High-throughput phenotyping and modeling to dissect the genetic architecture of plant plasticity and growth

By

James Ta

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Biophysics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Daniel Runcie, Chair

Hao Cheng

Julin Maloof

Committee in Charge

2022

Abstract

Plant development consists of a tight web of interconnected traits that are jointly affected by genetics and the environment. These traits not only are affected by internal and environmental cues but can influence and be influenced by other traits as well. In this dissertation, I describe two projects where I develop frameworks for examining potential mechanisms underlying phenotypic variation and show how perturbations of one trait may cause other traits to vary. In the first project, I use the shade avoidance response as a model system for understanding the genetic architecture of late development plasticity. I leverage multiple data sets - bolting time, rosette biomass, inflorescence growth, etc. - in conjunction with path models to better describe cascading effects of QTL to explain colocalizations of QTL for different traits. I show how combinations of direct and indirect QTL effects can lead to variation in plasticity. In the second project, I take advantage of dense time series data to develop a model for measuring relative growth rates from growth curves called `SplineRGR`. `SplineRGR` is a parsimonious and flexible framework that demonstrates how perturbations in a latent trait - relative growth rate - can result in a wide range of patterns in overall growth dynamics. Overall, this dissertation illustrates that using complex models to leverage increasingly intricate data sets can provide insight into the genetic architecture of developmental traits and a more unified view of the genotype-phenotype map.

Contents

List of Figures

List of Tables

Acknowledgments

First, I would like to thank my advisor, Daniel Runcie, for his support and patience during my PhD journey. I would also like to thank the numerous people in the Runcie lab that have enriched my time working in the lab. I would like to thank Po-kai Huang, Taylor Crow, Sarah Odell, Saghi Nojoomi, and Haixiao Hu for their assistance and advice. This PhD would also have been impossible without help from the Maloof lab at UC Davis. I would like to thank Julin Maloof, Christine Palmer, and Cody Markelz for their wonderful guidance in the projects. I would especially like to thank Christine for her tremendous work and help in these projects. I would like to thank Hao Cheng for his advice and research suggestions. I would also like to thank Cynthia Weinig, Marcus Brock, Stephen Welch, and Nan An for their collaboration in our projects.

I am thankful for the funding from the Plant Sciences Departmental GSR at UC Davis that has supported the projects in this dissertation.

I would also like to thank the people who were pivotal in the experiments and analysis in my work. I would like to thank Cody Markelz and Kazu Nozue for their help in the chamber experiments. I would like to thank Nan An and Gaurav Ghosal for their computer vision work that facilitated the rosette measurements in my projects.

Lastly, I would like to thank my family for their love and support during the hardest times of my PhD. Going through my PhD would have been impossible without them.

Chapter 1

Introduction

Plant development is the complex process of cell division, expansion, and differentiation that results in various structural organs and maturation. Though studies tend to focus on one developmental stage due to research constraints, plant development spans a network of associated developmental stages and traits that are correlated with each other due to a shared genetic or physiological basis. These shared connections give rise to a multitude of life history strategies across different plant species, including fast life history paths focused on quick maturation and seed setting, and slower paths with larger seeds and long-lived leaves (Adler et al. 2014). However, studies often look at end-point phenotypes, viewing traits individually and in isolation without considering the effects of earlier or intermediate traits. By ignoring developmental dynamics, we lose information in terms of the genetics that are active at different times of plant development. Addressing this issue can result in better crop models for optimizing yield, as well as creating a more thorough understanding of plant development.

The interplay between plant developmental traits, their effects on plant growth, and the underlying genetics that contribute to variation in these traits are what motivates this disser-

tation. The central question this dissertation attempts to address is: How do the associations between traits and the genetic effects transmitted across development jointly influence phenotypic variation?

The approach I took to addressing this question was to measure both a wide variety of traits (i.e. bolting time, rosette and inflorescence biomass), as well as measuring one trait (rosette size) densely through time. With this data, I could address multiple facets of the central question: 1) How do genetics affect phenotypic variation through multiple developmental stages? And 2) How do genetics affect phenotypic variation at more transient time-scales?

Below, I give an overview of the model system *Arabidopsis*, discuss high-throughput phenotyping, and then describe how I use quantitative genetics and various data types for my dissertation.

1.1 *Arabidopsis* as a model organism

Owing to its physical size, fast development cycle, as well as small genome, *Arabidopsis* has become a key model organism for plant research (Woodward and Bartel 2018). Hundreds of *Arabidopsis* seedlings can be grown in Petri dishes, and the process of germination to seed-producing plant can be as little as 6 weeks. The *Arabidopsis* genome is only $\sim$132 megabases long, much smaller than many other plant genomes (Initiative 2000; Cheng et al. 2017). These qualities of *Arabidopsis* have been taken advantage of and have led to broader understanding of numerous fields in plant biology, for example: deeper insight into the photorespiration pathway (Somerville 2001), into building the foundation of the ABC model of floral development (Coen and Meyerowitz 1991), and elucidating plastid function (Martin et al. 2002). While

research of *Arabidopsis* cannot broadly represent all plant species, common mechanisms underlying plant responses and genetics, such as the shade avoidance response (SAR) (Casal 2012), have been translated from *Arabidopsis* studies to crop research (Carriedo et al. 2016).

Due to a thorough understanding of the *Arabidopsis* genome, as well as its fast development cycle, *Arabidopsis* has served as a powerful tool for understanding the genetic basis of phenotypic variation (Simon et al. 2008; Balasubramanian et al. 2009; El-Soda et al. 2014). From recombinant inbred lines (RILs) (Simon et al. 2008) to advanced multiparent mapping populations (Kover et al. 2009; Brock et al. 2020), a variety of carefully constructed populations are widely available for research. In addition, recent development of high-throughput phenotyping technologies have made denser and non-destructive phenotyping possible (Fahlgren et al. 2015b). These tools have greatly expanded the power of quantitative trait loci (QTL) mapping, and serve as the basis for the studies in this dissertation.

1.2 High-throughput phenotyping

As sequencing genomic information has become cheaper and more accessible, the limiting factor in understanding the genetic architecture of plant development is now measuring plant phenotypes. Phenomics is an emerging field that pioneers the use of high-throughput phenotyping (HTP) methods to drastically increase the number of phenotypes that can be studied. These technologies, in conjunction with next-generation sequencing, will allow us to further dissect phenotypic variation into its genetic components in plants.

One significant development in phenomics has been the development of image phenotyping technologies, which can non-destructively measure plants over time. This advancement

has allowed us to monitor plant growth densely over time, sometimes over multiple stages of development. Autonomous imaging reduces the labor required and the amount of measurement error compared to manual measurements, while also providing us more opportunities to integrate data sets of different types. Consequently, extensive effort has been made to adapt these technologies to measure plants in a variety of conditions - from laboratory to natural environments (An et al. 2016, 2017; Volpato et al. 2021).

With autonomous and non-destructive phenotyping now made possible using imaging platforms, we can greatly expand the end-point phenotypes of previous studies. We can now track early and intermediate phases of development more easily, which facilitates understanding of treatment perturbations through time and how they influence hard-to-measure traits such as yield. However, many QTL mapping studies still analyze a single trait at a time, which can be difficult with HTP data. This is because repeating analyses for large numbers of traits becomes computationally intensive, and can lack statistical power as the number of data points increases the number of tests and *p*-value corrections (Kendziorski et al. 2002; Ma et al. 2002; Wu and Lin 2006). Consequently, making use of this information is the current challenge and there is now a need for more efficient and thorough analysis of this data. In Chapter 2, I jointly model four distinct traits from multiple developmental stages on thousands of plants - to probe into the genetic architecture of plant plasticity to shade. In Chapter 3, I use data collected with an automated, high throughput imaging platform to better understand the genetic constituents of variation in rosette size with dense time-series data. These chapters showcase two cases of how large data sets can help us better understand plant development - 1) with data from multiple developmental stages, and 2) with densely sampled longitudinal data.

1.3 Quantitative genetics for understanding genetic effects through time

Genetic perturbations and their effects through plant development can be tracked by growing genetically distinct individuals and measuring the differences between their phenotypes over time. With the addition of carefully constructed mapping populations and sequencing markers of interest, we can pinpoint which loci contribute to phenotypic variation in a particular mapping population (i.e. QTL mapping) (Mackay 2001). However, to understand how the effects of specific loci and the functional relationships among traits contribute to phenotypic variation over time we will need to model both in concert.

Path analysis is a statistical technique that uses structured relationships to answer questions about indirect effects through mediator variables. Previous research has shown that genetic effects on later developmental traits can be explained as indirect effects arising from direct effects on earlier developmental traits (Li et al. 2006; Fournier-Level et al. 2013). In Chapter 2, I used path models to differentiate direct from indirect genetic effects on traits from multiple developmental stages. The motivation for using path analysis is the functional relationship between traits. Our modeling is based on the idea that vegetative and reproductive traits are functionally connected because they share a carbon source (leaves). For example, larger rosettes can lead to increased numbers of flowers and seeds. Consequently, traits such as rosette and inflorescence biomass are correlated, and factors that affect carbon allocation such as the reproductive transition should also influence these relationships. Path analysis can provide deeper insight into how functional relationships can modulate genetic effects through development.

An alternative to path analysis is to build a dynamic model of trait relationships using differential equations. In Chapter 3, I used differential equations to model the changes in plant size through time, using densely sampled measurements of leaf area collected with a HTP platform. Specifically, I modeled the instantaneous growth rate of a plant as a function of its current size. The ratio between instantaneous growth rates and current size is called the relative growth rate (RGR) and is widely used to compare growth rates in plants because it has a clear link to biological mechanisms (larger plants capture more light to fix more carbon to use for growth), and tends to be more constant than absolute growth rates as plants grow. However, it is well known that RGRs are not constant through time, so the key innovation in Chapter 3 is the use of splines to allow the RGRs of individual plants to change smoothly during their development to accommodate subtle changes in developmental dynamics among traits.

