later confirmed using astrometric follow up.

5.2.2 Object Lightcurves

While the real-time alert stream is optimized for events with timescales of days to weeks, black hole microlensing events occur over months or even years. Fitting for microlensing events requires data from both a photometric baseline outside of the transient event and the period during which magnification occurs ($t_0 \pm 2t_E$). The data product therefore most relevant to our search for long-duration microlensing events are the data release lightcurves containing photometric observations for all visible objects within the ZTF footprint.

Filter	Accessible Sky-Coverage	$N_{\text{lightcurves}}$	$N_{\text{lightcurves}}$ with $N_{\text{obs}} \geq 20$
g-band	97.66%	1,226,245,416	582,677,216
r-band	98.31%	1,987,065,715	1,107,250,253
i-band	51.88%	346,398,848	78,425,164

Table 5.1: Object Statistics for Public Data Release 5. Additional details can be found at <https://www.ztf.caltech.edu/page/dr5>.

Data release lightcurves are seeded from the point spread function (PSF) source-catalogs measured from co-added reference images. PSF photometry measurements from each single-epoch image are appended to these seeds where they occur in individual epoch catalogs. A reference image can only be generated where at least 15 good-quality images are obtained, limiting the locations in the sky where lightcurves can be found to parts of the northern sky observed during good weather at sufficiently low airmass. This sets the limit on the declination at which the Galactic plane is visible in DR lightcurves and consequently our opportunity to find microlensing events in these areas of the sky. Statistics regarding the lightcurve coverage for nearly a billion objects have been collated from the DR5 overview into Table 5.1.

Observation fields are tiled over the night sky in a primary grid, with a secondary grid slightly shifted to cover the chip gaps created by the primary grid. The lightcurves are written into separate lightcurve files, one for each field of the ZTF primary and secondary grids and spanning approximately $7^\circ \times 7^\circ$ each. In each one of these lightcurve files is a large ASCII table with a single row detailing each lightcurve metadata (including right ascension, declination, number of epochs, and so forth), followed by a series of rows containing the time (in heliocentric modified julian date, or hmjd), magnitude, magnitude error, linear color coefficient term from photometric calibration, and photometric quality flags for each measurement for each single epoch measurement. These lightcurve files are served for bulk download from the IPAC web server and total approximately 8.7 terabytes.

Continuing to discuss the structure of the lightcurve files requires a consistent vocabulary for describing this data product that we introduce here for clarity throughout this paper. A diagram of the terms introduced here is shown in Figure 5.1. An object refers to a collection of photometric measurements of a single star in a single filter, in either the primary or secondary grid (but not both). Objects are written in the lightcurve files grouped by their readout channel. A source refers to the total set of objects within a single lightcurve file that are all taken at the same sky location and are assumed to arise from the same astrophysical origin. There can be a maximum of three objects per source for the three ZTF filters. Crowding in dense galactic fields will often result in different astrophysical signals falling in the same PSF and therefore the same object. We address this issue during our microlensing fitting. Each of the objects within a single source are siblings to one another and have mutually exclusive filters. Siblings must share the same readout channel within a lightcurve file. A star refers to all sources of the same astrophysical origin found in all of the lightcurve

files. While many of these measurements will be not stars but in fact several stars blended together, galaxies, or other astrophysical phenomena, we adopt this nomenclature to serve our purpose of searching for microlensing events mainly in the galactic plane. Each star contains one or more sources, with each source containing one or more objects, with each object containing one lightcurve.




Figure 5.1: Diagram of how observational information is stored in the PUZZLE pipeline. An object is a set of observations at the same sky position taken on the same field, readout channel and filter. A source can contain multiple objects from the same field and readout channel that are each different filters. This distinction exists because observations in different filters are assigned different identification numbers in the ZTF public data releases. All sources must be in the same lightcurve file. Finally all sources are cross-referenced with each other by sky location and together form a single star.

There are several limitations inherit to the construction of these lightcurves that we have addressed in our work. First, the lightcurves for each filter are created from independent extraction of each filter’s single epoch flux measurement at a given location of the sky. Flux measurements in the g-band, r-band and i-band filters from a single astrophysical origin will have three separate object IDs at three different locations in the lightcurve file with no association information present relating the three objects to each other. In terms of our nomenclature, it is unknown whether a source will have one, two or three objects until the entire lightcurve file is searched row by row.

Second, objects found in different fields in the ZTF observing grid are also not associated with each other despite coming from the same astrophysical origin. The ZTF secondary grid intentionally overlaps with the primary grid to increase coverage in the gaps between readout channels and fields found in the primary grid. However this results in single filter measurements from the same star being split into two primary and secondary grid lightcurve files. We have found instances where telescope pointing variations has resulted in the same star having sources located in two different primary grid lightcurve files. Therefore it is unknown whether a single star will contain one or more sources without searching through both the primary and secondary grid lightcurve files, or searching adjacent lightcurve files for those stars that fall on the edges of each field.

Third, the format of the lightcurve files are raw ASCII tables. While this format is exchangeable across multiple platforms and therefore reasonably serves the purpose of a public data release, ASCII is suited to neither high-speed streaming of the data from disk

to memory nor searching for objects with particular properties without reading through an entire file. If we are to search for microlensing events throughout the entire ZTF catalog then we must either ingest these lightcurves files into a more traditional database structure or create a method for efficiently reading these lightcurves files that provides the benefits of a database structure.

5.3 Software

We have developed a set of software tools to address the outstanding issues with the DR lightcurve data product. These tools are a mixture of public open source codes and internal codes that enable efficient access to the data contained in the lightcurve files despite their ASCII format. The functionality and availability of these tools proved essential to scaling out microlensing search method to the entire ZTF lightcurve catalog. However these tools and approaches would prove generally useful to any science that seeks to execute a large search across the entire set of bulk lightcurves. We will therefore outline our approach in a level of technical detail that would aide another researcher attempting to similarly carry out such a large-scale search using the publicly available ZTF lightcurves.

5.3.1 zort: ZTF Object Reader Tool

An initial approach to reading the lightcurve files would reasonably be to ingest all of the data into a relational database where objects could be cross-referenced against each other to construct sources and sources cross-referenced against each other to construct stars. Flux measurements could be organized by object ID and an association table created to match these measurements to the metadata for each object. However this approach has two major drawbacks. For science cases where the lightcurve files are used not as a search catalog but instead an historical reference for a singular object of interest it is entirely overkill to ingest all data into an almost 8 terabyte database. Additionally, our approach aimed to leverage the computational resources available at the NERSC Supercomputer located at the Lawrence Berkeley National Laboratory. Here we can simultaneously search and analyze data from different locations in the sky by splitting our work into parallel processes. We would experience either race conditions or the need to exclusively lock the table for each process' write if all of our processes were attempting to read and write to a singular database table. Relational database access is a shared resources at NERSC, compounding this problem by putting a strict upper limit on the number of simultaneous connections allotted to single user when connecting to a database on the compute platform. These bottlenecks would negate the benefits of executing our search on a massive parallel supercomputer. We therefore sought a different method for accessing our data that would keep the data on disk and allow for so-called embarrassingly parallel data access.

The ZTF Object Reader Tool, or **zort**, is an open source Python package that serves as an access platform into the lightcurve files that avoids these bottlenecks ([Medford 2021b](#)).

It is a central organizing principle of `zort` that keeping track of the file position of an object makes it efficient to locate the metadata and lightcurve of that object. To enact this principle `zort` requires four additional utility products for each lightcurve file. These products are generated once by the user to initialize their copy of a data release by running the `zort-initialize` executable in serial or parallel using the python `mpi4py` package.

During initialization, each line of a lightcurve file (extension `.txt`) containing the metadata of an object is extracted and placed into an objects file (extension `.objects`), with the file position of the object within the lightcurve file appended as an additional piece of metadata. This file serves as an efficient way to blindly loop through all of the objects of a lightcurve file. For each object in the objects file, `zort` will jump to the object's saved file position and load all lightcurve data into memory when requested. Additionally, a hash table with the key-value pair of each object's object ID and file position is saved to disk as an objects map (extension `.objects_map`). This enables near-immediate access to any object's metadata and lightcurve simply by providing `zort` with an object ID.

Next, a k-d tree is constructed from the sky positions for all of the objects in each readout channel and filter within a lightcurve file and consolidated into a single file (extension `.radec_map`). This k-d tree enables the ability to quickly locate an object with only sky position. Lastly, a record of the initial and final file position of each filter and readout channel within a lightcurve file is saved (extension `.rcid_map`). Lightcurve files are organized by continuous regions of objects that share a common filter and readout channel, and by saving this readout channel map we are able to limit searches for objects to a limited set of readout channels if needed.

`zort` presents to the user a set of Python classes that enables additional useful features. Lightcurve files are opened with the `LightcurveFile` class that supports opening in a `with` context executing an iterator construct for efficiently looping over an objects file by only loading objects into memory as needed. This `with` context manager supports parallelized access with exclusive sets of objects sent to each process rank, as well as limiting loops to only specific readout channels using the readout channel map. Objects loaded with the `Object` class contain an instance of the `Lightcurve` class. This `Lightcurve` class applies quality cuts to individual epoch measurements as well as color correction coefficients if supplied with an object's PanSTARRS g-minus-r color. Objects contain a `plot_lightcurves` method for plotting lightcurve data for an object alongside the lightcurves of it's siblings. Each object has a `locate_siblings` method that uses the lightcurve file's k-d tree to locate coincident objects of a different filter contained within the same field. The package has a `Source` class for keeping together all sibling objects of the same astrophysical origin from a lightcurve file. Sources can be instantiated with either a list of object IDs or by using the utility products to locate all of the objects located at a common sky position. `zort` solves all problems related to organizing and searching for objects across the right ascension polar transition from 360° to 0° by projecting instrumental CCD and readout channel physical boundaries into spherical observation space and transforming coordinates.

`zort` has been adopted by the ZTF collaboration and TESS-ZTF project as a featured tool for extracting and parsing lightcurve files. `zort` is available for public download as a

GitHub repository (<https://github.com/michaelmedford/zort>) and as a pip-installable PyPi package (<https://pypi.org/project/zort/>).

5.3.2 PUZLE: Pipeline Utility for ZTF Lensing Events

Our search for microlensing events sought to combine all flux information from a single astrophysical origin by consolidating all sources into a single star. This required a massive computational effort as flux information was scattered across different lightcurve files as different objects with independent object IDs. First we identified all sources within each lightcurve file by finding all of the siblings for each object within that file. Then sources in different lightcurve files at spatial coincident parts of the sky were cross-referenced with each other to consolidate them into stars. To execute this method, as well as apply microlensing search filters and visually examine the results of this pipeline via a web interface, we constructed the python Pipeline Utility for ZTF Lensing Events or PUZLE. Similar to the motivations for constructing `zort`, this package needed to take advantage of the benefits of massive supercomputer parallelization without hitting the bottlenecks of reading and writing to a single database table.

The PUZLE pipeline began by dividing the sky into an grid of adaptively sized cells that contain approximately equal number of stars. Grid cells started at $\delta = -30^\circ, \alpha = 0^\circ$ and were drawn with a fixed height of $\Delta\delta = 1^\circ$ and a variable width of $0.125^\circ \leq \Delta\alpha \leq 2.0^\circ$ at increasing values of α . We attempted to contain no more than 50,000 objects within each cell using a density determined by dividing the number of objects present in the nearest located lightcurve file by the file’s total area. The cells were drawn up to $\alpha = 360^\circ$, incremented by $\Delta\delta = 1.0^\circ$, and repeated starting at $\alpha = 0^\circ$, until the entire sky had been filled with cells. The resulting grid of 38,819 cells can be seen in Figure 5.2. Each of these cells represents a mutually exclusive section of the sky that was split between parallel processes for the construction of sources and stars.

A PostgreSQL database was created to manage the execution of these parallel processes. This table was vastly smaller than a table that would contain all of the sources and stars within ZTF and was therefore not subject to the same limitations as previously stated for applying relational databases to our work. We created two `identify` tables each containing a row for each cell in our search grid, one for identifying sources and one for identifying stars. Each row represented an independent `identify_job` that was assigned to a compute core for processing. Each row contained the bounds of the cell, and two boolean columns for tracking whether a job had been started and/or finished by a compute core. Historical information about the date and unique compute process ID was also stored in the row for debugging purposes. To ensure that per-user database connection limits were not exceeded, we used a on-disk file lock. A parallel process had to be granted permission to this file lock before it attempted to connect to the database, thereby offloading the bottleneck from the database to the disk.

Here we describe how each `identify_job` script worked, whether it is identifying sources or stars. An `identify_job` script was submitted to the NERSC supercomputer requesting




Figure 5.2: Processing jobs for the PUZLE pipeline (top) and density of objects with a minimum number of observations (bottom) throughout the sky. Job cells were sized to enclose approximately the same number of sources resulting in a larger number of smaller sized jobs in the areas of the sky with more objects per square degree. All job cells had a fixed height of 1° and a variable width between 0.125° and 2.0° . The limitations of a Galactic survey from the Northern hemisphere are also seen in the lack of sources near the Galactic center.

multiple compute cores for a fixed duration of time. A script began by each compute core in the script fetching a job row from the appropriate `identify` table that was both un-started and un-finished and marked the row as started. The script then identified all of the sources (or stars) within the job bounds as described below. The resulting list of sources (or stars) was then written to disk for later processing. Lastly the job row in the `identify` table was marked as finished. Each script continued to fetch and run `identify_jobs` for all of its compute cores until just before the compute job was set to expire. At this time it interrupted whichever jobs were currently running and reset their job row to mark them as un-started. In this way the `identify` table was always ready for any new `identify_job` script to be simultaneously run with any other script and ensure that only un-started and un-finished jobs were run. The `identify_job` script for stars could only request jobs where the associated script for sources had finished.

The `identify_job` for finding sources began by finding the readout channels within all of the lightcurve files that intersected with the spatial bounds of it's job by using the projected readout channel coordinates calculated by the `zort` package. Special attention was paid here to lightcurve files that cross the right ascension polar boundary between 360° and 0° to ensure that they were correctly included in jobs near these boundaries. Objects within the readout channels that overlap with the job were then looped over using the `zort LightcurveFile` class and those objects outside of the bounds of the job were skipped. Objects with less than 20 epochs of good quality observations were also skipped over as this cut is an effective way to remove spurious lightcurves arising from erroneous or non-stationary seeds (see Figure 3 on the [DR5 page](#)). The remaining objects had their siblings located using the `locate_siblings` method and all of the siblings were grouped together into a single source. Once all of the lightcurve files had been searched, the remaining list of sources was cut down to only include unique sources with a unique set of object IDs. This prevented two duplicate sources from being counted, such as when a g-band object pairs with a r-band sibling, and that same r-band object pairs with that same g-band object as it's sibling. Each source was assigned a unique source ID. The object IDs and file positions of all objects within each source were then written to disk as a source file named for that `identify_job`. In addition to this source file, a hash table similar to the `zort objects_map` file was created. This source map was a hash table with key-value pairs of the source ID and file position within the source file where that source was located.

The `identify_job` for finding stars began by loading all of the sources for that job directly from the on-disk file written by the source script. The sky coordinate of each source was loaded into a k-d tree. Each source's neighbors located within 2 arcseconds were found by searching through this tree. Sources were grouped together to form a star. A star's location was calculated as the average sky coordinate of it's sources and the star defined as the unique combination of source IDs which it contains. Each star was assigned a unique ID. The star id, sky location and list of source IDs of each star were written to disk in a star file. A star map similar to the source map described above was also written to disk.

5.3.3 Upgrading to a New Public Data Release

Our approach of using the file position to keep track of objects and sources was able to transition between different versions of the public data releases when available. We used an older DR as a seed release upon which we initially searched for sources and stars. For this microlensing search and by way of example we used DR3 as a seed release and updated our source and star files to DR5. We began by creating the necessary `zort` utility products for DR5 lightcurve files. We sought to append additional observations to our seed list of sources and stars. However we did not look to find any additional objects that did not exist in DR3. The convention for different DRs is to keep the same object IDs for reference sources found at the same sky coordinates and to add new IDs for additional objects found. Therefore when updating a row in the source file from DR3 to DR5 we only needed to use the source's object IDs and the DR5 object map to find the file location of the object in the DR5 lightcurve file.

We found that in this process less than 0.02% of the object IDs in a source file could not be located in the DR5 object map, resulting in a 99.98% conversion rate that was acceptable for our purposes. A new DR5 source map was then generated for this DR5 source file. Star files did not need to be updated and could be simply copied from an older to a newer data release. It should be noted that the names of lightcurve files slightly change between data releases, as the edge sky coordinates are printed into the file names and are a function of the outermost object found in each data release for that field. Any record keeping that involved the names of these lightcurve files needed to be adjusted accordingly.

5.4 Detection Pipeline

With our pipeline, we sought to measure the t_E distribution for long-duration microlensing events ($t_E \geq 30$ days) in order to search for a statistical excesses due to black hole microlensing. Second, we sought to generate a list of black hole candidates that could be followed up astrometrically by our group in future studies. We therefore made design choices optimized for increasing the probability of detecting these types of events, even at the expense of short-duration events. We designed our pipeline to remove the many false positives that a large survey like ZTF can generate, even at the expense of false negatives. Our detection pipeline had to be capable of searching through an extremely large number of objects, sources and stars and therefore had to be both efficient and computationally inexpensive to operate. These two priorities informed each of our decisions when constructing our microlensing detection pipeline.

5.4.1 Cutting on the von Neumann ratio, Star Catalogs and a Four-Parameter Microlensing Model

We created a `process` table containing a row for each cell in our search grid, just as we did for the sources and stars. However each row also had columns for detection statistics that kept track of throughout the pipeline’s execution. Metadata recording choices our algorithms automatically determined was also saved for later debugging and analysis. We reiterate that the cost of reading and writing to such a table from many parallel processes is not a constraining factor when each job needs to only read and write to this table once.

The pipeline continued with a `process_job` script selecting a job row from the `process` table that was un-started and un-finished. The job read in the star file for it’s search cell and the associated source maps for each lightcurve file from which a source could originate within it’s cell. These source maps were used to find the source file locations for each source ID of a given star. Each source row in the source file had the object IDs of the source’s objects and `zort` could use these object IDs to load all object and lightcurve data into memory. Throughout the pipeline we maintained a list of stars and matched that list with the list of sources associated with each of those stars. Our pipeline performed calculations and made cuts on objects, but our final visual inspections was done on the stars to which those objects

belong. We therefore needed to keep track of all associations between stars, sources and objects throughout the pipeline. All cuts described below were performed within the stars, sources and objects of each `process_job`.

The first cut we implemented was to remove any objects with less than 50 unique nights of observation. There were fields in our sample that were observed by the High-Cadence Plane Survey that contain many observation epochs but all within a few nights (Bellm et al. 2019b). We limited our search to those events with many nights of observation due to the long duration of black hole microlensing events as described in Section 5.1. While we had already performed a cut on objects with less than 20 good quality observations, performing this cut removed those objects which only passed our initial cut due to a small number of nights being sampled multiple times. We were left with 1,011,267,730 objects and 563,588,562 stars in our level 1 catalog.

Price-Whelan et al. (2014) developed a detection method for finding microlensing events in large, non-uniformly sampled time domain surveys using Palomar Transient Factory (PTF) data. Their key insight was to use η , or the von Neumann ratio (also referred to as the Durbin-Watson statistic; von Neumann et al. 1941; Durbin & Watson 1971), to identify microlensing events. This statistics is an inexpensive alternative to the costly $\Delta\chi^2$ that measures the difference in χ^2 of fitting data to a flat model and to a microlensing model. Their pipeline was also biased toward removing false positives at the expense of false negatives by culling their data on statistical false positive rate thresholds for η . ZTF captures nearly an order of magnitude more sky coverage with each exposure than PTF with different systematics, motivating a different implementation of this statistic.

We calculated the von Neumann ratio η on all objects in the level 1 catalog:

$$\eta = \frac{\delta^2}{\sigma^2} = \frac{\sum_i^{N-1} (x_{i+1} - x_i)^2 / (N - 1)}{\sigma^2}. \quad (5.1)$$

η is the ratio of the average mean square difference between a data point x and its successor $x+1$ to the variance of that dataset σ . Highly correlated data has a small difference between successive points relative to the global variance and have a correspondingly small η . Gaussian noise has an average $\eta \approx 2$ with a smaller variance in the measurement for datasets with more points. Several example lightcurves and their associated η values are demonstrated in Figure 5.3.

We calculated our η not on an object’s epoch magnitudes but instead on an object’s nightly magnitude averages. Fields observed by the High-Cadence Plane Survey were observed every 30 seconds and had additional correlated signal due to being sampled on such short time scales relative to the dynamic timescale of varying stellar brightness. This biased the objects in a job cell toward lower η . Calculating η on nightly magnitude averages removed this bias and created an η distribution closer to that expected by Gaussian noise. This reduced the undue prominence of those objects observed by the High-Cadence Plane Survey when searching for long-duration microlensing events. We performed all η calculations on the dates and brightness of nightly averages.

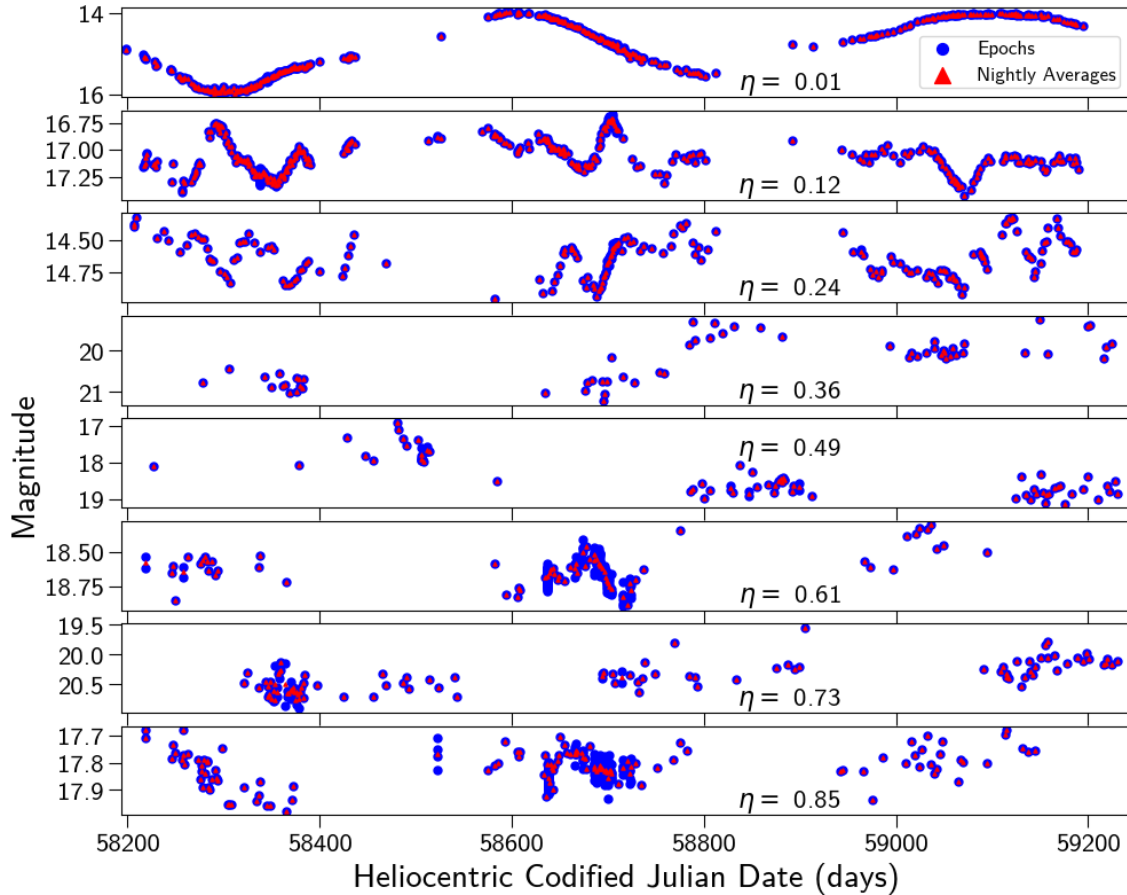


Figure 5.3: Example ZTF objects and their η values, with individual epochs (blue) and nightly averages (red). DR5 lightcurves contained various cadences and gaps in the data depending on their visibility throughout the year and which ZTF surveys were executed at their location. This resulted in a heterogeneous dataset that required a flexible approach. The η statistic was able to capture inter-epoch correlation despite these different observing conditions, with smaller η signaling more correlated variability in the lightcurve.

A cut on η required a false positive threshold that separated those events with significant amounts of correlated signals and those without. We made two alterations to our calculation of this threshold than that found in [Price-Whelan et al. \(2014\)](#). First we chose to calculate a threshold not from determining the 1% false positive recovery rate for scrambled lightcurves, but instead by finding the 1st percentile on the distribution of η . Initial attempts to set the threshold from scrambled lightcurves resulted in more than 1% of the objects passing our cut due to correlated noise not explained by the global variance. Only passing the 1% of objects with the lowest η guaranteed that this stage of the pipeline would remove 99% of events, significantly cutting down on the number of objects passing this stage of the pipeline.

Second we chose to bin our data by number of observation nights and calculated a separate

threshold for each of these bins. The variance of η is correlated with the number of data points in its measurement and therefore so too was the 1st percentile of η correlated with the number of observations. Objects in our level 1 catalog span from 50 nights of observation to nearly 750 nights due largely to the presence of both primary and secondary grid lightcurves, as well as the mixture of ZTF public and partnership surveys, falling into the same job cell. A single threshold calculated from all lightcurves would be biased toward passing short duration lightcurves and removing long duration lightcurves. Our binning significantly dampened the effect of this bias by only comparing lightcurves with similar numbers of observation nights to each other. In order to efficiently divide the lightcurves into these bins, we determined the cumulative distribution function for number of observations nights and found the values of η for which the cumulative distribution function equaled 0.33 and 0.66. These bin edges were unique for each job and were recorded in the `process` table. Each lightcurve was then compared to these bin edges and assigned to the appropriate bin. For each bin we calculated the 1st percentile of η and removed all lightcurves with η greater than this threshold. We also calculated the 90th percentile of η in each bin and saved it for a later stage in our pipeline.

Our next cut was a star-galaxy cut on those sources which we were confident are not astrophysical stars (not to be confused with our nomenclature for the word “star”). We used the Probabilistic Classifications of Unresolved Point Sources in PanSTARRS1 (PS1-PSC) that classified ~ 1.5 billion PanSTARRS1 sources as either extended sources or point sources using a machine learning model (Tachibana & Miller 2018). Each remaining object in a job was queried against this catalog to find a corresponding PS1-PSC score at that location in the sky. Lightcurves were retained that had a PS1-PSC score greater than or equal to 0.645, or did not have a corresponding score. This threshold is the value at which PS1-PSC labeled sources with a `rKronMag` < 21 (which captures nearly all ZTF sources) as astrophysical stars with a 96.4% true positive rate and only costs a false positive rate of 1.0%. This cut retained 96.4% of the astrophysical stars in our sample while only permitting 1.0% of galaxies to pass. To avoid the same database bottlenecks previously described, we downloaded the entire PS1-PSC catalog onto disk (private communication Adam Miller) broken into individual files for separate sections of the sky. We generated k-d trees for each of these files and all spatially coincident PS1-PSC catalog files were loaded into memory at the beginning of each `process_job`, enabling a fast PS1-PSC score look up for each object at run-time.

Lastly we fit a four-dimensional microlensing model to the daily average magnitudes of each remaining object. Microlensing models are multi-dimensional and non-linear which results in costly fitting that would be prohibitive to a search of this scale. Kim et al. (2018b) outlined an analytical representation of microlensing events that circumvented this issue by only attempting to fit microlensing events in the high magnification ($u_0 \lesssim 0.5$) and low magnification ($u_0=1$) limits. The microlensing model representation they deduced for these

Table 5.2: Boundaries for four-parameter microlensing model.

	t_0 (hmjd)	t_{eff} (days)	f_0 (flux)	f_1 (flux)
Low Bound	$\min(t) - 50$	0.01	$-\infty$	0
High Bound	$\max(t) + 50$	5000	∞	∞

limits was

$$\begin{aligned}
 F(t) &= f_1 A_j [Q(t; t_0, t_{\text{eff}})] + f_0; \\
 Q(t; t_0, t_{\text{eff}}) &\equiv 1 + \left(\frac{t - t_0}{t_{\text{eff}}} \right)^2; \\
 (j &= 1, 2), \\
 A_{j=1}(Q) &= Q^{-1/2}; \\
 A_{j=2}(Q) &= [1 - (Q/2 + 1)^{-2}]^{-1/2},
 \end{aligned} \tag{5.2}$$

with $j = 1, 2$ corresponding to $u_0 \lesssim 0.5$ and $u_0=1$ limits respectively. In this construction f_0 and f_1 no longer have a physical interpretation but are instead simply parameters of the fit. t_{eff} is a non-physical substitution for t_E . At this stage in our pipeline we had sufficiently small numbers of events as to permit the simultaneous four dimensional fit of f_0 , f_1 , t_0 , and t_{eff} for both the low and high magnification limits. Unlike the grid search that [Kim et al. \(2018b\)](#) performed for a solution, our two fits were performed on each object with bounds (Table 5.2) and least squares minimization with the Trust Region Reflective algorithm ([Branch et al. 1999](#)). The average time to fit each object to both the low and high magnification solution was 60 ± 5 milliseconds. For each of these two fits the $\Delta\chi^2$ between the microlensing model and a flat model was calculated and the solution with the largest $\Delta\chi^2$ was kept as the best fit. This model was subtracted from the lightcurve and η_{residual} was calculated on these residuals. If η_{res} was low then the microlensing model had failed to capture the variable signal that allowed the lightcurve to pass the first η cut and additional non-microlensing variability remained. We only retained those lightcurves with little remaining variability in the residuals by removing all objects with η_{residual} less than the 90th percentile threshold calculated earlier in our first η cut.

At this stage we combined our different epoch bins to obtain a single list of objects that had passed all of our cuts, as well as the record of the sources and stars to which these objects belong. This smaller catalog could be saved into a more conventional relational database as it no longer required massive parallel compute power to perform further cutting and fitting. For each star, hereafter referred to as a candidate, all of the candidate information, including the list of all its source IDs, was uploaded to a `candidates` table. If there were multiple objects belonging to the candidate that passed all cuts, the pipeline data of the object with the most number of nights was saved in the candidate database row. In this way each candidate contains the information of the star from which it was derived but has a single object upon which further cuts can be performed. The job row in the `process` table was marked as finished and updated with metadata relating to the execution of the job.

These cuts reduced the number of objects from 563,588,562 to 7,457,583 level 2 candidates, as outlined in Table 5.3.