1.4 Overview of Chapter 2, Results and Implications

In Chapter 2, I took advantage of the extensive research on the shade avoidance response (SAR) in *Arabidopsis*. The SAR is a set of developmental and physiological behaviors that plants display when shaded, and serves a powerful model for understanding plant plasticity (Schmitt et al. 1999, 2003; Casal 2012). I investigate the genetic architecture of the SAR, and look into the possible mechanisms underlying variation in plasticity. One reason I use the SAR as a model system is because while the genetic architecture of the SAR in early seedlings is well-known, the genetic architecture underlying the SAR beyond the hypocotyl is less understood (Ciolfi et al. 2013; Nozue et al. 2015, 2018). Accordingly, the purpose of this

study was two-fold: 1) to identify QTL for later development SAR traits, and 2) to identify and propose several mechanisms maintaining variation in the SAR. One main result of this study was that we observed variation in direct/indirect QTL effects among the 7 different populations. This suggests that differing developmental strategies, combined with differences in underlying genetics, resulted in variation in plasticity. This knowledge can build a more unified understanding of the SAR, through assimilating knowledge of genetic architecture, environmental effects, and phenotypic relationships.

1.5 Overview of Chapter 3, Results and Implications

In Chapter 3, I developed a method, Spline modeling of Relative Growth Rates (or `SplineRGR`), that combines aspects of both parametric and non-parametric models to describe differences among growth curves. The main benefits of `SplineRGR` are its mechanistic relevance and simplicity. `SplineRGR` models RGR, where changes in RGR serve a mechanism to explain effects on size at numerous time points. In our simulations, `SplineRGR` accurately estimated RGR from various types of growth curves and had higher power to detect QTL than other methods in certain conditions. When applied to real data, `SplineRGR` did not find the most QTL compared to other methods, but was able to detect a QTL not found by other methods in one data set. This suggests that `SplineRGR` can find novel QTL associated with RGR. While our results concern genetic architecture and QTL, `SplineRGR` is generalizable to a variety of situations looking at different groups of plants: different treatments, different species, different genotypes, etc.

Chapter 2

Multiple loci control variation in plasticity to foliar shade throughout development in *Arabidopsis thaliana*

2.1 Abstract

The shade avoidance response is a set of developmental changes exhibited by plants to avoid shading by competitors, and is an important model of adaptive plant plasticity. While the mechanisms of sensing shading by other plants are well-known and appear conserved across plants, less is known about the developmental mechanisms that result in the diverse array of morphological and phenological responses to shading. This is particularly true for traits that appear later in plant development. Here we use a nested association mapping (NAM) population of *Arabidopsis thaliana* to decipher the genetic architecture of the shade avoidance response in late-vegetative and reproductive plants. We focused on four traits: bolting time,

rosette size, inflorescence growth rate, and inflorescence size, found plasticity in each trait in response to shade, and detected 17 total QTL; at least one of which is a novel locus not previously identified for shade responses in *Arabidopsis*. Using path analysis, we dissected each colocalizing QTL into direct effects on each trait and indirect effects transmitted through direct effects on earlier developmental traits. Doing this separately for each of the seven NAM populations in each environment, we discovered considerable heterogeneity among the QTL effects across populations, suggesting allelic series at multiple QTL or interactions between QTL and the genetic background or the environment. Our results provide insight into the development and variation in shade avoidance responses in *Arabidopsis*, and emphasize the value of directly modeling the relationships among traits when studying the genetics of complex developmental syndromes.

2.2 Introduction

Because plants are sessile organisms and require light for energy, their ability to monitor and adjust to their light environment is essential to their fitness (Schmid 1992; Gratani 2014). Consequently, plants have photoreceptors to sense changes in the light environment, and have developmental and physiological responses to optimize fitness under non-optimal light conditions (Kami et al. 2010). The shade avoidance response (SAR) is a characteristic suite of responses to the proximity of nearby plants in competition for light, and is widely cited as a primary example of adaptive plasticity (Schmitt et al. 2003; Keuskamp et al. 2010; Bongers et al. 2014). Green plant tissues absorb red light and reflect far-red light, so a change in the ratio of red to far-red light, called the red:far-red ratio (R:FR), signals the presence of nearby

vegetation and elicits a SAR in receptive plants (Franklin and Whitelam 2005). The SAR is widely cited as an example of adaptive plant plasticity because the morphological and physiological changes are dramatic, and the adaptive benefit has been demonstrated in multiple populations and species (Morgan and Smith 1979; Dudley and Schmitt 1995; Schmitt et al. 2003). The SAR also has detrimental effects on yield in crops, and its genetics and management are important targets for optimizing yield (Ballaré et al. 1997; Carriedo et al. 2016; Wille et al. 2017).

Although the SAR is generally triggered by light quality (i.e. R:FR ratio), the specific morphological changes caused by shade differ across tissues and developmental stages, and also depend on the persistence and intensity of the light quality signal (Casal 2012, 2013). Typical SAR characteristics include changes in phenology, physiology and growth resulting in taller plants, but with reduced biomass, which helps a plant escape competition. Phenological changes usually include delayed germination, accelerated flowering, and accelerated seed set. Delaying germination allows a seed to optimize the light environment upon emergence when shading is temporary, while accelerating flowering is generally a strategy for cutting losses and making some seed when shading is persistent (Casal 2012). Elongated and more up-right organs – such as hypocotyls, petioles and stems – are common responses to reduced R:FR, and this response can help plants overtop neighbors and increase light capture (Casal 2012). However, not all organs display elongation, and adult plants often show other responses such as reduced branching and smaller biomass (Casal 2012; Carriedo et al. 2016). These contrasting SAR characteristics suggest that distinct mechanisms mediate the SAR across plant development, and recent research suggests that there are separate regulatory pathways for the SAR between the seedling and adult life stages (Nozue et al. 2015). Differentiating

the genetic mechanisms of the SAR among developmental stages is a central goal, as they remain less understood.

Not only is there variation in shade effects across developmental stages, but variation in the SAR is also observed across different species and among populations within the same species. For instance, the timing of bud outgrowth in response to shade is accelerated in silver birch (*Betula pendula*), delayed in white clover (*Trifolium repens*), and not affected in *Arabidopsis* (Demotes-Mainard et al. 2016). Similarly, a population of *Stellaria longipes* from a prairie environment dramatically elongated stems in response to shading, while a population from an alpine environment showed only a slight increase (Alokam et al. 2002). These within and among-species differences are thought to be adaptive (Schmitt et al. 2003). For example, elongated stems may help *Stellaria* plants outcompete neighboring vegetation in a prairie, but may not be beneficial in areas that lack crowding and overtopping by other plants (i.e. alpine environments) (Alokam et al. 2002). Clinal variation in other environmental variables, such as temperature and precipitation, have also been associated with variation in the SAR across *Arabidopsis* populations (Botto 2015). These results suggest that the SAR can evolve and that populations may harbor useful variation for genetically dissecting and manipulating the SAR in different species.

Despite variability in the SAR among species, many genetic mechanisms involved in sensing and responding to shading by other plants appear to be conserved across species. The phytochromes have been established as a mediator of the SAR in *Arabidopsis* (Franklin et al. 2003a,b; Franklin and Whitelam 2005), sorghum (Kebrom et al. 2006), maize (Sheehan et al. 2007), and tomato (Weller et al. 2000; Schrager-Lavelle et al. 2016). There are also similar genetic and hormonal mechanisms that control axillary bud growth in response to shade for both

Arabidopsis and crops. For example, shade repression of axillary bud growth is controlled by the transcription regulator *TB1* in sorghum, and its homologs *BRC1* and *BRC2* in *Arabidopsis* (Carriedo et al. 2016). The plant hormones auxin, cytokinin, and strigolactone are known to regulate axillary bud growth in *Arabidopsis* and sorghum (Carriedo et al. 2016). Auxin-related genes are upregulated in stem transcriptome profiles in tomato in shade conditions (Cagnola et al. 2012). Given the extensive genomic resources available in the model species *Arabidopsis*, studies of the SAR in this species can rapidly identify genes and mechanisms that could be useful for controlling the SAR in crops. For instance, insight on phytochrome function from *Arabidopsis* was used to repress the SAR in tobacco (Robson et al. 1996) and potato (Boccalandro et al. 2003), leading to increased harvest index and tuber yield, respectively.

Extensive variation in the SAR has been reported for *Arabidopsis*. The SAR for hypocotyl elongation and flowering time showed high genetic variation among 157 *Arabidopsis* accessions studied by Botto and Smith (2002) (Botto and Smith 2002). Botto (2015) additionally examined shade effects in 60 genotypes of *Arabidopsis* across 15 different populations and found that the shade plasticity for some reproductive traits was significantly different across populations and was correlated with environmental differences among populations (Botto 2015). The genetic basis of variation in several SAR traits in *Arabidopsis*, including hypocotyl length, petiole length, bolting time, and rosette diameter, has been studied by QTL mapping and GWAS (Jiménez-Gómez et al. 2010; Coluccio et al. 2011; Filiault and Maloof 2012). Studies of natural variation can complement mutation experiments for discovering novel SAR genes. For example, the circadian clock gene *ELF3* was first implicated in the SAR in *Arabidopsis* in a QTL mapping study (Jiménez-Gómez et al. 2010; Coluccio et al. 2011).