5.4.2 Simulated Microlensing Events

Reducing the number of candidates further required knowing how potential microlensing events within the sample would be affected by particular cuts. We therefore generated artificial microlensing events and injected them into ZTF data. Following the method outlined in Medford et al. (2020a), we ran 35 PopSyCLE simulations (Lam et al. 2020) throughout the Galactic plane and imposed observational cuts mimicking the properties of the ZTF instrument. The Einstein crossing times and Einstein parallaxes of artificial events were drawn from distributions that were fit for each PopSyCLE simulation. The impact parameters were randomly drawn from a uniform distributions between $[-2, 2]$. The source flux fraction, or the ratio of the flux originating from the un-lensed source to the total flux observed from the source, lens and neighbors, was randomly drawn from a uniform distributions between $[0, 1]$. These parameters were then run through a point-source point-lens microlensing model with annual parallax and only sets of parameters with an analytically calculated maximum source amplification greater than 0.1 magnitudes were kept.

We sampled real lightcurves from ZTF objects with at least 50 nights of observation located within the footprint of each corresponding PopSyCLE simulation. This real lightcurve equals, in our model, the total flux from the lens (F_L), neighbors (F_N), and source (F_S) summing to a total F_{LNS} outside of the microlensing event. The light from the lens and neighbor can be written, using the definition of the source flux fraction $b_{\text{sff}} = F_S/F_{LNS}$, as

$$\begin{aligned} F_{LN} &= F_{LNS} - F_S , \\ F_{LN} &= F_{LNS} - b_{\text{sff}} \cdot F_{LNS} , \\ F_{LN} &= (1 - b_{\text{sff}}) \cdot F_{LNS} . \end{aligned} \tag{5.3}$$

An artificial microlensing lightcurve (F_{micro}) has amplification applied to only the flux of the source,

$$F_{\text{micro}} = A \cdot F_S + F_{LN} . \tag{5.4}$$

Equations 5.3 and 5.4 can be combined to obtain a formula for the microlensing lightcurve using only the original ZTF data and the amplification of the microlensing model,

$$F_{\text{micro}} = A \cdot b_{\text{sff}} \cdot F_{LNS} + (1 - b_{\text{sff}}) \cdot F_{LNS} . \tag{5.5}$$

A value of t_0 was randomly selected in the range between the first and last epoch of the lightcurve. Events were then thrown out if they didn't have (1) at least three observations of total magnification greater than 0.1 magnitudes, (2) at least three nightly averaged magnitudes observed in increasing brightness in a row, and (3) have at least three of those nights be three-sigma brighter than the median brightness of the entire lightcurve. This selection

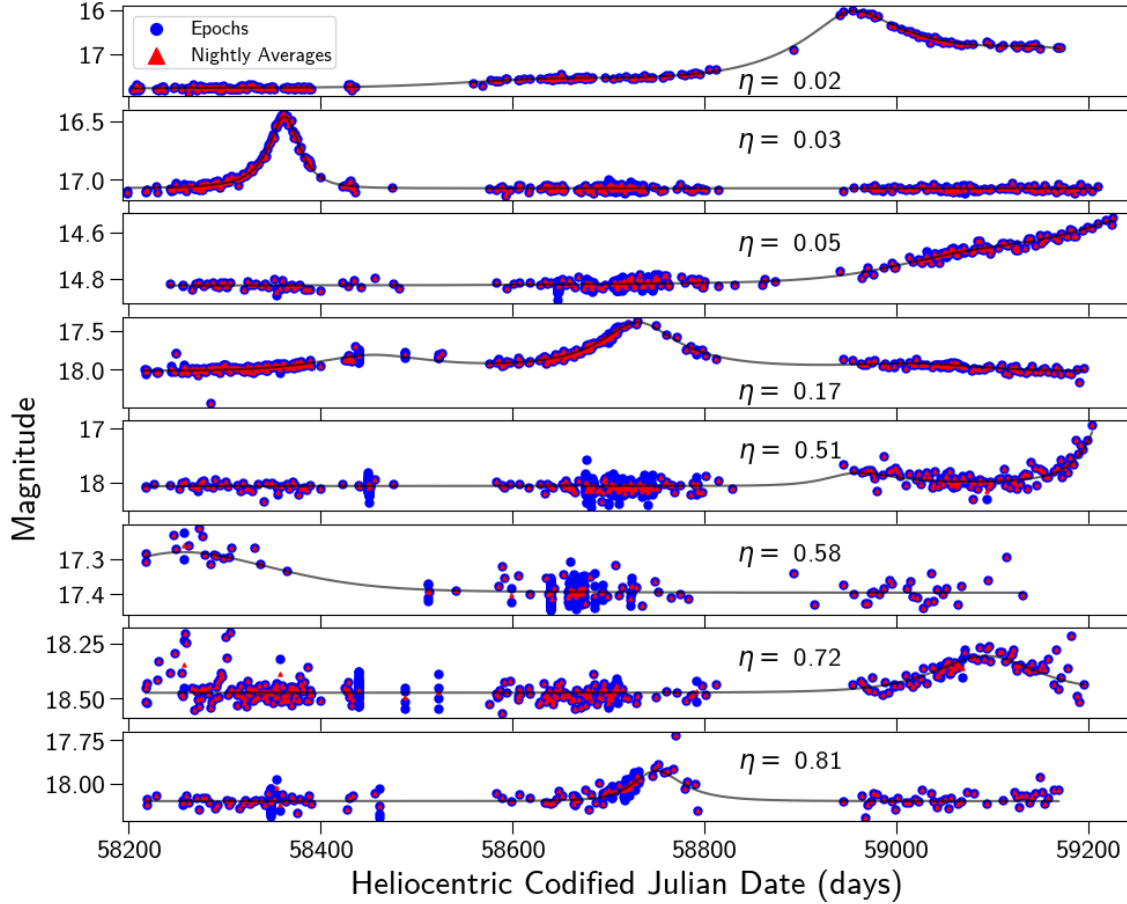


Figure 5.4: Example simulated microlensing lightcurves and their η values, with individual epochs (blue) and nightly averages (red). Each model (black line) was generated from distributions set by PopSyCLE simulations throughout the Galactic plane. The models were then injected directly into DR5 lightcurves, enabling calculations and cuts on these lightcurves to represent the diversity of cadence and coverage seen throughout the DR5 dataset.

process yielded 63,602 simulated events with varying signal-to-noise as shown in Figure 5.4.

η and η_{residual} were calculated on all simulated microlensing events. Most of the simulated microlensing events had low values of η , due to the correlation between subsequent data points as compared to the sample variance, and high values of η_{residual} , due to the lack of correlation after subtracting a successful microlensing model. Figure 5.5 shows the location of the simulated microlensing events in the $\eta - \eta_{\text{residual}}$ plane.

For each level 2 candidate we selected the object with the most nights of observation from among the candidate’s objects. The η , η_{residual} values of these 7,457,583 best lightcurves is plotted in Figure 5.5 as well. There existed a clear distinction between the location of most of our candidates in this plane as compared to the simulated events. This distinction was

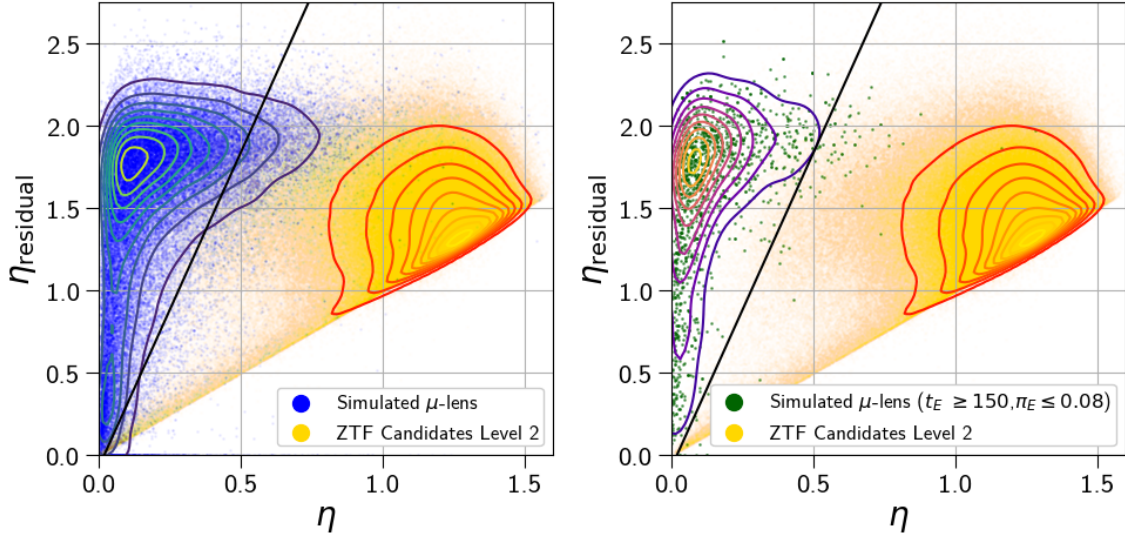


Figure 5.5: Microlensing events were clearly delineated in the level 2 catalog by calculating η on lightcurves and the residuals of those lightcurves after fitting them a four-parameter microlensing model. η was calculated on the nightly averages of ZTF candidates (yellow), simulated microlensing lightcurves (blue) and simulated microlensing lightcurves in the black-hole search space (green). This measurement of η was made on both the entire lightcurve (x-axis) and on a lightcurve with the four-parameter model subtracted (y-axis). Microlensing events tended toward smaller η values due to their correlated lightcurves and larger η_{residual} when microlensing is the only source of variability. This enabled a cut (black line) in this η - η_{residual} space that retained 83.6% of simulated events and 95% of simulated black events, while removing 98.8% of the ZTF candidates from our sample.

even stronger when we limited our simulated sample to events with true values of $t_E > 150$ and $\pi_E < 0.08$. These cuts have previously been used as selection criteria for identifying microlensing events as black hole candidates. We ran a grid of proposed cuts with the criteria $\eta_{\text{residual}} \geq m \cdot \eta + b$ and found that $m = 3.62, b = 0.01$ retained 83.6% of all microlensing lightcurves and 95% of the black hole microlensing lightcurves, while removing 98.8% of our level 2 candidates. This left us with 92,201 candidates in our level 3 catalog.

5.4.3 Cutting on a Seven-Parameter Microlensing Model

The nightly averaged magnitudes of the best lightcurve of each candidate were fit with a seven-parameter microlensing model for further analysis. The `scipy.optimize.minimize` routine (Virtanen et al. 2020) using Powell’s method (Powell 1964) fit each lightcurve to a point-source point-lens microlensing model including the effects of annual parallax. The average time to fit each object with this model was 1.2 ± 0.3 seconds. 603 candidates failed to be fit with this model and were cut. All simulated microlensing events were also fit with

this model for comparison. The results of these fits for both populations are shown in Figure 5.6.

Several cuts were then made on candidates using these fits. Our first cut was to remove candidates with excessively large $\chi_{\text{reduced,model}}^2$ values because they could not be well fit by our microlensing model. The threshold of $\chi_{\text{reduced,model}}^2 = 4.805$ was set by calculating the 95th percentile of the simulated data and removed 43,525 candidates. We next required that the candidate’s selected object had a range of epochs observed equal to at least $2t_E$ nights outside of the event ($t_0 \pm 2t_E$), removing another 34,301 candidates. $\chi_{\text{reduced,flat}}^2$ was determined by calculating the χ_{reduced}^2 comparing the brightness to the average magnitude in the region outside of the fit event for each lightcurve. This statistic and its 95th percentile value were also calculated on the simulated events and candidates above this percentile were removed. A cut was performed on candidates above the 95th percentile of π_E to eliminate candidates where excessive parallax was producing variability in the data. Lastly a cut was applied to t_0 , limiting candidates to only those that peaked at least 1 t_E into the survey data, removing candidates that peaked too early to be well constrained. The combination of these cuts resulted in a list of 8,619 candidates that were either completed or ongoing events in our level 4 catalog.

5.4.4 Cutting on a Bayesian Microlensing Model

All of these candidates are interesting objects worthy of further investigation. However we divided our sample into candidates with peaks during the survey ($t_0 + t_E < 59243$) and those candidates with peaks after the survey was completed ($t_0 + t_E \geq 59243$). These sub-samples of 6,929 and 1,690 candidates respectively serve two different purposes. Calculating the statistical properties of microlensing candidates within ZTF data requires a clean sample of completed events. We therefore sought to further remove contaminants from the 6,929 candidates that had peaked within the survey. The 1,690 candidates that had yet to peak need to be further monitored while they are rising in brightness until they reach their peak amplification. At this point they could be fit with a microlensing model to determine whether they are good events and perhaps even black hole candidates worthy of astrometric followup.

Each of the 6,929 completed candidates was successfully fit with a nested Bayesian sampler to a point-source point-lens microlensing model with annual parallax locked to the sky location of the candidate. We perform this more sophisticated fit to obtain error bars on our microlensing parameter measurements. Gaussian processes were also included in the fitter to model correlated instrument noise using the procedure outlined in Golovich et al. (2020). This fitter fit the three most observed objects within the candidate with at least 20 nights of data to the same model simultaneously. This model had five parameters shared by all lightcurves ($t_0, t_E, u_0, \pi_{E,E}, \pi_{E,N}$), two photometric parameters fit for each lightcurve ($m_{\text{base}}, b_{\text{sff}}$), and four Gaussian process parameters for each lightcurve ($\sigma, \rho, \omega_0, S_0$). Therefore each candidate was fit with either 11, 15 or 19 parameters if it contained 1, 2 or 3 objects with at least 20 nights of data. The median time to complete each of these fits (parallelized across 8 compute cores) was 82,575 and 3,515 seconds for 1, 2 or 3 lightcurves respectively.

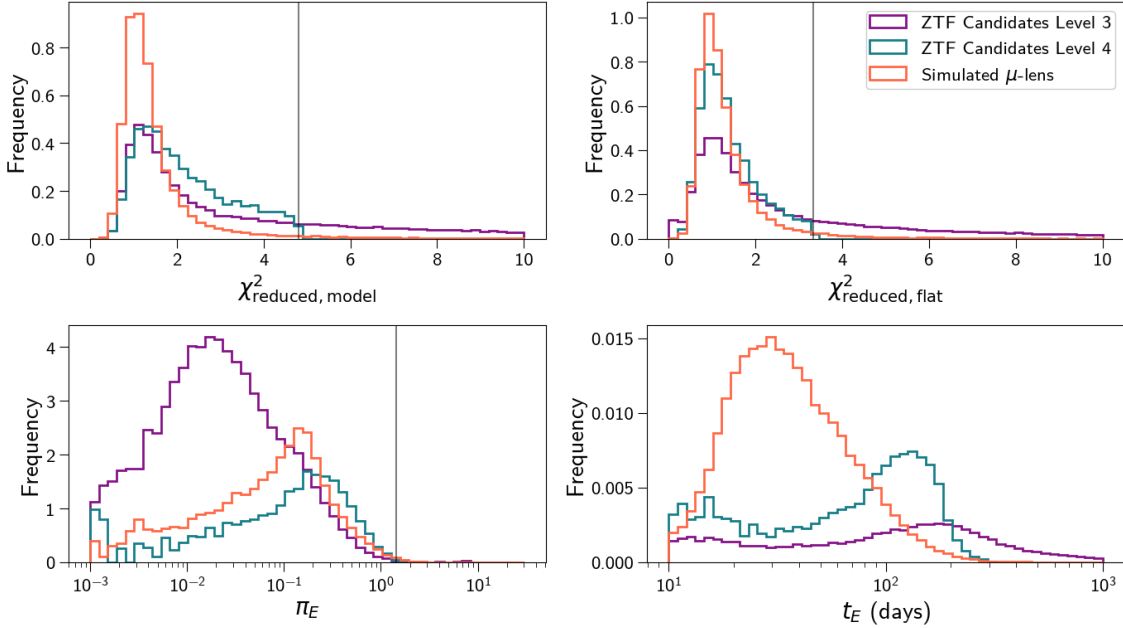


Figure 5.6: Fitting the seven-parameter model to both level 3 candidates (purple) and simulated microlensing lightcurves (orange) revealed cuts that removed many objects, leaving behind our level 4 catalog (blue). The $\chi^2_{\text{reduced,model}}$ of the entire lightcurve that was fit to a microlensing model (top-left) had a long tail in the level 3 candidates that was not present in the simulated microlensing events, motivating a cut on the 95th percentile (black line). Similarly the long tail for level 3 candidates of the $\chi^2_{\text{reduced,flat}}$ of the lightcurve outside of $t_0 \pm 2t_E$ fit to a flat model (top-right) was not seen in the microlensing samples and was removed with a cut on the 95th percentile. A small number of events with significantly larger Einstein parallaxes (bottom-left) were removed with a similar 95th percentile cut. These cuts, and others, removed events with Einstein crossing times (bottom-right) longer than are detectable by our survey.

Data quality cuts were applied to the resulting fits, as outlined in Table 5.3. The fractional error on t_E was required to be below 20% to ensure that a final distribution of Einstein crossing times would be well defined. Cuts on $u_0 \leq 1.0$ and $b_{\text{sff}} \leq 1.2$ are commonly accepted limits for the selection of high quality events. $\chi^2_{\text{reduced}} \leq 3.0$ for the data evaluated against the model fit was set to similarly remove lower quality events. The final three cuts were re-implementations of earlier cuts but with this more rigorous fit. The candidate was required to have $4t_E$ of observations outside of observations that were outside of significant magnification ($t_0 \pm 2t_E$). Lastly the candidate was required to have observations on both the rise ($t_0 - t_E \geq 58194$) and and fall ($t_0 + t_E \leq 59243$). All of these cuts reduced the sample down down to 669 candidates in our level 5 catalog. We additionally fit the 1,690 candidate to the same nested Bayesian sampler but with only the one single lightcurve containing the most nights of observations. This gave us the error on the t_0 and t_E measurements that we

could be used to determine which of these candidates should be astrometrically observed.

5.4.5 Manual Lightcurve Inspection

There were 669 candidates remaining in the level 5 catalog. There were several failure modes of our pipeline that could only be addressed by manually labeling each of these candidates. To facilitate this, we constructed a website that provided access to the candidate information contained within the PUZLE database tables and lightcurve plots generated by `zort`. This website was a docker container running a flask application that was served on NERSC’s Spin platform. The website connected inspectors to the NERSC databases and on-site storage. Each of the candidate labels was derived after inspecting the data by eye and grouping the candidates into common categories. The description and final number of candidates with each label after manual inspection are outlined in Table 5.4. Examples of each candidates with each label are shown in Figure 5.7.

The inspector was shown a display page with the four labels displayed above the `zort` lightcurves alongside information from the database. The model derived from the Bayesian fit was plotted onto the 1 to 3 objects that contributed to fitting the model and was absent from those objects within the candidate that did not. The user then selected which label best matched the candidate. The user either selected a scoring mode, where the page was automatically advanced to another unlabelled candidate, or a view mode, where the page remained on the candidate after selecting a label. An example page from the labelling process is shown in Figure 5.8.

5.5 Results

Our final catalog of level 6 candidates contains 128 clear microlensing events. The sky location of these events are shown in Figure 5.9. 93 (73%) of the events are within the Galactic plane ($|b| \leq 15^\circ$) and 35 events (27%) are outside of it. The distribution of events roughly follows the density of observable objects within DR5, also included in the figure. This indicates that our pipeline is finding microlensing events proportional to the stellar density in each part of the sky. A selected number of these events can be seen, divided by sky location, in Figure 5.10. There are a large number of events located in parts of the sky far outside of the Galactic plane. Microlensing from two spatially coincident stars is far less likely in these regions. We further discuss the possible origin of these events in Section 5.6. All 128 microlensing events have been posted online for public use (Medford 2021a), following the schema outlined in Table 5.5.

The distribution of Einstein crossing times in our level 6 events are shown in Figure 5.11. Short duration events as short as three days are recorded, increasing in number up to a peak at $t_E \approx 50$ days. At longer durations there is a flattened top to our distribution, followed by a sharp decline with no events found at $t_E \geq 140$ days. Several cuts within our pipeline require significant duration of observations outside of the microlensing event. These

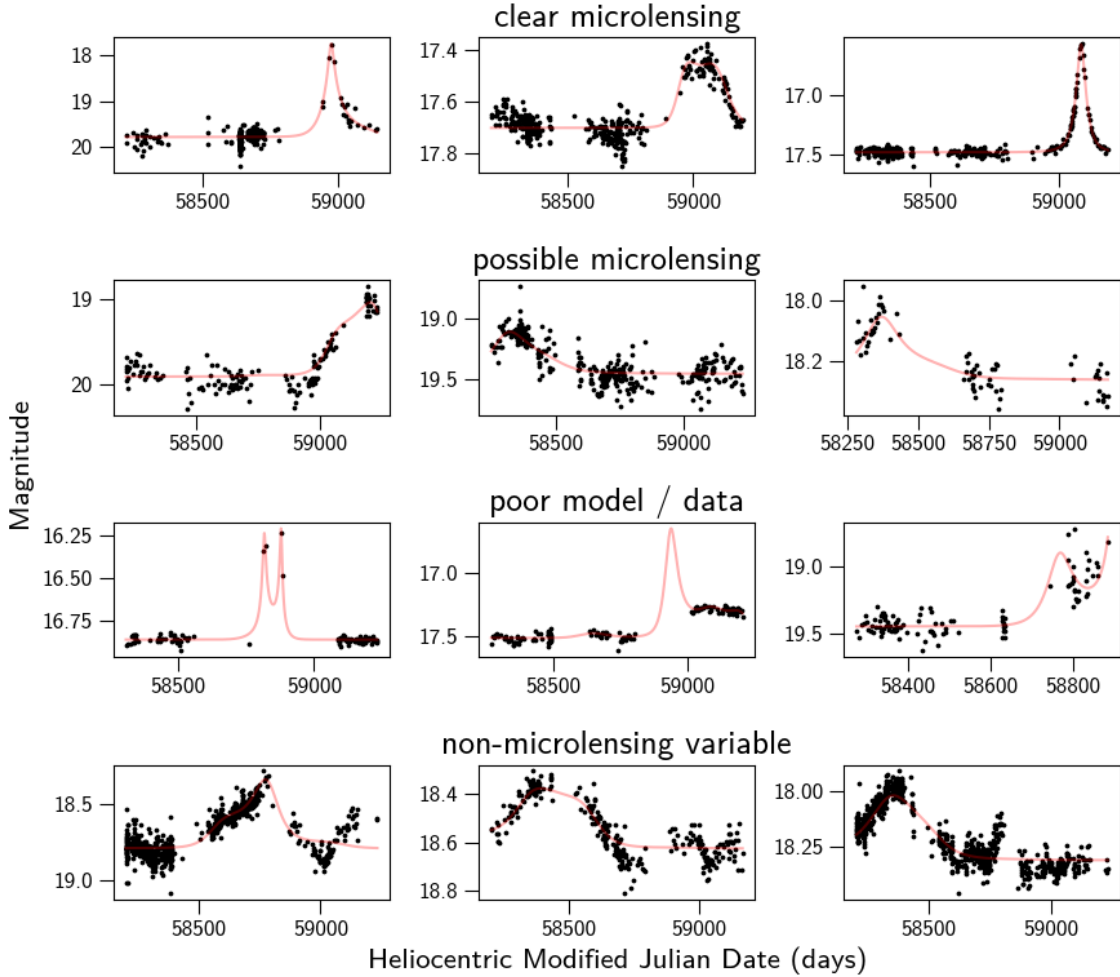


Figure 5.7: Lightcurves exemplifying the five labels manually assigned to level 5 candidates. Only the clear microlensing events were selected for our final level 6 catalog. The limitations of our method can be seen in the persistence of non-microlensing variables that had additional variation not explained by microlensing, as well as lightcurves where the model was not supported by the data labelled as poor model / data. The possible microlensing events had insufficient data to be confidently named microlensing. Those appearing to rise in brightness could be confirmed with further observations.

constraints, combined with the limited duration of DR5, place a strong upper limit on the timescale of events we can detect. We compare this distribution previous simulations and other surveys in Section 5.6.

The luminosity function of our level 6 events is shown in Figure 5.12. The magnitude shown is the median baseline magnitude, as determined by our Bayesian filter, for all lightcurves of the same filter belonging to an event. 127 of the events contain at least one lightcurve that was fit by the Bayesian fitter in the r-band. Only 110 of the events have

Level 5 Candidate 226146_9143

Glion Glat 89.38878 , 12.39568

RA Dec 301.31815 , 55.32261

Category clear_microlensing

[Level 3 Candidate Page](#)

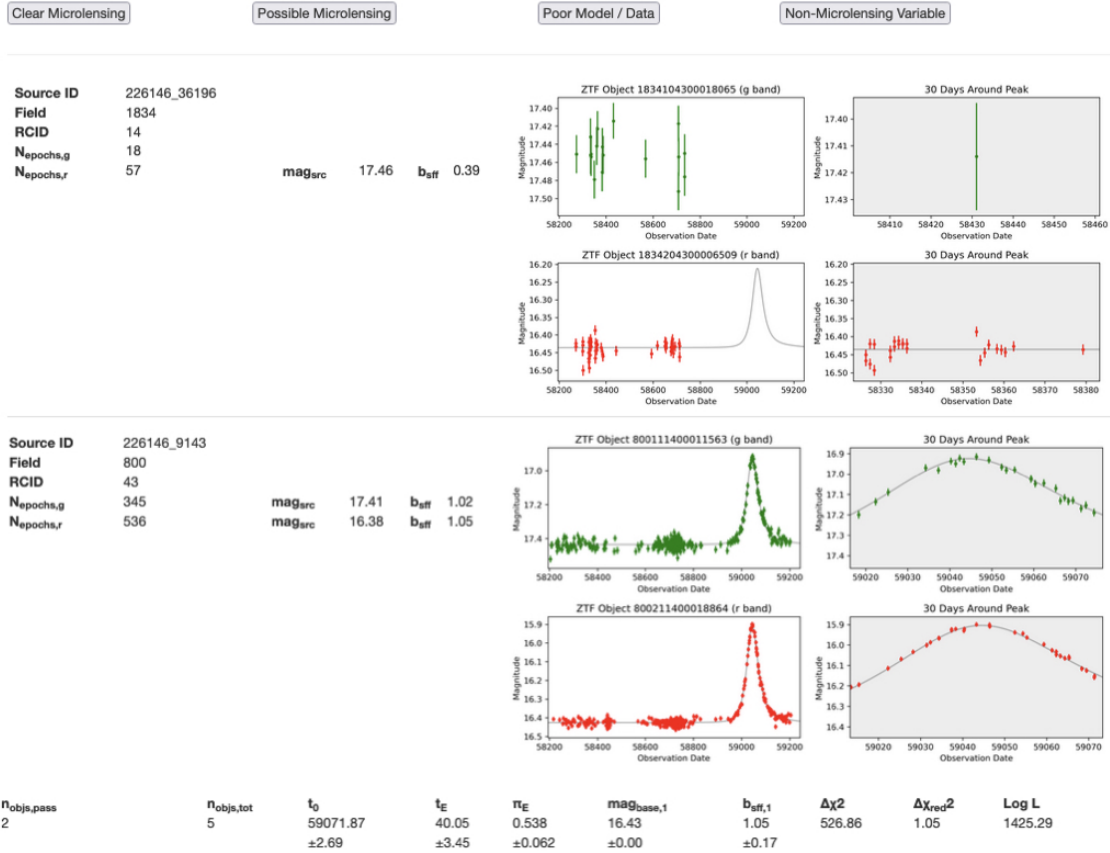


Figure 5.8: Level 5 candidates were manually screened by a human expert on the PUZLE website. Inspectors were presented with all objects in the candidate star grouped by source. Model curves were plotted on those lightcurves that were included in the Bayesian fit. The a maximum a posteriori probability and 1σ error bars of Bayesian fit parameters were also shown. The inspector then assigned one of four labels to the candidate.

fit data in g-band and 14 events possess fit i-band. All three filters peak in number at 19th magnitude, followed by a sharp decline. Our pipeline appears to be removing events fainter than 19th magnitude. The baseline magnitudes for g-band events also decline but at a slower rate. We compare the r-band luminosity function to simulation in Section 5.6.

Medford et al. (2020a) predicted that events observed throughout the outer Galactic plane would have larger source flux fractions (b_{src}) than events observed in the Galactic bulge. Fig-

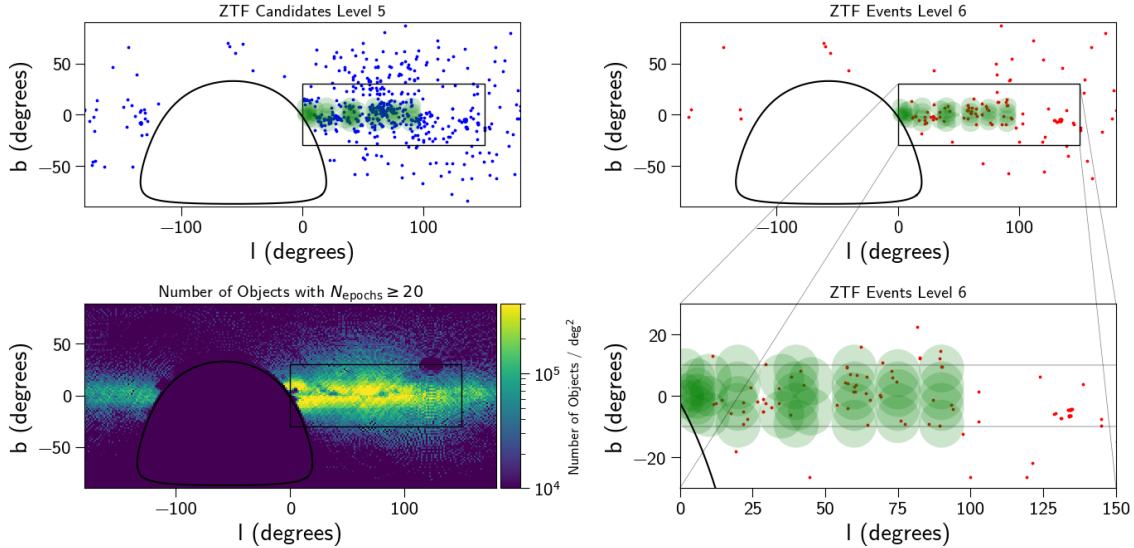


Figure 5.9: Identifying our final level 6 candidates by manual inspection produced a list of clear microlensing events scattered throughout the night sky. The 669 level 5 candidates (blue) were positioned throughout the sky (top-left), while the 128 level 6 events (red, top-right) are primarily located within the Galactic plane (bottom-right). 56 of the 128 level 6 events (44%) appear within with the area of the PopSyCLE simulations we performed (footprints in green). However there are a significant number of events at larger Galactic latitudes. This can be explained by the relatively larger number of objects with $N_{\text{epochs}} \geq 20$ observed at these locations in DR5 (bottom-left), contamination by long-duration variables, or perhaps the presence of MACHOs in the stellar halo.