However, previous QTL mapping studies on the SAR in *Arabidopsis* have been limited in

several ways. First, only one QTL mapping experiment has studied the SAR in adult plants (Jiménez-Gómez et al. 2010). Second, most studies have been done in single biparental populations, which harbor limited genetic diversity. Third, existing studies have mapped QTL for each trait separately, and have not taken into account the associations between traits. This can limit power when multiple traits are correlated, and can be misled by indirect effects transmitted from one trait to another trait due to developmental and physiological relationships between traits. For example, a higher leaf area index indirectly leads to increases in yield due to higher levels of photosynthesis and carbon assimilates for plant growth (Heuvelink et al. 2005). Weinig (2000) showed that the light environment modulated elongation in velvetleaf, and this has indirect effects on fecundity through biomass (Weinig 2000). Fournier-Level et al. (2013) revealed that both genetic background and planting location contribute to life history variation, and that planting location affected indirect QTL effect sizes (Fournier-Level et al. 2013). Accounting for trait relationships in QTL studies can help describe the similarities and differences among the underlying genetics of early and late developmental SARs in this species.

We use a nested association mapping population (NAM) to characterize the genetic architecture of the SAR in *Arabidopsis thaliana* for four traits: bolting days, inflorescence growth, rosette biomass, and inflorescence biomass (Yu et al. 2008). Compared to biparental populations, our NAM population has higher genetic diversity, which increases QTL mapping power and detects QTL that are broadly important across populations. Surprisingly, we find that while there is a shade effect, there is little genetic variation in later developmental SAR compared to earlier developmental SAR. However, we do find 17 SAR QTL among 4 traits, and evidence of an allelic series for many of our QTL. Among these, we find QTL on chromosomes 4 and 5 that colocalize for multiple phenotypes, suggesting pleiotropy for later developmental SAR. To

determine if these QTL are truly pleiotropic, we estimate the direct and indirect effects of colocalizing QTL on traits throughout developmental time using path analysis. Because shading involves accelerated flowering, which in turn is associated with smaller plant size and biomass, our hypothesis is that QTL effects on later developmental traits (e.g. biomass) should primarily be indirect. We find that trait associations and direct QTL effects on later developmental traits vary across populations and environments. This suggests that pleiotropy depends on both the genetic background and environment. These results highlight the importance of an integrated view of the genotype-phenotype relationship and the need to not only account for genetics and environment, but also phenotype relationships among traits throughout time.

2.3 Materials and Methods

Plant material

We used two mapping populations to study the genetics of the shade avoidance response in *A. thaliana*: a nested association mapping (NAM) population consisting of seven biparental populations with 1152 total recombinant inbred lines (RILs) (Brock et al. 2020), and a diversity panel consisting of ~ 100 diverse accessions (Table S1). Col-0 (186AV) was the recurrent parent of all seven NAM populations. Blh-1 (180AV), Bur-0 (172AV), Cvi-0 (166AV), Ita-0 (157AV), Jea (25AV), Oy-0 (224AV), and Sha (236AV) were the alternative parents. F8 generation RILs were created through single-seed descent, selfing, and bulk multiplication. We obtained seeds for each RIL from the Versailles Arabidopsis Stock Center, and the seeds for the diversity panel accessions from Magnus Nordborg.

Growth conditions

Seeds were stratified for four days at 4℃ in 0.15% agar solution and then planted in 4x4-potted trays (East Jordan Plastics: EJP804-200) filled with Sungrow Sunshine Mix #1. To improve germination rates, soil surfaces were flattened with a custom tamper before planting seeds. 2 - 3 seeds of the same RIL were planted in the center of each pot. Each pot was thinned to one plant after one week.

Plants were grown in the Controlled Environmental Facilities at UC Davis in five experiments from 05/13 - 08/15. Light was provided by fluorescent light bulbs at 100 μmol photosynthetic active radiation (PAR), and supplemented by LEDs with different red:far-red ratios (R:FR) to simulate sun (R:FR ratio > 1.0) and foliar shade (R:FR ratio $\sim$0.5) conditions (Franklin and Whitelam 2005). Daylength in both conditions was set to 16h light, 8h dark and the temperature set to 22℃. There were 3 shelves (blocks) for each treatment in each experiment, and in total between 4 - 10 replicates of each RIL per treatment grown over all experiments. Photosynthetic active radiation (PAR) and R:FR were checked using a spectrophotometer at the start of each experiment to verify lighting and sun and shade conditions. Because shade affects germination rate, shade-treated plants were germinated in sun conditions (R:FR > 1.0) for approximately one week to ensure comparable germination rates between sun and shade-treated plants.

Trays were watered with 200 - 300 mL Hoagland solution and rotated 3 times per week. For each block, temperature and humidity were measured continuously using HOBO environmental loggers. Plants were sprayed to prevent and treat diseases and pests whenever necessary.

Traits measurements

Bolting time, inflorescence height, and dry rosette and inflorescence biomass were measured on each plant. Plants were scored 3 times a week for bolting (BD, measured as days after planting). Inflorescence height was measured from the base of the inflorescence to the tip of the main inflorescence, and was measured approximately right after being scored for bolting, and the first and second weeks after bolting. Because not all inflorescence height measurements were taken at the same time, we estimated the growth rate of the main inflorescence (IG) by taking the difference in height between the first and last inflorescence height measurements and dividing by the number of days between the first measurement and the last measurement. Whole rosettes and inflorescences were harvested two weeks after bolting (immediately after the last inflorescence height measurement), dried, and weighed to obtain dry biomass (RB and IB, respectively for dry rosette and inflorescence biomass).

Data were scanned for obviously erroneous data and measurement error, which were excluded from the subsequent statistical analyses.

Statistical analyses: QTL mapping

Traits were transformed using the Box-Cox procedure and subsequently z-transformed to satisfy the linear model assumptions of normality and constant variance (transformed data in File S1). We estimated shade responses for each line with the follow mixed linear model:

$$P_{ijkl} = SHELF_i + TRT_j + RIL_k + RIL:TRT_{ij} + e_{ijkl} \tag{2.1}$$

where *P* is the phenotype, *SHELF* refers to spatial block, *TRT* is light treatment (sun or shade), *RIL* is the genotype (Recombinant Inbred Line), *RIL:TRT* is the genotype-by-environment in-

teraction, and *e* is the error. *SHELF* and *TRT* were modeled as fixed effects, while *RIL* and *RIL:TRT* modeled as random effects. We fit the model as a Bayesian linear mixed model using the *brms R* package (Buerkner 2017). We used the student family of residuals when fitting the Bayesian mixed models to reduce the influence of potential outliers.

We fit models separately for each of the seven populations to estimate the percentage of total phenotypic variance explained (PVE) by genotype main effect (G-PVE) and gene-environment interactions (i.e. GxE-PVE) for each trait. PVE was calculated by dividing the respective random effect variances by the total phenotypic variance (i.e. the sum of the genetic variance, GxE variance, and residual variance). We then reported the average G-PVE and GxE-PVE over all populations for each trait. We also defined the coefficient of genetic variation in plasticity (CV_p) as the standard deviation of GxE for each trait standardized by the absolute value of the population mean plasticity, which is an alternative measure of the amount of genetic variation in plasticity in a population. Plasticity in this case refers to the differences in the genotype means between the simulated sun and shade conditions. We estimated 95% credible intervals for the shelf fixed effects and PVE estimates (Table S2) as the 2.5% and 97.5% quantiles of the samples from its posterior distribution.

We used the posterior means of the line GxE effects as phenotypes for QTL mapping (posterior means in File S2). QTL mapping was run using the *GridLMM* package (Runcie and Crawford 2019). *GridLMM* provides the flexibility of joint QTL mapping in multi-parent populations using linear mixed models, and can also prevent proximal contamination of markers, which improves QTL mapping power (Lippert et al. 2011). We developed a forward stepwise algorithm using *GridLMM* functions to fit multiple-QTL models to our data. By adopting a stepwise approach, we gain greater power to detect QTL by controlling for additional QTL

elsewhere in the genome. We first generated genotype probabilities for all markers (obtained from Brock et al. (2020)) using the *R/qtl* package (Broman et al. 2003). We then performed QTL scans for each shade response trait using the Haley-Knott regression approach (Haley and Knott 1992; Broman et al. 2003), including a random effect to account for genetic background effects based on genotypes at all markers > 10cM from the testing marker. Since nearby markers were highly correlated, we ran QTL scans using a reduced set of 464 markers by iteratively dropping pairs of markers with a correlation $> 99\%$ (full marker set in File S3). QTL models were run separately for each population. We combined results across the seven populations for joint QTL mapping by summing the log-likelihoods from each population at the testing marker, and then subtracting from this total the sum of log-likelihoods of null models fit to each population. This log-likelihood ratio was compared to a chi-sq distribution with 6 degrees of freedom for hypothesis testing. We generated a *p*-value threshold by permuting genotypes within each biparental population 1000 times (Cheng et al. 2010) and used the 95% quantile of the largest -log10(*p*) values per permutation as the entry *p*-value threshold to control the type I error rate at $\alpha = 0.05$.

We estimated uncertainty in QTL positions using the full set of 10,688 markers by calculating 95% confidence intervals for each QTL using an approach modeled on TASSEL's stepwise regression method (Bradbury et al. 2007). Briefly, we determined confidence bounds around each peak marker by sequentially adding a nearby marker to the QTL model at a greater and greater distance to the QTL peak. We defined the confidence interval bounds as the nearest marker positions that resulted in the QTL peak's *p*-value being $\leq \alpha$. The only difference in our method relative to TASSEL is that we determined the confidence intervals on just the first confidence interval scan with no subsequent scans. These QTL confidence intervals were

then annotated with known shade avoidance genes (combined list of genes from Nozue et al. (2015); Sellaro et al. (2017)). Intervals lacking annotated genes are considered to be novel SAR QTL and are likely to contain novel SAR genes.