Figure 5.13 shows that our level 6 candidates have brighter source flux fractions in line with this prediction. Very few events are seen with $b_{\text{sff}} \leq 0.2$ as the relatively small stellar densities of the outer Galactic bulge and the stellar halo prevent many neighboring stars from appearing in the instrument’s observational aperture. This is in contrast to the findings of most Galactic bulge surveys that find a bi-modal distribution with a peak around $b_{\text{sff}} \lesssim 0.2$. Both g-band and r-band objects have an increasing number of events at larger source flux fractions and peak at approximately $b_{\text{sff}} \approx 1$. However both bands and particularly the g-band shows an excess of events at $b_{\text{sff}} \geq 1$. This indicates that there is an overestimate of the background noise present in the photometric measurements. Reprocessing the ZTF observations at the location of these events with calibration parameters tuned to the particulars of each event’s fields could improve the estimates of these backgrounds and reduce the number of events with excessive source flux fractions.

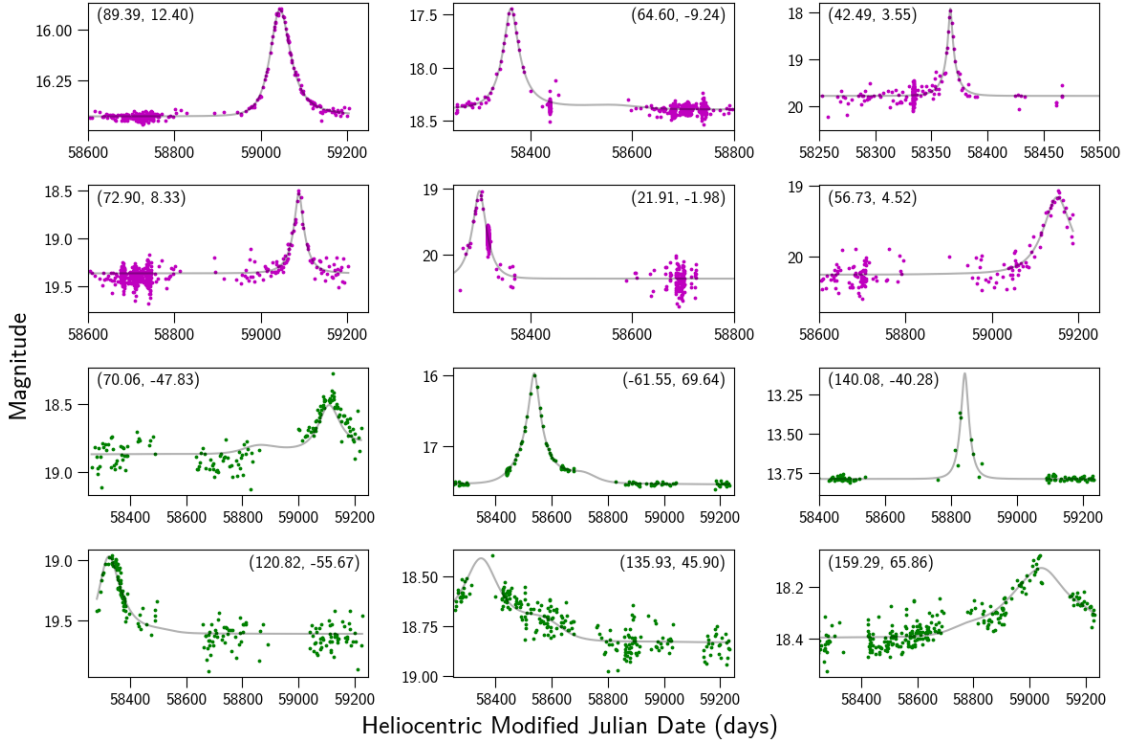


Figure 5.10: Our level 6 catalog contains 93 Galactic plane microlensing events (purple) and 35 events outside of the plane (green). These sample lightcurves show the variety of timescales and magnitudes identified by our pipeline. The Galactic longitude and latitude of each event is printed in each corner. These 35 events nearly triple the total number of microlensing events yet discovered outside of the Galactic plane and bulge.

5.6 Discussion

The microlensing catalog resulting from this work is the largest collection of events discovered by an all-sky synoptic survey not dedicated to microlensing. Throughout this work we have developed new techniques for efficiently finding events in the massive ZTF dataset. Simulations in the same areas of the sky give us expectations against which to compare our results. We also can compare our results to surveys designed specifically for microlensing to get a handle on how well our non-traditional methods have performed.

Our catalog of microlensing events contains objects in regions of the sky where previous microlensing campaigns have not observed and simulations have not been run. We therefore need to limit our level 6 catalog to a sub-sample that overlaps with previous OGLE observations and PopSyCLE simulations to be able to make comparisons. The locations of our PopSyCLE microlensing simulations are shown in the right panels of Figure 5.9. For this comparison we select observable PopSyCLE events there are required to have $u_0 \leq 1.0$, a maximum amplification $\delta m_r \geq 0.3$ and a baseline magnitude $m_r \leq 21.5$. These requirements

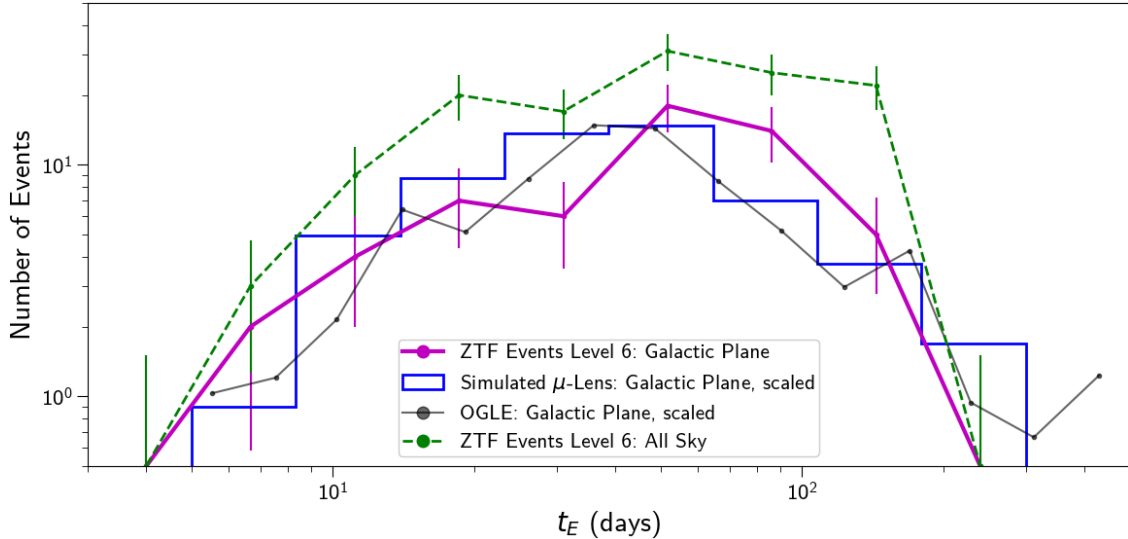


Figure 5.11: The Einstein crossing times for the level 6 event catalog across the entire sky (purple) and those found within our simulated PopSyCLE Galactic plane fields (green) appear similar. Each curve peaks at around 50 days and remains flat within the Poisson errors out to 80 days before sharply declining. These peaks are at longer crossing times than expected from the observable events contained within the overlapping PopSyCLE simulations (blue) scaled to match the number of ZTF Galactic plane level 6 events. The OGLE-IV Galactic plane fields (black), also scaled to match these events, largely agrees with our simulations. However it also shows the limits of our small sample size and short-duration survey to be able to detect the longest duration events.

produce events which are observable to ZTF, albeit without a correction for cadence or time span. We then scale the distribution of Einstein crossing times of these events to match the total number of events in our Galactic plane sub-sample for comparison in Figure 5.11.

The short-duration slope of our Galactic plane sub-sample is approximately equal to the simulation, with a deviation around 30 days. Both our sub-sample and the PopSyCLE catalogs peak at approximately 50 days. However our sub-sample appears to have an excess of long-duration events after this peak. The significance of this excess is marginal relative to the Poisson errors on our binned dataset. However, increased ZTF survey time could more than double the microlensing event sample size and greatly improve the significance of this signal. At the longest durations our pipeline is unable to recover events. The cut on sufficient baseline applied to level 4 candidates required $4t_E$ of observations outside of the event ($t_0 \pm 2t_E$). Applying this requirement for $8t_E$ of data to a three year survey puts an upper limit on our data of approximately $t_E \sim 135$ days. This explains the sharp drop-off in long-duration events around this point.

Mróz et al. (2020c) conducted the largest Galactic plane search for microlensing events prior to this work, discovering 630 events in the OGLE survey observing 3000 square degrees

of the Galactic plane between 2013 and 2019. ZTF observes an even larger footprint of the Galactic plane, with 6900 square degrees between $-15^\circ \leq b \leq 15^\circ$ and either $10^\circ \leq l \leq 180^\circ$ or $-180^\circ \leq l \leq -120^\circ$. However many of these fields are point toward or near the Galactic anti-center where stellar densities are low compared to the Galactic bulge fields observed by Mroz. OGLE also observed their fields for three more years than the ZTF DR5 dataset. We include in Figure 5.11 the distribution of Einstein crossing times found by Mroz for the Galactic plane ($|l| \geq 20^\circ$), re-scaled to match the number of events in our sample. The PopSyCLE and OGLE distributions are in good agreement and predict a relatively equal number of short and long duration events around a peak of approximately 50 days. While the PopSyCLE fields are run at different Galactic longitudes than the OGLE fields, this is in line with the findings of Mroz that the Einstein crossing time distribution within the plane is independent of longitude for $|l| \geq 20^\circ$.

Lam et al. (2020) found that black holes could be identified by placing microlensing events in the t_E - π_E space and looking for the events with the largest t_E and the shortest π_E . In Figure 5.14 we place our events in this space, alongside all observable events from within our PopSyCLE simulations. The simulated events are separated between those that have stellar lenses and those with a black hole lens. Level 6 events appear within parameter space most often occupied by events with stellar lenses. This can be caused by several factors. It is possible that all of our events are caused by stellar lenses. The relatively short timescale of DR5 combined with our strict requirements for data outside $2t_E$ of our event also places a limit on the longest events we are able to detect. And finally, small values of π_E are difficult to constraint with the quality of photometry that ZTF produces. The result is that we are unable to confidently claim that any of our microlensing events are black hole candidates.

Medford et al. (2020a) produced an estimate for the number of events detectable as a function of the faintest observable magnitude of ZTF imagery. They predicted that ZTF would discover ~ 500 events in the outer Galaxy for a r-band limiting magnitude of 20.6, and ~ 135 events in the same region for a r-band limiting magnitude of 19. In Figure 5.12 we compare the cumulative distribution function of our luminosity function to these results. Our all-sky level 6 catalog closely follows the yields predicted by Medford et al. (2020a) up until 19th magnitude. The sub-sample of level 6 events that fall within the Medford et al. (2020a) estimate footprint ($10^\circ \geq l \geq 100^\circ$, $-10^\circ \geq b \geq 10^\circ$) has a similar shape to the predictions. However we detect only 51 microlensing events within the Medford et al. (2020a) estimate footprint, 35 of which have $m_{\text{base,r}} \leq 19$.

We attribute the gap between the predictions of Medford et al. (2020a) and our estimate footprint sub-sample to several factors. The cut on level 2 candidates removed 16.4% of simulated microlensing event, and the three threshold cuts on level 3 candidates each removed 5% of simulated events. This accounts for missing 28% of possibly observable events. The level 4 catalog are all candidates that are well fit by a microlensing model and the cuts applied to them are to get a sample of high quality events. Therefore the order of magnitude drop from 5,029 candidates to 669 candidates likely removes many true events. Strict requirements on the amount of observable data outside of the event removes long duration events from the final catalog.

The level 2 and level 3 cuts retained 72% of possible observable events. The 128 events we discovered is 26% of the ~ 500 predicted by Medford et al. (2020a). We would have to assume that the level 4 cuts removed an additional 65% of true events from our catalog to match these predictions. While this is possible, given that the level 4 cuts selected only the highest quality events from a list of possible microlensing candidates, the difference between our yields and the predictions does suggest notable incompleteness to our method. However Medford et al. (2020a) did not take into account the effect of gaps in the data due to either weather or seasonality. This likely makes the true number of microlensing events lower than their prediction.

The lack of fainter sources in the level 6 catalog indicates that the ZTF limiting magnitude of 20.6 used in Medford et al. (2020a) may have been too large. This result also suggests that our pipeline selectively removed fainter stars with relatively larger photometric errors. This finding also points to the number of additional microlensing events that could be discovered if ZTF were reprocessed by combining subsequent observations into co-additions. Medford et al. (2020a) show the non-linear gains that are predicted to be achieved by this method, and our results corroborate that claim.

The level 6 catalog is not ideal for searching for black hole candidates due to the limit on t_E that our cuts impose. However, there are numerous microlensing events within ZTF-I that are still on-going. These were cut in level 4 due to the requirement that all candidates have a time of closest approach within $1t_E$ of the end of DR5 observations. This requirement limits our level 5 and 6 samples to completed events. 1690 candidates that had passed all level 4 cuts except this completeness requirement were instead classified as “level 4 ongoing”. If black hole microlensing events are contained within ZTF observations, they would most likely be long duration and therefore still ongoing.

The Einstein crossing times versus times of closest approach of level 4 ongoing candidates are shown in Figure 5.15. The majority are modeled to have already passed peak brightness by the end of DR5 data. However 128 candidates have $t_E \geq 150$ days and are projected to hit peak brightness during the ZTF-II campaign. These sources could be continuously observed to see if any decline in brightness is in agreement with a microlensing model. If found to do so, astrometric follow up could be combined with photometric measurement to weigh the mass of the lens and possibly make a detection of an isolated black hole. All 1690 level 4 ongoing candidates have been posted online for public use at <https://portal.nersc.gov/project/uLens/PUZLE>, following the schema outlined in Table 5.5.

Our search for microlensing in the public data releases was not performed in real-time. This approach is limited when attempting to find targets for astrometric follow-up. However, all of our microlensing event selection methods have been designed to be computationally inexpensive and quick to execute. Our methods could be applicable to a real-time data stream that triggers ongoing follow-up.

Our final level 6 catalog contains 35 events outside of the Galactic plane ($|b| \geq 15^\circ$), with several examples shown in 5.10. Previously attempts to detect microlensing with lenses in the stellar halo by observing background stars in the Magellanic Clouds and the Andromeda galaxy have yielded only a few events. In Wyrzykowski et al. (2011b) and Wyrzykowski

[et al. \(2011c\)](#), the OGLE survey discovered 2 and 3 events over 8 years toward the Large Magellanic Clouds and Small Magellanic Clouds respectively. [Alcock et al. \(2000\)](#) discovered 13 to 17 events in the 5.7 years that the MACHO collaboration observed the Large Magellanic Clouds. [Tisserand et al. \(2007\)](#) detected 1 event in the both Magellanic Clouds in the 6.7 years of the EROS-2 survey. Surveys toward the Andromeda galaxy have detected even fewer events, with [Novati et al. \(2014\)](#) detecting 3 events over the 4 year PLAN campaign and [Niikura et al. \(2019\)](#) discovering 1 short duration event after 7 hours of dense sampling. Our pipeline’s 35 events is more than all 20 to 25 previously discovered events, nearly tripling the total number of microlensing events discovered outside of the Galactic plane and bulge.

There are several explanations for the presence of microlensing events in significant numbers outside of the Galactic plane in our catalog. The probability of two stars appearing to co-align in lower stellar densities is small, but such yields could be possible by searching across ZTF’s all-sky footprint. Our sample could be contaminated by variations in stellar brightness that are well fit by a microlensing model but are instead other long-duration variability. The most exotic explanation would be the presence of MACHOs that exist in large enough numbers to lens background stars. While previous works claim to have eliminated the possibility of large numbers of MACHOs ([Alcock et al. 2001](#); [Wyrzykowski et al. 2011c](#)), more recent simulations that replace a monochromatic mass with an extended mass distribution have re-opened the door for MACHOs ([Carr et al. 2016](#); [Calcino et al. 2018](#)). Our detections of events toward the stellar halo could be interpreted as evidence for these claims, although further validation is needed.

There are several aspects of our pipeline that could be improved to increase the accuracy and yields of our results. Our DR3 to DR5 conversion method leaves out additional i-band observations that wouldn’t have been seen in earlier versions. A future implementation could keep an ordered set that tracks the `object` IDs that have been used and then uses each `radec_map` to push the new objects into an existing `source`. Simulated microlensing lightcurves could have been injected into our sample at the beginning of the pipeline to better measure completeness throughout the process. These lightcurves could also be mixed into the web portal candidates before expert scoring to measure the effect of human bias in identifying “clear microlensing” events. More careful comparison to simulations after completeness correction would also more accurately assess the effectiveness of our pipeline.

Several large synoptic surveys are due to see first light in the next few years, including the Rubin Observatory’s Legacy Survey of Space and Time and the Nancy Grace Roman Space Telescope. The techniques developed in this work are designed to accommodate the massive datasets that these surveys will generate. Previous pipelines included significant amounts of manual operation and expert assessment. While extremely successful for the surveys they operate on, these techniques cannot scale to the next generation of synoptic surveys. Our pipeline was designed to remove large numbers of events quickly and without significant computational cost. We fit progressively more complicated models onto the data in order to prevent spending resources fitting events that are less likely to be true microlensing. We eliminate all but the last few candidates before requiring human intervention. This approach, successfully executed on ZTF data, can unlock the potential for the next generation of

massive all-sky surveys to become microlensing machines.

Acknowledgments

This work is based on observations obtained with the Samuel Oschin Telescope 48-inch and the 60-inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy. We acknowledge support from the DOE under grant DE-AC02-05CH11231, Analytical Modeling for Extreme-Scale Computing Environments. J.R.L. and C.Y.L. acknowledge support by the National Science Foundation under Grant No. 1909641 and the National Aeronautics and Space Administration (NASA) under contract No. NNG16PJ26C issued through the WFIRST (now Roman) Science Investigation Teams Program.

5.7 PUZLE Database Tables

These tables record the database schema mentioned in Section 5.4.

- Table 5.6: source
- Table 5.7: star
- Table 5.8: source_ingest_job
- Table 5.9: star_ingest_job
- Table 5.10: star_process_job
- Table 5.11: candidate_level2
- Table 5.12: candidate_level3
- Table 5.13: candidate_level4 1
- Table 5.14: candidate_level4 2
- Table 5.15: candidate_level4 3

5.8 NERSC

5.8.1 Python Environment

The environment exists within the `miniconda3` executable located at `/global/cfs/cdirs/uLens/code/src/miniconda3`. Once this executable has been initialized and `./miniconda3/bin` has been prepended to the user's `PATH` environment variable, the user will be able to load the PUZLE conda environment.

`pymultinest` requires `LD_LIBRARY_PATH` to point to the `multinest.so` library file located at

`/global/cfs/cdirs/uLens/code/src/MultiNest/lib`. NERSC supports a module infrastructure that add or remove packages from a user's environment. To run `pymultinest`, users have to run `module load cmake/3.18.2` and `module unload craype-hugepages2M`.

The code for running this pipeline is located at <https://github.com/michaelmedford/puzle>. This is a private repository with access available upon request.

5.8.2 uLens Group

Users and their work at NERSC are organized into groups for the purpose of sharing storage and compute resources. All of the work for this paper is organized under the `uLens` group. The main storage directory is located at `/global/cfs/cdirs/uLens`. The DR5 sources and stars are stored on disk at `/global/cfs/cdirs/uLens/DR5`. Several libraries and additional executables are required to run this pipeline, in addition to python packages in the `ulens` or `PUZLE` environments are located in `/global/cfs/cdirs/uLens/code`. The source files for these executables are located in `/global/cfs/cdirs/uLens/code/src` and can be accessed by prepending `/global/cfs/cdirs/uLens/code/bin` to the `PATH` variable.

5.8.3 Database Access

The PUZLE database is hosted on the NERSC `postgresql` platform under the name `ulens`. This database is hosted on `nerscdb03.nersc.gov` and has users `ulens_user` and `ulens_admin`. An example command to access the database would be `psql -U ulens_admin -d ulens -h nerscdb03.nersc.gov` if the password credentials for `ulens_admin` were located in `/.pgpass`.

5.9 PUZLE Website

Candidates are able to be viewed, queried, commented and shared on the PUZLE website, located at <http://app.puzle.development.svc.spin.nersc.org/>. This website is written in flask

using the `jinja2` template engine and `alembic` migration environment. Code for this website is located at <https://github.com/michaelmedford/puzle>. This is a private repository with access available upon request.

5.9.1 Website Features

- Radial Search : Search for objects in the database by sky coordinates
- Filter Search : Search for objects in the database by candidate values
- Browse Candidates : Browse candidates and their microlensing parameters
- Profile : Personal page of candidates for each user
- Users : Links to the homepages of all the users on the site

Candidates can be linked to a ZTF alert ID through clicking on the Fetch ZTF IDs button. This sends a POST request to the Los Cumbres Observatory MARS (Make Alerts Really Simple) alert broker (<https://mars.lco.global/>) radial query API using the right ascension and declination of the candidate. The alert is then saved into the `_ztf_ids` column of the candidate's respective database table and displayed on the candidate's web page. A link is then available to send the user to the candidate's web page on the ALerCE (Automatic Learning for the Rapid Classification of Events) Explorer page (<https://alerce.online/>).

Users can also follow candidates that are then saved to their personal User page and see all of the users following a particular candidate. Users can leave comments on the web page of a candidate that are seen by all other collaborators looking on that page. The candidates that each user is following are available for all users to see on the Users page.

5.9.2 Hosting and Updating

The code is packaged into a `Docker` environment using the `Dockerfile` located in the main `PUZLE` directory. The build image is hosted on the NERSC image registry at registry.nersc.gov/ulens/puzle. Versions must be tagged and pushed along with the image name.

The website is hosted on NERSC's container-based platform `Spin` that enables web access to data stored in the `uLens` storage directories and databases. The `PUZLE` image is hosted on the the NERSC Rancher platform `uLens` production environment with current version `v1.7.3`. Edits to the website must be built into the `Docker` image, tagged with a new version, and pushed to the NERSC image registry. Then the version number must be updated on the NERSC Rancher platform and the instance redeployed. Documentation for `Spin` is located at <https://docs.nersc.gov/services/spin/>.

Cut Performed	Objects Remaining	Stars Remaining
Level 0		
ZTF DR5 lightcurves with $N_{\text{obs}} \geq 20$	1,768,352,633	-
$N_{\text{obs}} \geq 20$	1,744,425,342	702,431,964
$N_{\text{nights}} \geq 50$	1,011,267,730	563,588,562
Level 1		
$\eta \leq 1^{\text{st}}$ Percentile	-	10,227,820
PS1-PSC ≥ 0.645	-	8,987,330
Successful 4-parameter model fit	-	8,749,737
Duplicates between fields and filters removed	-	7,457,583
Level 2		
$\eta_{\text{residual}} \geq \eta * 3.82 - 0.077$	-	92,201
Level 3		
Successful 7-parameter model fit	-	91,598
$\chi_{\text{red}}^2 \text{ opt} \leq 4.805$	-	48,073
$2t_E$ baseline outside of $t_0 \pm 2t_E$	-	13,772
$\chi_{\text{red}}^2 \text{ flat} \leq 3.327$	-	12,195
$\pi_E \text{ opt} \leq 1.448$	-	11,489
$t_0 - t_E \geq 58194$	-	8,619
$t_0 + t_E \leq 59243$	-	6,929
Level 4		
Successful Bayesian model fit	-	6,929
$\sigma_{t_E}/t_E \leq 0.20$	-	5,029
$ u_0 \leq 1.0$	-	3,873
$b_{\text{sff}} \leq 1.2$	-	2,795
$\chi_{\text{red}}^2 \leq 3$	-	1,589
$4t_E$ baseline outside of $t_0 \pm 2t_E$	-	716
$t_0 - t_E \geq 58194$	-	683
$t_0 + t_E \leq 59243$	-	669
Level 5		
Manually assigned clear microlensing label	-	128
Level 6		

Table 5.3: Cuts for the PUZLE pipeline. Catalogs are defined as the collection of candidates remaining after previous cuts (i.e. 6929 candidates in the level 4 catalog). Cuts between level 0 and level 3 are applied to all objects within a star. Cuts between level 3 and level 4 are applied on the object within each star with the most number of observations. Cuts between level 4 and level 6 are applied to all objects within the star that are fit by the Bayesian model.

Label	Description	Number
clear microlensing	Model accurately follows a rise and fall in brightness with a clear region of uncorrelated non-lensing brightness measured outside of the event.	128
possible microlensing	Model accurately follows either a rise or fall in brightness, or does not have sufficient area of non-lensing outside of the event.	357
poor model / data	Model predicts a significant variation in brightness in areas without sufficient data.	113
non-microlensing variable	Correlated deviation from the model is present in either the non-lensing region outside of the event, or the lightcurve is similar in appearance to a supernova.	71

Table 5.4: Description of labels manually assigned to level 5 candidates.

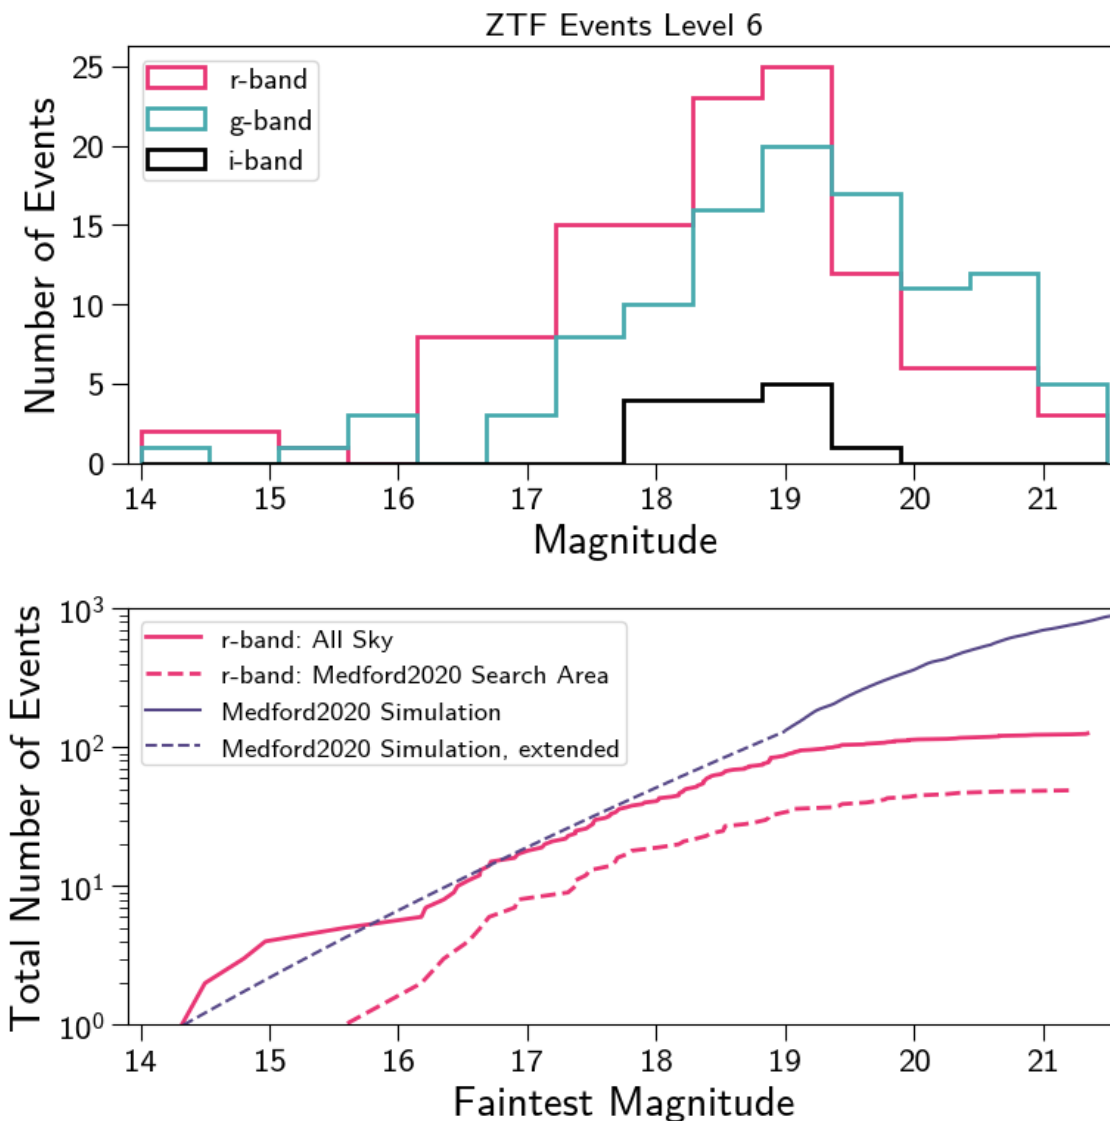


Figure 5.12: The baseline magnitudes of level 6 events are consistent across filters and follow the shape of the luminosity function predicted by [Medford et al. \(2020a\)](#) up to 19th magnitude. The 128 level 6 events contain 127 r-band, 110 g-band and 14 i-band baseline magnitudes that met the requirements to be fit by the Bayesian fitter. The distribution of these baseline magnitudes (top) appears increasingly plentiful at bright magnitudes, with all three filters dropping off to lower numbers past 19th magnitude. Comparing the cumulative number of events with r-band baseline magnitudes brighter than a given magnitude (bottom, red) to the [Medford et al. \(2020a\)](#) 3-year $l \geq 10^\circ$ predictions (purple) show a similar shape for magnitudes brighter than 19th, but with a clear deficit of events due to the selection effects of our pipeline. At r-band magnitudes fainter than 19th we are increasingly unable to observe any events. This suggests that executing a search on ZTF data reprocessed to combine images into co-additions could yield many more events as predicted by [Medford et al. \(2020a\)](#).




Figure 5.13: The source flux fractions of level 6 events drawn from across the sky are more uniform than results drawn from other works aimed toward the Galactic bulge. While the 127 r-band, 110 g-band and 14 i-band baseline magnitudes that met the requirements to be fit by the Bayesian fitter each have slightly different distributions, they all have very few events with low source flux fractions ($b_{\text{sff}} \lesssim 0.2$). This is expected in the relatively low stellar densities of the outer Galactic plane where the presence of neighboring stars is less prevalent than in the Galactic bulge. The g-band and r-band magnitudes peak at around $b_{\text{sff}} = 1$, where nearly all of the event flux originates in the source. The presence of sources with $b_{\text{sff}} > 1$ indicates that the background is overestimated, indicating the accuracy limits of the DR5 photometric catalogs.