To show that our pipeline gives results consistent with other methods, we repeated our analysis of the bolting day shade response (BD_SAR) with three other QTL mapping methods: GEMMA's LMM (Zhou and Stephens 2012), TASSEL (Bradbury et al. 2007), and QTL IciMapping (Meng et al. 2015). All methods were run across all populations jointly. For GEMMA and TASSEL, we used the full set of 182,314 SNPs; for QTL IciMapping we used the full set of 10,668 markers (both obtained from Brock et al. (2020)). Genotype, phenotype, covariate, and annotation data used for these methods are in File S4-S10. For GEMMA, we used the default settings to generate the kinship matrix and to run the linear mixed-model. For TASSEL, we used the default settings to run the stepwise algorithm but limited the maximum number of markers in the stepwise model to 10. We used the JICIM method of QTL IciMapping with the default settings and 1,000 permutations. For GEMMA and TASSEL we included population as a covariate. All methods find QTL on chromosomes 4 and 5, but there are differences in other QTL found (Figure S1 and Figure S2 in Supplementary Material R1, and Table S3). These differences might arise due to the statistical method used to find QTL; *GridLMM* estimates a separate effect of each marker for each population using Haley-Knott regression, while GEMMA and TASSEL use a GWAS approach that treats each SNP as bi-allelic. Overall, while there are discrepancies between the QTL found between methods, we used *GridLMM* because we were interested in comparing marker effects between populations. *GridLMM* can also provide an advantageous combination of controlling for population structure, reducing proximal contamination, and increasing QTL mapping power using a stepwise algorithm not

found in any other QTL mapping software.

Statistical analyses: path analysis

We used a QTL-path analysis to assess whether QTL that are shared between traits have separate direct effects on both traits, or if the QTL effect on one trait can be explained as an indirect effect on a trait expressed earlier in development. We built a path model to explain the developmental relationships among traits based on the time of measurement of each trait. We then fit a QTL-path model by performing a QTL scan for each trait starting with all possible paths from other traits included as fixed covariates. Trait order was determined by collection time and developmental timing: BD -> RB -> IG -> IB. The set of paths included for each later trait consisted of both main effects and plasticity effects of all earlier traits.

To create a final QTL-path model, we collected all colocalizing QTL from the QTL scans and built multi-QTL path models separately for the sun and shade conditions using the R package *lavaan* (Rossel 2012) with the multiple groups analysis (phenotype data used for path analysis in File S11). We then took the difference of QTL effects between environments to estimate the QTL effects of the shade response. QTL effects reported in the path analysis figures thus represent the differences in QTL effects between sun and shade conditions, unless otherwise specified. We used a backward elimination approach to reduce this model to only terms that were significant in either treatment. For each trait, all QTL and previous traits were included in the initial model as predictors. Non-significant terms ($p > 0.01$) in both treatments were removed through an iterative process: the term with the highest p-value was sequentially dropped from the model and then the model was re-fitted until all remaining predictor terms

were significant ($p < 0.01$) for that trait in either treatment. Model fit was evaluated according to the comparative fit index (CFI), the root mean square error of approximation (RMSEA), and the standardized root mean square residual (SRMR) (Hu and Bentler 1999). We then used mediation analysis in *lavaan* to estimate direct and indirect effects of QTL. A description of the equations for the QTL-path scans and the path analysis in *lavaan*, as well as an explanation of how direct and indirect effects are estimated, can be found in File S12.

Data availability

Scripts and analyses are available at `https://github.com/jkhta/sar_qtl`. The Bayesian mixed model and QTL mapping pipeline scripts were run on the FARM cluster at UC Davis. File S1 contains the transformed and standardized phenotype data. File S2 contains the posterior means for the genotype and GxE random effects used for QTL mapping. File S3 contains the markers used for the *GridLMM* analysis. File S4 contains the genotype file used for the GEMMA analysis. File S5 contains the phenotypes (bolting day shade responses) used for the GEMMA analysis. File S6 contains the family covariate used for the GEMMA analysis. File S7 contains the SNP annotation file used for the GEMMA analysis. File S8 contains the genotype file used for the TASSEL analysis. File S9 contains the phenotypes (bolting day shade responses) and family covariate data used for the TASSEL analysis. File S10 contains the phenotype and genotype information used for the QTL IciMapping analysis. File S11 contains the trait data in sun/shade conditions used for path analysis. File S12 has descriptions on the equations used in the QTL-path scans, and the estimation of direct and indirect effects in *lavaan*. Supplementary Material R1 contains Figures S1-S3. Table S1 contains the

names of about 100 natural accessions grown in addition to the NAM population. Table S2 contains the 95% credible intervals for the shelf fixed effects, G-PVE, and GxE-PVE. Table S3 contains the significant markers detected for the bolting day shade response BLUPs using TASSEL. Table S4 contains the summary statistics from the fitted Bayesian mixed models. Table S5 contains the trait effects - estimated from the path models - for the different populations in the sun condition. Table S6 contains the trait effects - estimated from the path models - for the different populations in the shade condition. Table S7 contains the QTL found for the genotype random effects using *GridLMM*. Supplemental material is available at figshare: https://doi.org/10.25387/g3.12108063.

2.4 Results

Variation in shade responses among populations

To determine the underlying genetics of SAR variation across a broad panel of *Arabidopsis* accessions, we quantified the genetic variation in shade responses of four later-staged developmental traits in a NAM population consisting of 7 biparental populations and a total of 1152 recombinant inbred lines. Plants showed the classic SAR syndrome: compared to sun conditions, plants in simulated shade bolted faster (-0.58 sd decrease), had faster inflorescence growth (0.27 sd increase), and had lower dry rosette biomass (-0.78 sd decrease) and lower dry inflorescence biomass (-0.15 sd decrease) (Table 2.1).

Overall, the variance in shade responses among genotypes was fairly small, with GxE-PVE ranging between 1.27% - 5.15%, which is an order of magnitude lower than the variances in

genetic main effects, which ranged between 13.81% - 52.04% (Table S4). However, coefficients of genetic variation in plasticity were moderate to large, ranging from 12% to 1596% (Table 2.1). We also estimated small GxE variances among the diversity panel, with GxE-PVE ranging between 0.39% - 4.94% (Table S4).

Trait	Intercept	Plasticity	CV_p
BD	0.57 (0.46 - 0.67)	-0.58 (-0.65 - -0.52)	0.12 (0.06 - 0.19)
IG	0.02 (-0.11 - 0.15)	0.27 (0.19 - 0.36)	1.57 (0.21 - 5.71)
RB	0.86 (0.72 - 0.99)	-0.78 (-0.86 - -0.7)	0.15 (0.09 - 0.24)
IB	0.38 (0.26 - 0.51)	-0.15 (-0.23 - -0.07)	15.96 (0.63 - 20.24)

Table 2.1: Posterior means of the intercept and treatment fixed effects (Plasticity), and the coefficient of genetic variation for plasticity (CV_p = $\sigma_{GxE} / |\mu_{Plasticity}|$) averaged over all populations for each trait. Values in parentheses next to each posterior mean are the 95% credible intervals for the means. BD, bolting days; IG, inflorescence growth; RB, dry rosette biomass; IB, dry inflorescence biomass.

Additive QTL

To determine the genetic architecture underlying SAR variation, we estimated shade responses for each line for each of the four traits (BD_SAR, RB_SAR, IG_SAR, and IB_SAR) and used these estimates as phenotypes for QTL mapping. We detected 17 SAR QTL across all shade response traits, with 2 - 8 QTL per trait (Table 2.2). Interestingly, we detect the most QTL for the bolting day shade response (BD_SAR) (8 QTL) and the least for dry inflorescence biomass shade response (IB_SAR) (2 QTL), even though the GxE-PVE for BD (1.27%) is lower than for IB (5.15%). Our QTL mapping results suggest that the genetic architecture underlying the SAR for later developmental shade responses is polygenic.

Trait	QTL	SNP PVE	QTL Marker	Chromosome	Left Bound	Right Bound	# Genes
BD_SAR	BD_SAR1_1	4.71	m_1_28847340	1	m_1_28607852	m_1_29478919	3
BD_SAR	BD_SAR1_2	3.16	m_1_29478919	1	m_1_29400200	m_1_29897126	3
BD_SAR	BD_SAR3_1	1.08	m_3_8066460	3	m_3_8040793	m_3_8658987	6
BD_SAR	BD_SAR4_1	10.98	m_4_41028	4	m_4_41028	m_4_527682	2
BD_SAR	BD_SAR4_2	2.00	m_4_9240644	4	m_4_7937660	m_4_9455527	9
BD_SAR	BD_SAR5_1	4.68	m_5_3142427	5	m_5_3062640	m_5_3475211	1
BD_SAR	BD_SAR5_2	1.37	m_5_7484984	5	m_5_7063023	m_5_8277645	3
BD_SAR	BD_SAR5_3	1.36	m_5_25961748	5	m_5_25950815	m_5_26346630	2
IG_SAR	IG_SAR2_1	2.07	m_2_11683361	2	m_2_11607434	m_2_12272151	0
IG_SAR	IG_SAR4_1	1.36	m_4_16640333	4	m_4_16212324	m_4_17289054	8
IG_SAR	IG_SAR5_1	3.72	m_5_4647184	5	m_5_3799350	m_5_5130837	7
IB_SAR	IB_SAR4_1	3.25	m_4_41028	4	m_4_41028	m_4_527682	2
IB_SAR	IB_SAR5_1	3.80	m_5_4110711	5	m_5_3799350	m_5_5018484	7
RB_SAR	RB_SAR4_1	6.32	m_4_41028	4	m_4_41028	m_4_527682	2
RB_SAR	RB_SAR4_2	2.14	m_4_8938713	4	m_4_8504098	m_4_9455527	7
RB_SAR	RB_SAR5_1	3.92	m_5_3142427	5	m_5_3062640	m_5_4251866	6
RB_SAR	RB_SAR5_2	2.07	m_5_25961748	5	m_5_25637221	m_5_26182104	2

Table 2.2: Quantitative trait loci (QTL) for the shade responses of each trait. SNP PVE, percent variance explained for the QTL; Left Bound, left marker of the 95% confidence interval; Right Bound, right marker of the 95% confidence interval. # Genes, number of annotated genes found for QTL. BD_SAR, bolting days shade response; IG_SAR, inflorescence growth shade response; IB_SAR, dry inflorescence biomass shade response. RB_SAR, dry rosette biomass shade response.