Figure 5.14: The Einstein crossing times and Einstein parallax of all level 6 events with 1σ errors (red) fall right on top of the region where the PopSyCLE catalogs contain events with stellar lenses (blue). We do not find any events within the region of $t_E - \pi_E$ space where PopSyCLE predicts black hole lenses (black). The quality of our data prevents us from constraining small values of π_E , and the duration of our survey prevents the detection of the longest duration events. The PopSyCLE events shown are all of the observable events within the simulations.

Column Name	Datatype	Unit
id	string	-
ra	float	degrees
dec	float	degrees
t0	float	hmjd
t0_err	float	hmjd
tE	float	days
tE_err	float	days
u0_amp	float	-
u0_amp_err	float	-
piE_E	float	-
piE_E_err	float	-
piE_N	float	-
piE_N_err	float	-
mag_base_r	float	magnitude
mag_base_err_r	float	magnitude
mag_base_g	float	magnitude
mag_base_err_g	float	magnitude
mag_base_i	float	magnitude
mag_base_err_i	float	magnitude
b_sff_r	float	-
b_sff_err_r	float	-
b_sff_g	float	-
b_sff_err_g	float	-
b_sff_i	float	-
b_sff_err_i	float	-

Table 5.5: Exported Candidate Schema. Exported data can be found at <https://portal.nersc.gov/project/uLens/PUZLE>.




Figure 5.15: There are 1690 candidates that passed all cuts up to level 4 but were still ongoing as of the end of the DR5 dataset. These “level 4 ongoing” candidates have a time of closest approach (t_0) within one Einstein crossing time (t_E) of the end of DR5 observations (black solid line) or later. Events with an Einstein crossing time greater than 150 days (green line) that are continuing throughout ZTF-II (black shaded) are good candidates for continued observation. If they are found to decline in brightness in agreement with a microlensing model, they may be followed-up astrometrically to search for black hole lenses.

Table 5.6: `source`

Column Name	Datatype	Units
<code>id</code>	character varying(128)	
<code>ra</code>	double precision	degrees
<code>dec</code>	double precision	degrees
<code>lightcurve_filename</code>	character varying(128)	
<code>ingest_job_id</code>	bigint	
<code>object_id_g</code>	bigint	
<code>object_id_r</code>	bigint	
<code>object_id_i</code>	bigint	
<code>lightcurve_position_g</code>	bigint	
<code>lightcurve_position_r</code>	bigint	
<code>lightcurve_position_i</code>	bigint	
<code>comments</code>	character varying(1024)	
<code>_ztf_ids</code>	character varying(256)	
Index Name	Index Type	Columns
<code>source_pkey</code>	primary key	<code>id</code>
<code>ix_puzzle_source_ingest_job_id</code>	binary tree	<code>ingest_job_id</code>
<code>ix_puzzle_source_lightcurve_filename</code>	binary tree	<code>lightcurve_filename</code>
<code>source_q3c_ang2ipix_idx</code>	q3c_ang2ipix binary tree	<code>ra</code> , <code>dec</code>

Table 5.7: `star`

Column Name	Datatype	Units	Constraints
<code>id</code>	character varying(128)		Not Null
<code>ra</code>	double precision	degrees	Not Null
<code>dec</code>	double precision	degrees	Not Null
<code>ingest_job_id</code>	bigint		
<code>comments</code>	character varying(1024)		
<code>_ztf_ids</code>	character varying(256)		
Index Name	Index Type	Columns	
<code>star_pkey</code>	primary key	<code>id</code>	
<code>ix_puzzle_star_ingest_job_id</code>	binary tree	<code>ingest_job_id</code>	
<code>star_q3c_ang2ipix_idx</code>	q3c_ang2ipix binary tree	<code>ra</code> , <code>dec</code>	

Table 5.8: `source_ingest_job`

Column Name	Datatype	Units
<code>id</code>	<code>bigint</code>	
<code>ra_start</code>	<code>double precision</code>	degrees
<code>ra_end</code>	<code>double precision</code>	degrees
<code>dec_start</code>	<code>double precision</code>	degrees
<code>dec_end</code>	<code>double precision</code>	degrees
<code>started</code>	<code>boolean</code>	
<code>finished</code>	<code>boolean</code>	
<code>uploaded</code>	<code>boolean</code>	
<code>datetime_started</code>	<code>timestamp</code>	
<code>datetime_finished</code>	<code>timestamp</code>	
<code>slurm_job_id</code>	<code>integer</code>	
<code>slurm_job_rank</code>	<code>integer</code>	
Index Name	Index Type	Columns
<code>source_ingest_job_pkey</code>	primary key	<code>id</code>
<code>source_ingest_job_q3c_ang2ipix_idx</code>	<code>q3c_ang2ipix</code> binary tree	<code>ra_start</code> , <code>dec_start</code>

Table 5.9: `star_ingest_job`

Column Name	Datatype	Units	Constraints
<code>id</code>	<code>bigint</code>		Auto-generated Sequence / Not Null
<code>source_ingest_job_id</code>	<code>bigint</code>		Foreign Key / Not NULL
<code>started</code>	<code>boolean</code>		Default <code>false</code> / Not Null
<code>finished</code>	<code>boolean</code>		Default <code>false</code> / Not Null
<code>uploaded</code>	<code>boolean</code>		Default <code>false</code> / Not Null
<code>datetime_started</code>	<code>timestamp</code>		
<code>datetime_finished</code>	<code>timestamp</code>		
<code>slurm_job_id</code>	<code>integer</code>		
<code>slurm_job_rank</code>	<code>integer</code>		
Index Name	Index Type	Columns	
<code>star_ingest_job_pkey</code>	primary key	<code>id</code>	

Table 5.10: `star_process_job`

Column Name	Datatype	Constraints
<code>id</code>	<code>bigint</code>	Auto-generated Sequence / Not Null
<code>source_ingest_job_id</code>	<code>bigint</code>	Foreign Key / Not NULL
<code>started</code>	<code>boolean</code>	Default <code>false</code> / Not Null
<code>finished</code>	<code>boolean</code>	Default <code>false</code> / Not Null
<code>uploaded</code>	<code>boolean</code>	Default <code>false</code> / Not Null
<code>datetime_started</code>	<code>timestamp</code>	
<code>datetime_finished</code>	<code>timestamp</code>	
<code>slurm_job_id</code>	<code>integer</code>	
<code>slurm_job_rank</code>	<code>integer</code>	
<code>priority</code>	<code>integer</code>	
<code>num_stars</code>	<code>integer</code>	
<code>num_stars_pass_n_days</code>	<code>integer</code>	
<code>num_objs</code>	<code>integer</code>	
<code>num_objs_pass_n_days</code>	<code>integer</code>	
<code>num_objs_pass_eta</code>	<code>integer</code>	
<code>num_stars_pass_eta</code>	<code>integer</code>	
<code>num_objs_pass_rf</code>	<code>integer</code>	
<code>num_stars_pass_rf</code>	<code>integer</code>	
<code>num_objs_pass_eta_residual</code>	<code>integer</code>	
<code>num_stars_pass_eta_residual</code>	<code>integer</code>	
<code>num_candidates</code>	<code>integer</code>	
<code>epoch_edges</code>	<code>json</code>	
<code>eta_thresholds_low</code>	<code>json</code>	
<code>eta_thresholds_high</code>	<code>json</code>	
Index Name	Index Type	Columns
<code>star_process_job_pkey</code>	primary key	<code>id</code>

Table 5.11: candidate_level2

Column Name	Datatype	Units
id	character varying(128)	
ra	double precision	degrees
dec	double precision	degrees
ingest_job_id	bigint	
source_id_arr	character varying(128)[]	
color_arr	character varying(8)[]	
pass_arr	boolean[]	
num_objs_pass	boolean	
num_objs_tot	boolean	
idx_best	integer	
eta_best	double precision	
rf_score_best	double precision	
eta_residual_best	double precision	
eta_threshold_low_best	double precision	
eta_threshold_high_best	double precision	
t_E_best	double precision	days
t_0_best	double precision	hmjd
f_0_best	double precision	
f_1_best	double precision	
a_type_best	character varying(256)	
chi_squared_flat_best	double precision	
chi_squared_delta_best	double precision	
comments	character varying(1024)	
_ztf_ids	character varying(256)	
Index Name	Index Type	Columns
candidate_level2_pkey	primary key	id
candidate_level2_q3c_ang2ipix_idx	q3c_ang2ipix binary tree	ra, dec

Table 5.12: candidate_level3

Column Name	Datatype	Units
id	character varying(128)	
ra	double precision	degrees
dec	double precision	degrees
ingest_job_id	bigint	
source_id_arr	character varying(128)[]	
color_arr	character varying(8)[]	
pass_arr	boolean[]	
num_objs_pass	boolean	
num_objs_tot	boolean	
idx_best	integer	
eta_best	double precision	
eta_residual_best	double precision	
tE_best	double precision	days
t0_best	double precision	hmjd
u0_amp_best	double precision	
mag_src_best	double precision	mags
b_sff	double precision	
piE_E_best	double precision	
piE_N_best	double precision	
chi_squared_ulens_best	double precision	
num_epochs_best	integer	
num_days_best	integer	
comments	character varying(1024)	
_ztf_ids	character varying(256)	
Index Name	Index Type	Columns
candidate_level3_pkey	primary key	id
candidate_level3_q3c_ang2ipix_idx	q3c_ang2ipix binary tree	ra, dec

Table 5.13: `candidate_level4`

Column Name	Datatype	Units
<code>id</code>	character varying(128)	
<code>ra</code>	double precision	degrees
<code>dec</code>	double precision	degrees
<code>source_id_arr</code>	character varying(128)[]	
<code>color_arr</code>	character varying(8)[]	
<code>pass_arr</code>	boolean[]	
<code>num_objs_pass</code>	integer	
<code>num_objs_tot</code>	integer	
<code>idx_best</code>	integer	
<code>num_epochs_arr</code>	double precision[]	
<code>num_days_arr</code>	double precision[]	
<code>eta_arr</code>	double precision[]	
<code>eta_residual_arr</code>	double precision[]	
<code>t0_arr</code>	double precision[]	hmjd
<code>u0_amp_arr</code>	double precision[]	
<code>tE_arr</code>	double precision[]	days
<code>mag_src_arr</code>	double precision[]	magnitude
<code>b_sff_arr</code>	double-precision[]	
<code>piE_E_arr</code>	double-precision[]	
<code>piE_N_arr</code>	double-precision[]	
<code>chi_squared_ulens_arr</code>	double-precision[]	
<code>chi_squared_flat_arr</code>	double-precision[]	
<code>chi_squared_flat_outside_arr</code>	double-precision[]	
<code>num_3sigma_peaks_inside_arr</code>	integer[]	
<code>num_3sigma_peaks_outside_arr</code>	integer[]	
<code>num_5sigma_peaks_inside_arr</code>	integer[]	
<code>num_5sigma_peaks_outside_arr</code>	integer[]	
<code>num_days_inside_arr</code>	integer[]	
<code>num_days_outside_arr</code>	integer[]	
<code>chi_squared_flat_inside_arr</code>	double-precision[]	
<code>delta_hmjd_outside_arr</code>	double-precision[]	days

Table 5.14: candidate_level4

Column Name	Datatype	Units
num_pspl_gp_fit_lightcurves	integer	
source_id_arr_pspl_gp	character-varying(128)[]	
color_arr_pspl_gp	character-varying(8)[]	
slurm_job_id	integer	
node	character-varying(64)	
pspl_gp_fit_datetime_finished	timestamp without time zone	
pspl_gp_fit_datetime_started	timestamp without time zone	
pspl_gp_fit_finished	boolean	
pspl_gp_fit_started	boolean	
fit_type_pspl_gp	character-varying(128)	
chi2_pspl_gp	double-precision	
rchi2_pspl_gp	double-precision	
logL_pspl_gp	double-precision	
t0_pspl_gp	double-precision	hmjd
t0_err_pspl_gp	double-precision	hmjd
u0_amp_pspl_gp	double-precision	
u0_amp_err_pspl_gp	double-precision	
tE_pspl_gp	double-precision	days
tE_err_pspl_gp	double-precision	days
piE_E_pspl_gp	double-precision	
piE_E_err_pspl_gp	double-precision	
piE_N_pspl_gp	double-precision	
piE_N_err_pspl_gp	double-precision	
piE_pspl_gp	double-precision	
piE_err_pspl_gp	double-precision	
b_sff_pspl_gp	double-precision	
b_sff_err_pspl_gp	double-precision	
b_sff_arr_pspl_gp	double-precision[]	
b_sff_err_arr_pspl_gp	double-precision[]	
mag_base_pspl_gp	double-precision	magnitudes
mag_base_err_pspl_gp	double-precision	magnitudes
mag_base_arr_pspl_gp	double-precision[]	magnitudes
mag_base_err_arr_pspl_gp	double-precision[]	magnitudes

Table 5.15: `candidate_level4`

Column Name	Datatype	Units
<code>gp_log_sigma_pspl_gp</code>	double-precision	
<code>gp_log_sigma_err_pspl_gp</code>	double-precision	
<code>gp_rho_pspl_gp</code>	double-precision	
<code>gp_rho_err_pspl_gp</code>	double-precision	
<code>gp_log_omega04_S0_pspl_gp</code>	double-precision	
<code>gp_log_omega04_S0_err_pspl_gp</code>	double-precision	
<code>gp_log_omega0_pspl_gp</code>	double-precision	
<code>gp_log_omega0_err_pspl_gp</code>	double-precision	
<code>gp_log_sigma_arr_pspl_gp</code>	double-precision[]	
<code>gp_log_sigma_err_arr_pspl_gp</code>	double-precision[]	
<code>gp_rho_arr_pspl_gp</code>	double-precision[]	
<code>gp_rho_err_arr_pspl_gp</code>	double-precision[]	
<code>gp_log_omega04_S0_arr_pspl_gp</code>	double-precision[]	
<code>gp_log_omega04_S0_err_arr_pspl_gp</code>	double-precision[]	
<code>gp_log_omega0_arr_pspl_gp</code>	double-precision[]	
<code>gp_log_omega0_err_arr_pspl_gp</code>	double-precision[]	
<code>delta_hmjd_outside_pspl_gp</code>	double-precision	days
<code>level5</code>	boolean	
<code>category</code>	character-varying(128)	
<code>ongoing</code>	boolean	
<code>comments</code>	character-varying(1024)	
<code>_ztf_ids</code>	character-varying(256)	
<code>ogle_target</code>	character-varying(128)	
Index Name	Index Type	Columns
<code>candidate_level4_pkey</code>	primary key	id
<code>candidate_level4_q3c_ang2ipix_idx</code>	q3c_ang2ipix binary tree	ra, dec

Chapter 6

Conclusion

The last 30 years have seen microlensing detections transform from rare astrophysical phenomenon to a relatively common tool used to constrain Galactic morphology and stellar populations. Dedicated surveys are designed and executed for the sole purpose of detecting these events to great effect. Thousands of events are observed annually throughout the Galactic bulge, including detections of binary lenses and a growing number of exoplanets.