Most SAR QTL were found on chromosomes 4 and 5, and four QTL confidence intervals overlapped for multiple traits, suggesting pleiotropy. A region of $\approx 500,000$ bp on the top of chromosome 4 (*SAR4_1*, around 41,028 bp) was associated with BD_SAR, IB_SAR, and RB_SAR, and explained between 3.25% - 10.98% of the variation in the SAR found in this population (Table 2.2). A region of $\approx 1,500,000$ bp in the middle of chromosome 4 (*SAR4_2*, around 8,938,713 bp) was associated with the BD_SAR and RB_SAR. A region on the top of chromosome 5 (*SAR5_1*, between 3,000,000 and 5,000,000 bp) was associated with all four shade response traits, and explained between 3.72% - 4.68% of the variation among traits. A region at the end of chromosome 5 (*SAR5_2*, around 25,961,748 bp) was detected for BD_SAR and RB_SAR, and explained 1.36% - 2.07% of the variation. Not all SAR QTL were associated with multiple traits: markers m_2_11683361 and m_4_16640333 were detected only for IG_SAR. This suggests that there are both unique and shared aspects of genetic architecture between later developmental traits in the SAR.

Evidence for allelic series

One of the advantages of a NAM population is that the effect sizes of QTL can be compared across populations. We found clear evidence of multiple functionally distinct alleles at several QTL (Figure 2.1). For example, at BD_SAR4_1, 3 parents contributed alleles that increased BD_SAR relative to Col-0, 2 contributed alleles that decreased BD_SAR relative to Col-0, and the remaining parents contributed alleles that were similar to Col-0. At other QTL (e.g. BD_SAR5_1), only one or two parents contributed an allele that differed significantly from the Col-0 common reference (Figure S3 in Supplementary Material R1). We did not observe any QTL

where the Col-0 allele was different from every other parent.

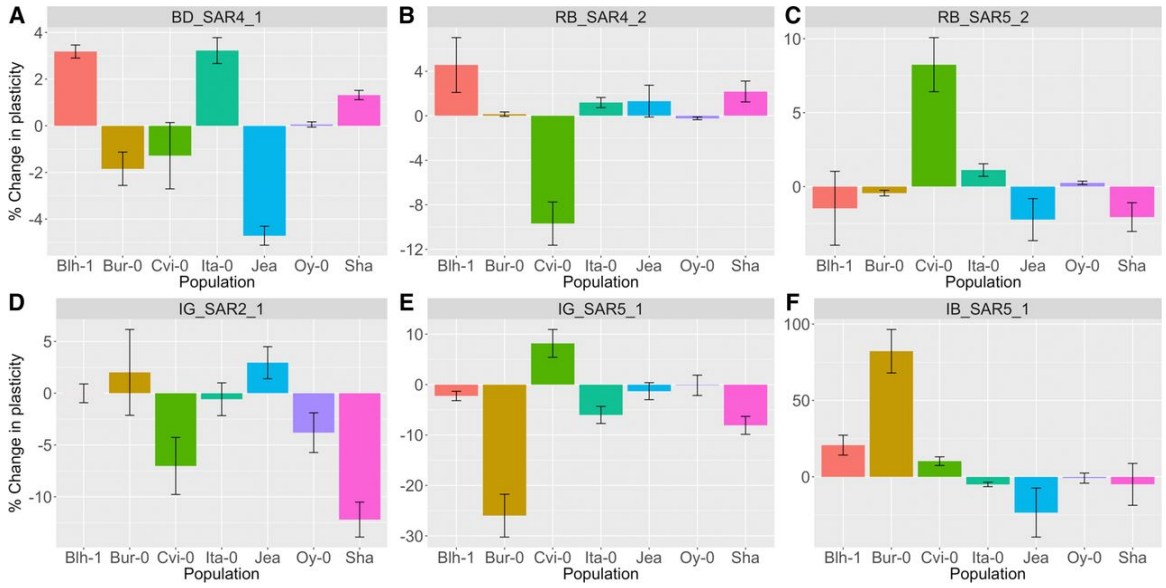

Figure 2.1: Allelic series among selected SAR QTL. % changes in plasticity relative to the Col-0 allele (allelic-specific change in plasticity / average plasticity) are plotted for selected SAR QTL. Each panel represents a different SAR QTL (panel title). Each bar represents the effect of the non-Col-0 allele in one of the seven different biparental populations. Blh-1, Blh-1 x Col-0; Bur-0, Bur-0 x Col-0; Cvi-0, Cvi-0 x Col-0; Ita-0, Ita-0 x Col-0; Jea, Jea x Col-0; Oy-0, Oy-0 x Col-0; Sha, Sha x Col-0. Error bars represent one standard error of the estimated allele substitution effect.

Gene annotation

We annotated each QTL region using a list of genes previously associated with the SAR in *Arabidopsis* from Nozue et al. (2015) and Sellaro et al. (2017), and listed the number of SAR genes under each QTL (Table 2.2). Many of these QTL have candidate genes that have been implicated in the mechanism of the SAR through mutant knockouts; however, several have not been shown to vary among natural accessions for the SAR. Additionally, we found 1 SAR QTL that does not contain any previously identified SAR genes: IG_SAR2_1. This

region represents a novel SAR QTL, and may provide new insight into the mechanisms of this plasticity.

Path analysis

Next, we used QTL-path analysis to determine if QTL effects on later-staged traits could be explained as indirect effects caused by direct effects of the QTL on earlier traits (in each environment), or earlier shade responses (differences between environments). QTL-path analysis identified a slightly different set of QTL (Figure 2.2B) as compared to the non-path analysis (Figure 2.2A). When mapping with earlier traits and shade responses as covariates (Figure 2.2B), we detected similar QTL on top of chromosomes 4 (for BD_SAR) and chromosome 5 (BD_SAR, IG_SAR, IB_SAR) as compared to the single-trait analyses (Figure 2.2A). However, the QTL on the top of chromosome 4 is no longer significant for IB_SAR and RB_SAR, and the QTL in the middle of chromosome 4, on the top of chromosome 5, and at the end of chromosome 5 are no longer significant for RB_SAR. These results suggest that these QTL have indirect effects on IB_SAR and RB_SAR.

To determine if QTL for later-development traits could be explained as indirect effects of colocalized QTL for earlier-development traits, we quantified the direct and indirect effects of each QTL. We modeled QTL effects in sun and shade conditions separately, and then estimated the difference in their effects between sun and shade to determine the effect on shade responses. We then used path analysis to compare the magnitudes of direct and indirect QTL effects among the seven RIL populations (File S12). Fit indices for our models implied adequate fits to the data (average CFI > 0.97, average RMSEA < 0.08, and average SRMR < 0.08

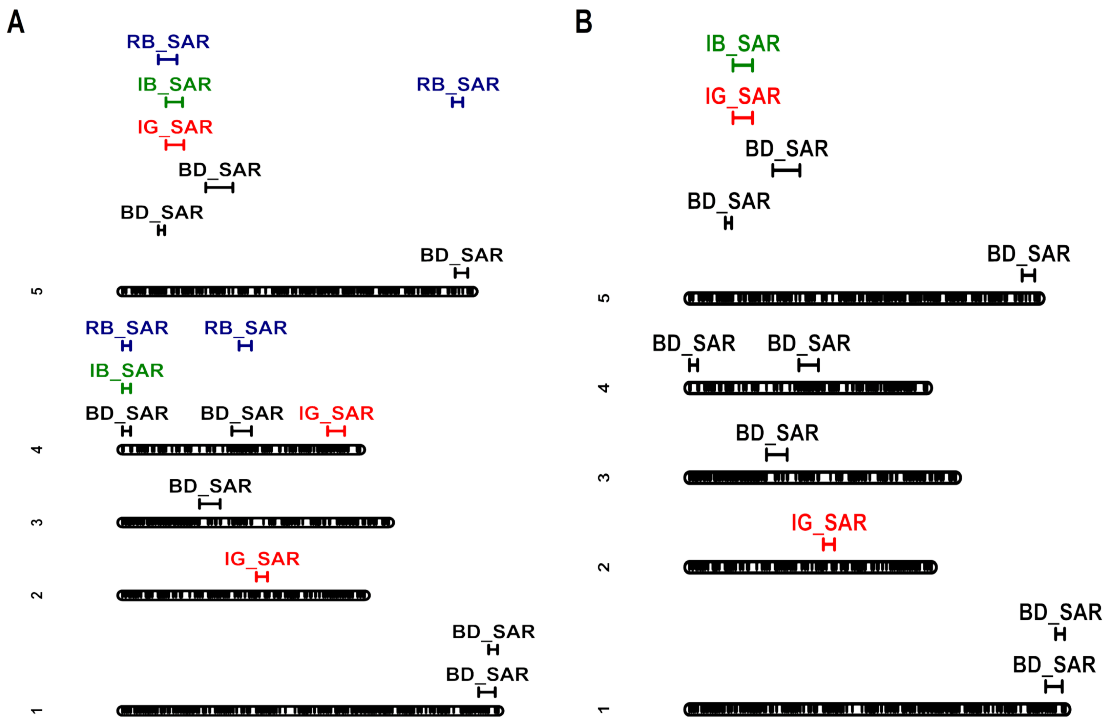

Figure 2.2: GBS-based single nucleotide polymorphism linkage map of *Arabidopsis*. Roughly 10,668 markers are distributed across 5 chromosomes. (A) Additive quantitative trait loci (QTL) detected without earlier traits as covariates and (B) with earlier traits as covariates. 95% confidence bounds for each QTL are also shown. Overlapping confidence interval bounds suggest colocalization of QTL. BD_SAR, bolting days shade response; RB_SAR, dry rosette biomass shade response; IG_SAR, inflorescence growth shade response; IB_SAR, dry inflorescence biomass shade response.

for all models). A conceptual illustration of the path models is shown in Figure 2.3.

We treated the multiple QTL found on the top of chromosomes 5 for the different shade responses as a single QTL region in our path analysis. This is because the confidence bounds for RB_SAR5_1 overlap with the confidence bounds for BD_SAR5_1, IG_SAR5_1, and IB_SAR5_1.

Most colocalizing QTL had significant effects in only a subset of the populations (Figure 2.4). For *SAR4_1*, only populations created with Blh-1, Ita-0, Jea, and Sha showed differences

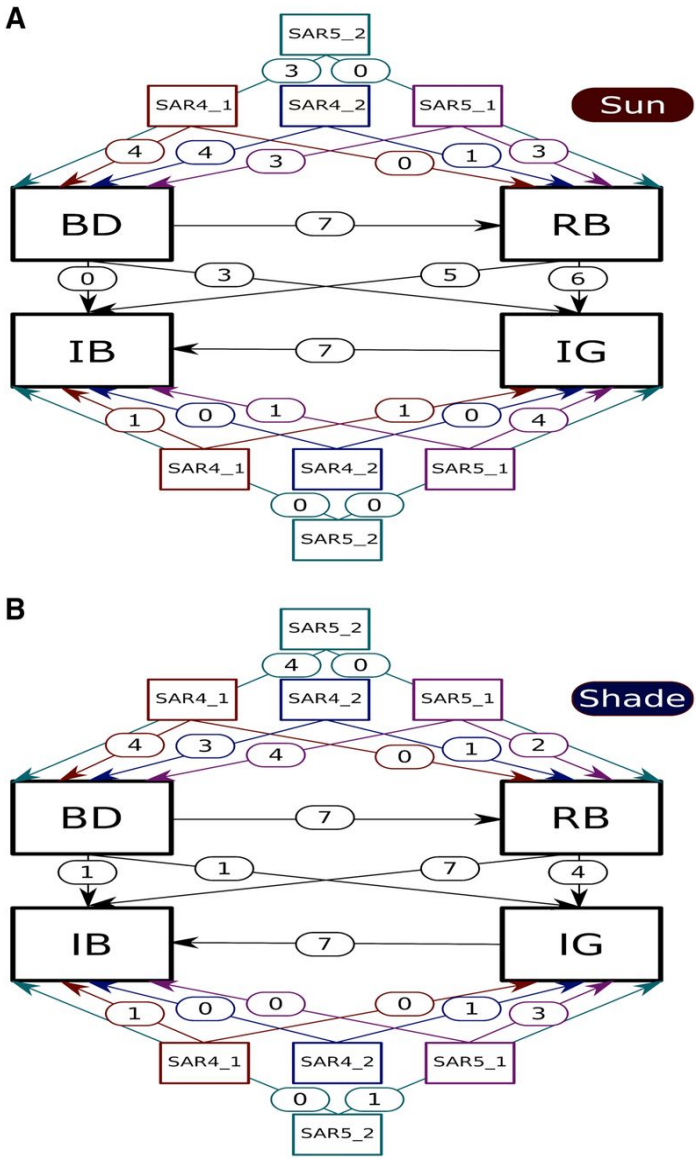

Figure 2.3: Representation of the fitted path models for sun and shade conditions. Directed arrows represent direct effects. The numbers within the arrows are the number of significant associations across populations ($p < 0.01$). BD, bolting days; IG, inflorescence growth; RB, dry rosette biomass; IB, dry inflorescence biomass. (A) Number of significant path effects in the simulated sun environment. (B) Number of significant path effects in the simulated shade environment. *SAR4_1*, the QTL at the top of chromosome 4 that colocalized for multiple traits. *SAR4_2*, the QTL in the middle of chromosome 4. *SAR5_1*, the QTL at the top of chromosome 5. *SAR5_2*, the QTL at the end of chromosome 5.

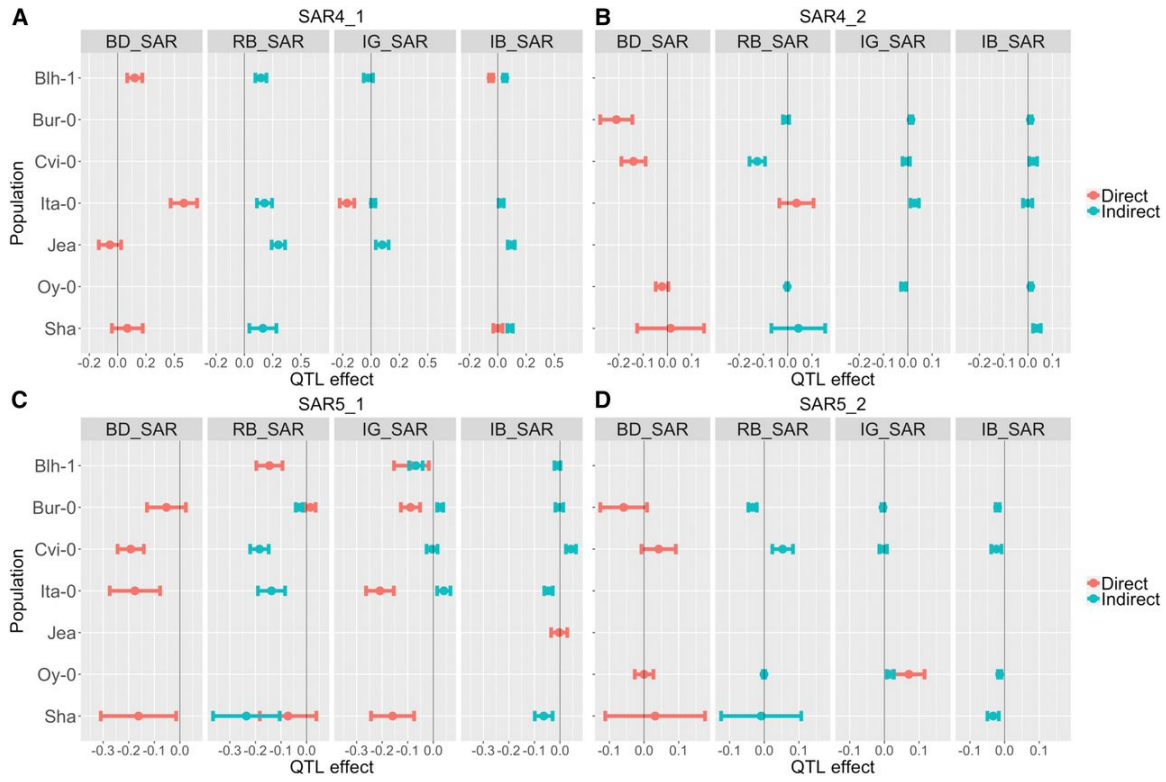

Figure 2.4: Direct and indirect effects of colocalizing quantitative trait loci (QTL) across populations and traits. Each panel represents a different shade response, going from earlier development (left) to later development (right). The y-axis depicts the different biparental populations, denoted by the non-recurrent parent of the biparental population. Blh-1, Blh-1 x Col-0; Bur-0, Bur-0 x Col-0; Cvi-0, Cvi-0 x Col-0; Ita-0, Ita-0 x Col-0; Jea, Jea x Col-0; Oy-0, Oy-0 x Col-0; Sha, Sha x Col-0. Direct effects are in orange while indirect effects are in teal. Each point represents the estimated QTL effect and the bars represent one standard error of the mean. Non-significant direct effects are not shown; consequently, downstream indirect effects from non-significant direct effects are not shown. BD_SAR, bolting days shade response; RB_SAR, dry rosette biomass shade response; IG_SAR, inflorescence growth shade response; IB_SAR, dry inflorescence biomass shade response. (A) QTL effects for *SAR4_1*, the QTL at the top of chromosome 4 that colocalized for multiple traits. (B) QTL effects for *SAR4_2*, the QTL in the middle of chromosome 4. (C) QTL effects for *SAR5_1*, the QTL at the top of chromosome 5. (D) QTL effects for *SAR5_2*, the QTL at the end of chromosome 5.

in QTL effects between sun and shade conditions. In the Blh-1 population we observed a positive direct QTL effect on the response to shade for BD_SAR; for the Bur-0 population, however, the direct effect of *SAR4_1* was non-significant. In later developmental traits, indirect effects for *SAR4_1* were non-zero in some, but not all, populations. For example, indirect effects of *SAR4_1* on RB_SAR and IB_SAR were positive in the Blh-1, Ita-0, Jea, and Sha populations.

For *SAR4_2*, we observed direct effects on BD_SAR and RB_SAR, but only indirect effects on IG_SAR and IB_SAR. In contrast, *SAR5_1* shows more direct effects on later developmental traits; including negative direct effects on RB_SAR and IG_SAR in the Blh-1, Ita-0, and Sha populations. Lastly, *SAR5_2* had direct effects on IG_SAR for the Oy-0 population, and indirect effects on RB_SAR and IB_SAR. Interestingly, though we do not detect *SAR5_2* for IG_SAR in our QTL mapping (Figure 2.2A), we find that *SAR5_2* has direct effects on IG_SAR (Figure 2.4); this discrepancy might be due to the more stringent significance thresholds in our QTL mapping method compared to our path analysis modeling.

These differences in direct and indirect QTL effects across populations potentially arise due to different trait and QTL effects in different environments. For instance, BD generally had a larger effect on RB in shade conditions than in sun conditions across populations (Table S5 and Table S6). BD effects on RB (RB $\sim$ BD) ranged between 0.10 - 0.73 in sun and between 0.11 - 0.84 in shade.