The next 30 years promises to reveal an unforeseeable number of cosmic mysteries. Massive synoptic surveys will observe an ever growing number of stars throughout the Milky Way Galaxy. These surveys promise measurements made with ever finer spatial and temporal resolutions across larger numbers of wavelengths. The Zwicky Transient Facility serves as an introduction to the challenges and opportunities of these large-scale surveys over the next few decades.

The approaches for discovering and characterizing events that heralded the age of microlensing will themselves require transformation to harness the power of these new datasets. This thesis has presented a series of tools and techniques that facilitate microlensing modelling and detection in this new era of observation. The methodologies contained within this thesis will hopefully contribute to continuing the tradition of innovation and discovery in microlensing science.

Bibliography

- Abdurrahman, F. N., Stephens, H. F., & Lu, J. R. 2021, On the Possibility of Stellar Lenses in the Black Hole Candidate Microlensing Events MACHO-96-BLG-5 and MACHO-98-BLG-6, [arXiv:2103.09923 \[astro-ph.SR\]](#)
- Afonso, C., Albert, J. N., Andersen, J., et al. 2003, [Astronomy & Astrophysics](#), **400**, 951
- Agol, E., Kamionkowski, M., Koopmans, L. V. E., & Blandford, R. D. 2002, [The Astrophysical Journal](#), **576**, L131
- Alcock, C., Allsman, R. A., Alves, D. R., et al. 2000, [The Astrophysical Journal](#), **542**, 281
- . 2001, [The Astrophysical Journal](#), **550**, L169
- Arredondo, J. N., Lu, J., Abdurrahman, F., et al. 2019, american Astronomical Society, AAS Meeting 233, id.369.02
- Aubourg, E., Bareyre, P., Bréhin, S., et al. 1993, [Nature](#), **365**, 623
- B. Paczyński, title = Gravitational microlensing by the galactic halo, j. . T. 1986, 304, 1
- Baraffe, I., Homeier, D., Allard, F., & Chabrier, G. 2015, [Astronomy & Astrophysics](#), **577**, A42
- Batista, V., Gould, A., Dieters, S., et al. 2011, [Astronomy & Astrophysics](#), **529**, A102
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2018, [Publications of the Astronomical Society of the Pacific](#), **131**, 018002
- Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019a, [Publications of the Astronomical Society of the Pacific](#), **131**, 068003
- . 2019b, [Publications of the Astronomical Society of the Pacific](#), **131**, 068003
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019c, [Publications of the Astronomical Society of the Pacific](#), **131**, 018002
- Bennett, D. P., Becker, A. C., Quinn, J. L., et al. 2002, [The Astrophysical Journal](#), **579**, 639
- Bernstein, G. M., Abbott, T. M. C., Desai, S., et al. 2017, [Publications of the Astronomical Society of the Pacific](#), **129**, 114502
- Bertin, E. 2006, in Astronomical Society of the Pacific Conference Series, Vol. 351, Astronomical Data Analysis Software and Systems XV, ed. C. Gabriel, C. Arviset, D. Ponz, & S. Enrique, 112
- Bertin, E. 2011, in Astronomical Society of the Pacific Conference Series, Vol. 442, Astronomical Data Analysis Software and Systems XX, ed. I. N. Evans, A. Accomazzi, D. J. Mink, & A. H. Rots, 435
- Bertin, E., & Arnouts, S. 1996, [Astronomy and Astrophysics Supplement Series](#), **117**, 393

- Bertin, E., Mellier, Y., Radovich, M., et al. 2002, in *Astronomical Society of the Pacific Conference Series*, Vol. 281, *Astronomical Data Analysis Software and Systems XI*, ed. D. A. Bohlender, D. Durand, & T. H. Handley, 228
- Branch, M. A., Coleman, T. F., & Li, Y. 1999, *SIAM Journal on Scientific Computing*, 21, 1
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *Astronomy & Astrophysics*, 564, A125
- Calchi Novati, S. 2018, *National Academy Committee on an Exoplanet Science Strategy*
- Calcino, J., García-Bellido, J., & Davis, T. M. 2018, *Monthly Notices of the Royal Astronomical Society*, 479, 2889
- Carr, B., Kühnel, F., & Sandstad, M. 2016, *Physical Review D*, 94
- Cassan, A., Kubas, D., Beaulieu, J.-P., et al. 2012, *Nature*, 481, 167
- Castelli, F., & Kurucz, R. L. 2003, in *IAU Symposium*, Vol. 210, *Modelling of Stellar Atmospheres*, ed. N. Piskunov, W. W. Weiss, & D. F. Gray, A20
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv e-prints, arXiv:1612.05560
- Choi, J., Dotter, A., Conroy, C., et al. 2016, *The Astrophysical Journal*, 823, 102
- Damineli, A., Almeida, L. A., Blum, R. D., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 463, 2653
- Dekany, R., Smith, R. M., Riddle, R., et al. 2020, *Publications of the Astronomical Society of the Pacific*, 132, 038001
- Derue, F., Afonso, C., Alard, C., et al. 2001, *Astronomy & Astrophysics*, 373, 126
- Dong, S., Mérand, A., Delplancke-Ströbele, F., et al. 2019, *The Astrophysical Journal*, 871, 70
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *The Astrophysical Journal*, 696, 870
- Durbin, J., & Watson, G. S. 1971, *Biometrika*, 58, 1
- Einstein, A. 1936, *Science (New York, N.Y.)*, 84, 506
- Evans, N. W., & Belokurov, V. 2002, *The Astrophysical Journal*, 567, L119
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *Monthly Notices of the Royal Astronomical Society*, 398, 1601
- García-Bellido, J., & Clesse, S. 2017, Constraints from microlensing experiments on clustered primordial black holes, [arXiv:1710.04694 \[astro-ph.CO\]](https://arxiv.org/abs/1710.04694)
- Gardner, J. P., Mather, J. C., Clampin, M., et al. 2006, *Space Science Reviews*, 123, 485
- Gaudi, B. S. 2012, *Annual Review of Astronomy and Astrophysics*, 50, 411
- Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *The Astronomical Journal*, 157, 98
- Godines, D., Bachelet, E., Narayan, G., & Street, R. 2019, *Astronomy and Computing*, 28, 100298
- Golovich, N., Dawson, W. A., Bartolić, F., et al. 2020, A Reanalysis of Public Galactic Bulge Gravitational Microlensing Events from OGLE-III and IV, [arXiv:2009.07927 \[astro-ph.GA\]](https://arxiv.org/abs/2009.07927)
- Gould, A. 1992, *The Astrophysical Journal*, 392, 442
- Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *Publications of the Astronomical Society of the Pacific*, 131, 078001
- Green, A. M. 2016, *Physical Review D*, 94

- . 2017, *Physical Review D*, 96
- Halko, N., Martinsson, P., & Tropp, J. 2011, SIAM review, 53, 217
- Hamadache, C., Guillou, L. L., Tisserand, P., et al. 2006, *Astronomy & Astrophysics*, 454, 185
- Husser, T.-O., von Berg, S. W., Dreizler, S., et al. 2013, *Astronomy & Astrophysics*, 553, A6
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111
- Johnson, H., Mitchell, R., Iriarte, B., & Wisniewski, W. 1966, *Communications of the Lunar and Planetary Laboratory*, 4, 99
- Jolliffe, I. T., & Cadima, J. 2016, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374, 20150202
- Jr., M. W. H., Lu, J. R., Lam, C. Y., et al. 2020, *The Astronomical Journal*, 160, 143
- Kains, N., Bramich, D. M., Sahu, K. C., & Calamida, A. 2016, *Monthly Notices of the Royal Astronomical Society*, 460, 2025
- Kalirai, J. S., Hansen, B. M. S., Kelson, D. D., et al. 2008, *The Astrophysical Journal*, 676, 594
- Kerins, E. J. 1995, *Monthly Notices of the Royal Astronomical Society*, 276, 785
- Kim, D.-J., Kim, H.-W., Hwang, K.-H., et al. 2018a, *The Astronomical Journal*, 155, 76
- . 2018b, *The Astronomical Journal*, 155, 76
- Kroupa, P. 2001, *Monthly Notices of the Royal Astronomical Society*, 322, 231
- Lam, C. Y., Lu, J. R., Hosek, M. W., Dawson, W. A., & Golovich, N. R. 2020, *The Astrophysical Journal*, 889, 31
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *Publications of the Astronomical Society of the Pacific*, 121, 1395
- Lu, J. R., Lam, C. Y., Medford, M., Dawson, W., & Golovich, N. 2019, *Research Notes of the AAS*, 3, 58
- Lu, J. R., Sinukoff, E., Ofek, E. O., Udalski, A., & Kozłowski, S. 2016, *The Astrophysical Journal*, 830, 41
- Mao, S., & Paczyński, B. 1996, *The Astrophysical Journal*, 473, 57
- Martinsson, P.-G., Rokhlin, V., & Tygert, M. 2011, *Applied and Computational Harmonic Analysis*, 30, 47
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2018, *Publications of the Astronomical Society of the Pacific*, 131, 018003
- Medford, M. S. 2021a, PUZLE Microlensing Catalogs <https://doi.org/10.5281/zenodo.5120987>
- . 2021b, zort: ZTF Object Reader Tool <https://github.com/michaelmedford/zort>
- Medford, M. S., Lu, J. R., Dawson, W. A., et al. 2020a, *The Astrophysical Journal*, 897, 144
- . 2020b, *The Astrophysical Journal*, 897, 144
- Medford, M. S., Lu, J. R., & Schlafly, E. F. 2020c, *Research Notes of the AAS*, 4, 38
- Medford, M. S., & Nugent, P. 2019, fringeZ: Correct Atmospheric Fringes in ZTF i-Band Images <https://github.com/michaelmedford/fringeZ>
- Medford, M. S., Nugent, P., Goldstein, D., et al. 2021, *Publications of the Astronomical Society of the Pacific*, 133, 064503

- Miyamoto, M., & Yoshi, Y. 1995, [The Astronomical Journal](#), 110, 1427
- Mróz, P., Udalski, A., Skowron, J., et al. 2017, [Nature](#), 548, 183
- . 2019, [The Astrophysical Journal Supplement Series](#), 244, 29
- Mróz, P., Street, R. A., Bachelet, E., et al. 2020a, [Research Notes of the AAS](#), 4, 13
- . 2020b, [Research Notes of the AAS](#), 4, 13
- Mróz, P., Udalski, A., Szymański, M. K., et al. 2020c, [The Astrophysical Journal Supplement Series](#), 249, 16
- Navarro, M. G., Minniti, D., Pullen, J., & Ramos, R. C. 2020, [The Astrophysical Journal](#), 889, 56
- Navarro, M. G., Minniti, D., & Ramos, R. C. 2017, [The Astrophysical Journal](#), 851, L13
- Niikura, H., Takada, M., Yasuda, N., et al. 2019, [Nature Astronomy](#), 3, 524
- Novati, S. C., Bozza, V., Paolis, F. D., et al. 2009, [The Astrophysical Journal](#), 695, 442
- Novati, S. C., Bozza, V., Bruni, I., et al. 2014, [The Astrophysical Journal](#), 783, 86
- Nucita, A. A., Licchelli, D., Paolis, F. D., et al. 2018, [Monthly Notices of the Royal Astronomical Society](#), 476, 2962
- Paczynski, B. 1996, [ARA&A](#), 34, 419
- Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2018, [Publications of the Astronomical Society of the Pacific](#), 131, 018001
- Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2019, [PASP](#), 131, 018001
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, [Journal of Machine Learning Research](#), 12, 2825
- Penny, M. T., Gaudi, B. S., Kerins, E., et al. 2019, [The Astrophysical Journal Supplement Series](#), 241, 3
- Powell, M. J. D. 1964, [The Computer Journal](#), 7, 155
- Price-Whelan, A. M., Agüeros, M. A., Fournier, A. P., et al. 2014, [The Astrophysical Journal](#), 781, 35
- Prince, T., & Zwicky Transient Facility Project Team. 2018, in [American Astronomical Society Meeting Abstracts](#), Vol. 231, American Astronomical Society Meeting Abstracts #231, 348.18
- Rahal, Y. R., Afonso, C., Albert, J.-N., et al. 2009, [Astronomy & Astrophysics](#), 500, 1027
- Raithel, C. A., Sukhbold, T., & Özel, F. 2018, [The Astrophysical Journal](#), 856, 35
- Refsdal, S., & Bondi, H. 1964, [Monthly Notices of the Royal Astronomical Society](#), 128, 295
- Ren, W., Cao, Q., Wu, D., et al. 2017, [Opt. Express](#), 25, 17039
- Renn, J. 1997, [Science](#), 275, 184
- Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, [Astronomy & Astrophysics](#), 409, 523
- Rowe, B. T. P., Jarvis, M., Mandelbaum, R., et al. 2015, [Astronomy and Computing](#), 10, 121
- Rybicki, K. A., Wyrzykowski, Ł., Klencki, J., et al. 2018, [Monthly Notices of the Royal Astronomical Society](#), 476, 2013
- Sajadian, S., & Poleski, R. 2019, [ApJ](#), 871, 205
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, [The Astrophysical Journal](#), 500, 525

- Sharma, S., Bland-Hawthorn, J., Johnston, K. V., & Binney, J. 2011, [The Astrophysical Journal](#), 730, 3
- Skilling, J. 2006, [Bayesian Analysis](#), 1, 833
- Spergel, D., Gehrels, N., Breckinridge, J., et al. 2013, arXiv e-prints, arXiv:1305.5422
- Stefano, R. D., & Esin, A. A. 1995, [The Astrophysical Journal](#), 448
- Sukhbold, T., Ertl, T., Woosley, S. E., Brown, J. M., & Janka, H.-T. 2016, [The Astrophysical Journal](#), 821, 38
- Sumi, T., Kamiya, K., Bennett, D. P., et al. 2011, [Nature](#), 473, 349
- Sumi, T., Bennett, D. P., Bond, I. A., et al. 2013, [The Astrophysical Journal](#), 778, 150
- Tachibana, Y., & Miller, A. A. 2018, [Publications of the Astronomical Society of the Pacific](#), 130, 128001
- Tisserand, P., Guillou, L. L., Afonso, C., et al. 2007, [Astronomy & Astrophysics](#), 469, 387
- Udalski, A., Szymański, M., & Szymański, G. 2015a, [Acta Astronomica](#), 65
- Udalski, A., Szymański, M. K., & Szymański, G. 2015b, OGLE-IV: Fourth Phase of the Optical Gravitational Lensing Experiment, [arXiv:1504.05966 \[astro-ph.SR\]](#)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, [Nature Methods](#), 17, 261
- von Neumann, J., Kent, R. H., Bellinson, H. R., & Hart, B. I. 1941, [The Annals of Mathematical Statistics](#), 12, 153
- Wyrzykowski, Ł., Kozłowski, S., Skowron, J., et al. 2011a, [Monthly Notices of the Royal Astronomical Society](#), 413, 493
- . 2011b, [Monthly Notices of the Royal Astronomical Society](#), 413, 493
- Wyrzykowski, Ł., Skowron, J., Kozłowski, S., et al. 2011c, [Monthly Notices of the Royal Astronomical Society](#), 416, 2949
- Wyrzykowski, Ł., Rynkiewicz, A. E., Skowron, J., et al. 2015, [The Astrophysical Journal Supplement Series](#), 216, 12
- Wyrzykowski, Ł., Kostrzewa-Rutkowska, Z., Skowron, J., et al. 2016, [Monthly Notices of the Royal Astronomical Society](#), 458, 3012
- Wyrzykowski, Ł., Mróz, P., Rybicki, K. A., et al. 2020, [Astronomy & Astrophysics](#), 633, A98