FRI and FLC may underlie the QTL on top of chromosomes 4 and 5

We detected strong QTL on top of chromosomes 4 (*SAR4_1*) and 5 (*SAR5_1*) for multiple shade response traits, including bolting time main effects (averaged over the two environments) (Table S7). These QTL overlap the major flowering repressor genes *FRIGIDA* (*FRI*) and *FLOWERING LOCUS C* (*FLC*), respectively. In low R:FR conditions, flowering is known to be accelerated because the repression of the floral transition by *FRI* and *FLC* is bypassed (Wollenberg et al. 2008). Therefore, *FRI* and *FLC* are likely candidate genes for *SAR4_1* and *SAR5_1*.

However, while *FRI* and *FLC* are within *SAR4_1* and *SAR5_1*, respectively, the 95% confidence intervals for these QTL span several Mb, so other loci in these regions may also be involved in these populations. On the other hand, since *FRI* and *FLC* have been extensively studied in *Arabidopsis*, the alleles of these genes have previously been characterized in the majority of the NAM parents in our study. Therefore, if *FRI* and *FLC* are the major causal genes underlying these QTL, the effect sizes and directions across populations should follow predictable patterns.

For instance, Col-0, Cvi-0, and Oy-0 have a non-functional *FRI* allele, while Blh-1, Bur-0, Ita-0, and Sha have a functional *FRI* allele (Lovell et al. 2013). Consequently, if *FRI* was the main driver of variation at the BD4_1 QTL, we expect that the QTL effects on BD to be close to zero for the Cvi-0 and Oy-0 alleles, and positive for the Blh-1, Bur-0, Ita-0, and Sha alleles. We find that the Cvi-0 and Oy-0 alleles do not delay bolting (QTL effect close to 0), while the Blh-1, Bur-0, Ita-0, and Sha alleles delay bolting (positive QTL effect) (Figure 2.5A). These results suggest that variation at *FRI* is the main driver of variation at the BD4_1 QTL.

Similarly, Col-0 and Blh-1 have a functional *FLC* allele while Bur-0 and Sha have either a weak or non-functional *FLC* allele (Gazzani et al. 2003; Werner et al. 2005; Simon et al. 2008). Consequently, we expect QTL effects at BD5_1 on BD to be close to 0 for the Blh-1 allele, and to be negative for the Bur-0, and Sha alleles. We find that the Blh-1 allele only slightly delays bolting (positive QTL effect close to 0), while the Bur-0 and Sha alleles accelerate bolting (negative QTL effect) (Figure 2.5B). These results then suggest that *FLC* is the main driver of variation at the BD5_2 QTL.

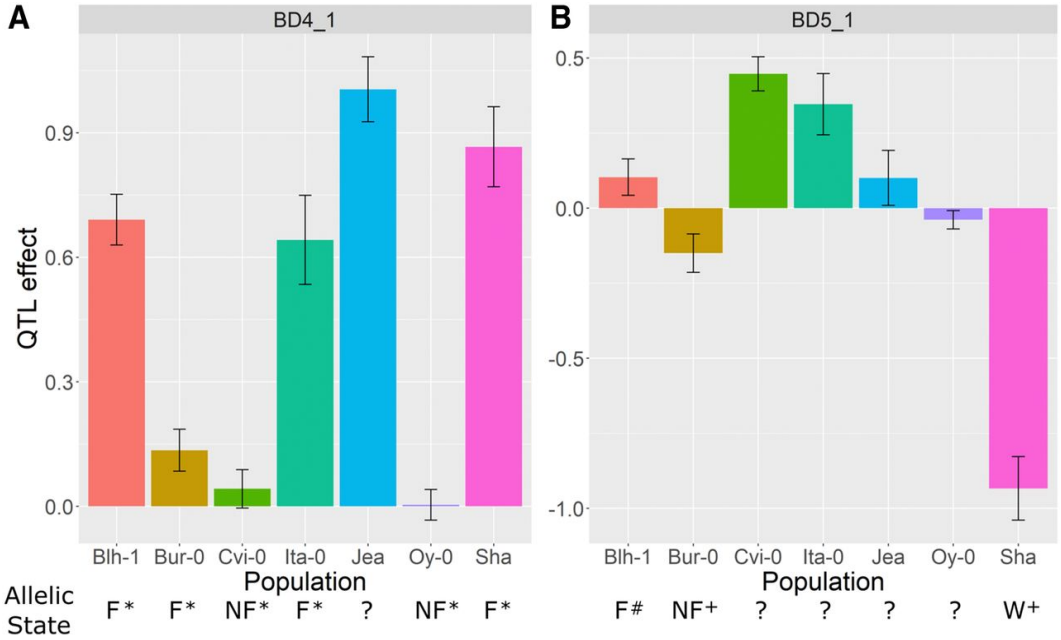

Figure 2.5: Effects of the BD4_1 and BD5_1 QTL for BD across populations. BD4_1 covers *FRI* and BD5_1 covers *FLC*. QTL effects relative to the Col-0 allele are plotted. Allelic state is either functional (F), non-functional (NF), unknown (?), or weak (W). Symbols next to allelic states represent the references where the information was collected: * (Lovell et al. 2013); # (Simon et al. 2008); + (Werner et al. 2005). Bars represent one standard error of the estimated effect of the non-Col-0 allele.

2.5 Discussion

General findings

We used a nested association mapping (NAM) population to investigate the diversity and genetic basis of variation in developmental responses to shade in *Arabidopsis thaliana*. Our study is the first in *Arabidopsis* to search for quantitative trait loci (QTL) for the shade responses of several late-development traits, including inflorescence growth (IG), rosette biomass (RB), and inflorescence biomass (IB), and includes a much greater sampling of genetic diversity than previous QTL mapping studies of the shade avoidance response (SAR) in this species (Jiménez-Gómez et al. 2010; Coluccio et al. 2011). Because of the large size of our study, the power of the NAM population, and the assessment of new SAR traits, we find at least one novel QTL that may be useful for future fine-mapping studies to discover genes involved in SAR regulation. Of the 17 SAR QTL we detect, only a few overlap with those found by Jiménez-Gómez et al. (2010), who also measured the SAR for later developmental traits. Similar to Jiménez-Gómez et al. (2010), we detect QTL near the end of chromosome 5. However, we also detect QTL on chromosome 1, on chromosome 4, and on top of chromosome 5 that were not detected by Jiménez-Gómez et al. (2010). This might be due to the greater genetic variation in the NAM population compared to the Bay-0 x Sha population, but the differences in detected QTL might also be due to the differences in the measures used for the shade responses. Jiménez-Gómez et al. (2010) used subtraction and residual indices on untransformed data, while we used genotype plasticity estimated from mixed models on transformed data. These discrepencies: the accessions that are represented by the RIL populations, the measure used as the SAR, and the use of untransformed or transformed data, could lead to

the contrast in QTL profiles seen between studies. We also use path analysis to determine the mechanisms of pleiotropy among QTL, and discovered that some QTL effects on later development can be explained as effects on earlier development. Fournier-Level et al. (2013) also reported increased indirect QTL effects on later developmental traits (Fournier-Level et al. 2013). However, this depends on the genetic background and environment. Overall, our work describes how foliar shade and genetics influence traits across developmental time.

Magnitudes of genetic variation in traits and trait plasticity

Our ability to detect QTL depends on the percentage of phenotypic variation that is due to genetic variation, which can be quantified by the percent variance explained (PVE) statistic. We find that variation in the SAR (GxE) among our NAM lines explained very little of the overall variation in any of the traits we measured (GxE-PVE ranged between 1.27% - 5.15%). This is an order of magnitude lower than the amount of variation explained by genotype main effects (G-PVE ranged between 13.81% - 52.04%), and also much lower than the amount of residual, or unexplained variation (E-PVE, which ranged between 46.69% - 81.04%, Table S4). We also observed lower GxE-PVE among our traits in this NAM population compared to the $\sim$ 15% GxE-PVE observed for hypocotyl elongation in a panel of 180 *Arabidopsis* accessions (Filiault and Maloof 2012). The low GxE-PVE was not a result of limited diversity among the eight NAM parents, as we observed similar GxE-PVE (< 5% across all traits) in a diversity panel of $\sim$100 accessions. This suggests lower variation in how *Arabidopsis* accessions respond to shade during later development when compared to the shade response in earlier development. Dechaine et al. (2014) also observed higher GxE variation for early internode elongation com-

pared to later internode elongation in *Brassica rapa*, suggesting that decreased GxE variation for later developmental traits is prevalent across multiple species (Dechaine et al. 2014). However, differences in chambers and lighting conditions compared to Filiault and Maloof (2012) could also contribute to differences in GxE-PVE (Filiault and Maloof 2012).

However, as a measure of the magnitude of plasticity variation, GxE-PVE can be misleading if the variation attributable to genotype main effects (G-PVE) is large (as this contributes to the total variation independently of GxE). We therefore also report the coefficient of genetic variation in plasticity (CV_p) as a metric of the magnitude of gene-environment interactions. CV_p compares the genetic variation in plasticity to the average plasticity across the populations. By this metric, traits for which some lines respond moderately to shade while others do not may be scored as showing higher genetic variation in plasticity than traits where all lines show strong plasticity, but vary in their magnitude. By the CV_p statistic, we observed considerable variation in the SAR of our traits (CV_p ranged between 12%-1596%).

Genetic diversity at key QTL

By using a multi-parent population, we were able to compare the effects of the same QTL across different donors. Our results provide evidence of allelic series for many of our SAR QTL. Allelic series have previously been observed in *Arabidopsis* for flowering time (Salomé et al. 2011) and seed dormancy (Kerdaffrec et al. 2016), and allelic variation has also been described for traits in response to shade. McNellis and colleagues (1994) described different allelic classes of *cop1* mutants and their effects on hypocotyl elongation in both simulated canopy shade and end-of-day far-red light treatments (McNellis et al. 1994). Previous QTL

mapping of the SAR in seedlings and adult plants have found two distinct alleles of *ELF3* that regulate hypocotyl elongation and bolting date in response to shade (Jiménez-Gómez et al. 2010; Coluccio et al. 2011). Similar to Jiménez-Gómez et al. (2010), we find different alleles for bolting time in response to shade as well as other later developmental shade responses. In our population, however, we detect more than two functionally distinct alleles for multiple QTL across all of our traits. The alleles vary in effect sizes, from small to moderate, and while most alleles change the plasticity less than |15%|, some alleles change plasticity as much as |80%|. This is similar to the magnitudes of the effects of polymorphisms on the shade response in genes like *COP1* and *ELF3* (McNellis et al. 1994; Jiménez-Gómez et al. 2010). In comparison, allelic effects on BD main effects were much higher (Table 2.1 and Figure 2.5), with alleles changing the average BD by as much as 0.95/0.57 = 163%. Our results suggest that while the range of allelic effects and their effect sizes on the SAR are small-to-moderate, allelic series are still important for variation in the SAR.

An allelic series can be caused by several possible mechanisms. 1) Multiple functionally distinct alleles may be present at the same gene among the 8 NAM parents, such as strong, weak, and non-functional versions of the same gene. 2) The causal variants in each of the 7 NAM families may reside in different genes, but we are unable to resolve multiple QTL due to the limited mapping resolutions within each family (our average QTL width was 0.82 Mb). 3) Even if there are only two functionally distinct alleles at the locus, the average effect of an allele may differ among the NAM families due to differences in genetic background, such as epistatic interactions with variants at other regions of the genome.

Distinguishing among these alternative models will require fine-mapping of each QTL across the NAM families and is beyond the scope of this study. However, our path anal-

ysis of the relationships among QTL and traits provides evidence that genetic background effects may be important. We observed several cases of colocalizing QTL among multiple traits, including QTL on chromosomes 4 (*SAR4_1*) and 5 (*SAR5_1* and *SAR5_2*). Using path analysis, we demonstrated that at least some of the QTL on later traits could be explained as indirect effects of the QTL effects on earlier traits during development. However, the breakdown between direct and indirect QTL effects varied among populations and between the sun and shade environments. If the functional relationships among traits vary among populations, then even if a QTL has the same effect on an early developmental trait among populations, the indirect effect of the QTL on a later trait may vary. This would then appear as an allelic series for the later trait. In this study, we only measured four later-development traits. Had we been able to observe many more traits throughout development, we would have been able to further characterize colocalizing QTL to distinguish allelic series of direct effects from allelic series that are the result of different indirect effects through trait relationships.

FRI and FLC as candidate genes

We found two colocalizing QTL on chromosomes 4 and 5 (*SAR4_1* and *SAR5_1*) for multiple shade responses and provided evidence that *FRI* and *FLC* are the drivers of variation at these loci. *FRI* and *FLC* are flowering repressor genes that control the initiation of flowering, and previous studies have estimated that they are responsible for over 70% of natural variation in flowering time in *Arabidopsis* (Lempe et al. 2005; Shindo et al. 2005). However, under shade conditions, the effects of *FRI* and *FLC* are bypassed and flowering is accelerated (Wollenberg et al. 2008). Because of the association of *FRI* and *FLC* with accelerated flowering in shade,

as well as the correlations of flowering time with plant size and inflorescence height (Mitchell-Olds 1996; Gnan et al. 2017), it is not surprising that we detect loci that overlap with *FRI* and *FLC* for our traits since our populations carry functionally distinct alleles of both genes (Werner et al. 2005; Simon et al. 2008; Lovell et al. 2013).

However, this logic suggests that *SAR4_1* and *SAR5_1* should only affect the later developmental traits indirectly through its effects on bolting time. But this is not supported by our results. *SAR4_1* and *SAR5_1* have direct effects on rosette biomass and inflorescence growth in some populations, even after correcting for flowering time (Figure 2.4), indicating that *FRI* and *FLC* directly influence variation in other traits besides flowering. Consistent with these results, Deng et al. (2011) showed that *FLC* binds to genes that regulate vegetative development (e.g. *SPL15* and *SPL3*) in addition to genes involved in the floral transition and floral patterning pathways (Deng et al. 2011). Similarly, allelic variation in *FRI* has pleiotropic effects on growth rate, flowering time, and water-use efficiency (McKay et al. 2003, 2008; Lovell et al. 2013). However, another possibility is that the effect of bolting time on later developmental traits is not entirely linear, and our path analysis only accounts for the linear relationship between traits.

Future work

The SAR is a widely studied example of plant plasticity, and has important implications in plant breeding and agriculture due to its negative effects on yield. Using natural variation to identify important loci for the SAR can help identify genes that both improve our understanding of the mechanisms underlying the SAR and may attenuate the SAR to improve yield in crops. Our study provides insight into the genetic architecture of the SAR in adult plants, and found

at least one novel SAR locus. The loci that we describe represent opportunities for future fine-mapping studies to identify new casual variants. Furthermore, several of the previously identified genes located within our other SAR QTL have not been implicated in the natural variation of the SAR and may be worth further study.

Our path analysis results also show a complex, temporal element to the underlying genetic architecture of the SAR, where QTL directly affect earlier – but not later – developmental traits. For instance, *SAR4_2* had direct effects on the shade responses of BD and RB but not IG and IB. An intriguing future direction would be to investigate the temporal dynamics of the SAR development in mature plants. Shade effects on hypocotyl elongation in response to shade are detectable within hours (Cole et al. 2011). Our traits were measured over days or weeks so we could not measure short time-scale effects. However, the SAR in adult plants may be amenable to high-throughput phenotyping studies, which could capture genetic changes at hourly (or even finer) time-scales. Numerous studies have used imaging pipelines and time-series data to capture the genetic architecture of plant growth (Zhang et al. 2017; Knoch et al. 2020), and studies that leverage the same technology to study the genetic architecture of plant plasticity over time are emerging (Honsdorf et al. 2014; Marchadier et al. 2019). The SAR can thus serve as a system for future high-throughput phenotyping studies to expand our understanding of natural variation in a plastic and adaptive trait throughout time.

2.6 Acknowledgments

JT and DER were funded by NSF grants IOS-1546719 and IOS-1558090. JT is funded by a Plant Sciences Departmental GSR at UC Davis. DER is supported by United States Depart-

ment of Agriculture (USDA), National Institute of Food and Agriculture (NIFA), Hatch project 1010469. This research was also supported by NSF grant IOS-0923752 to CW and JNM and NSF grant IOS-1444571 to CW, MTB, and JNM. JNM was funded by the U.S. Department of Agriculture NIFA project (CA-D-PLB-7226-H). We would additionally like to thank Cody Markelz and Kazu Nozue for help with planting.

Associated Publication

Chapter 2 was published November 1, 2020:

Ta, J., Palmer, C., Brock, M., Rubin, M., Weinig, C., Maloof, J., & Runcie, D. (2020). Multiple loci control variation in plasticity to foliar shade throughout development in Arabidopsis thaliana. G3: Genes, Genomes, Genetics, 10(11), 4103-4114.

Chapter 3

SplineRGR: A flexible modeling strategy for estimating relative growth rates

3.1 Abstract

High-throughput phenotyping technologies that record finely-sampled longitudinal measurements of plant size through time hold great promise for advancing our understanding of plant growth. However, optimally leveraging these data in analyses of growth rate variation remains a challenge. We have developed a method that we call `SplineRGR`, for Spline modeling of Relative Growth Rates. This method combines the interpretable nature of mechanistic models with the flexibility of non-parametric models. `SplineRGR` can capture key features of plant growth more parsimoniously and flexibly than other methods. We demonstrate the advantages of `SplineRGR` in QTL mapping applications with dense longitudinal data. In simulated data, `SplineRGR` can show up to a 30% higher true-positive to false-positive rate than other methods. In real data collected from *Arabidopsis*, we also detect a QTL for growth rate on chromosome

2 that was not detected by any other method. `SplineRGR` is a powerful, intuitive, and robust tool for modeling variation in plant growth, and can detect types of variation in growth rates that other parametric models cannot.

3.2 Introduction

Plant growth is a dynamic process that involves cell division and cell expansion, the development and expansion of different tissues (stems and leaves), and changes in the relative positions of organs in 3D space. Mathematical models have been widely used to simulate and describe the fundamental aspects of plant growth. Such plant growth models range from simple parametric models like the logistic model, which broadly captures the common phenomenon that plant growth rates start slow when plants are small, accelerate as they grow, but then slow again as plants reach their maximum size, to complex crop growth models (CGMs) that integrate multiple interconnected modules to simulate growth and development of individual organs, phenology, and interactions with the environment. For example, the CSM-IXIM model is a CGM for maize that integrates information such as intercepted radiation, and can output measures of leaf area (Lizaso et al. 2011).

A fundamental unifying principle underlying most plant growth models is the idea that plant growth rates are roughly proportional to plant size. This principle accurately expresses how plants grow by capturing light through their leaves and use this light to generate carbon to grow. Consequently, larger plants capture more light and therefore grow faster. CGMs elaborate on this idea in more detail, additionally modeling changes in carbon allocation among different organs and in different developmental stages, and therefore can capture more of the growth

rate dynamics of a plant.

With the advent of phenomics, high-throughput phenotyping has emerged as a pre-eminent tool for measuring plant development, and provides more data for complicated models. For example, image capture of plant development can result in thousands of data points per plant, or gigabytes to terabytes of high-dimensional data per study (An et al. 2016; Rutkoski et al. 2016; An et al. 2017; Diaz-Garcia et al. 2018). Automated image capture by autonomous unmanned vehicles (e.g. drones) and other platforms can measure plants non-destructively, permitting repeated measurements of the same plant over time. This can provide an unprecedented amount of data for parameterizing plant growth models. For instance, high-throughput phenotyping platforms involving cameras have been able to measure genetic effects on size throughout time in wheat, *Arabidopsis*, and *Brassica napus* (Marchadier et al. 2019; Knoch et al. 2020; Lyra et al. 2020). However, CGMs require more than densely sampled longitudinal data; parameterizing them requires data that can only be obtained through destructive sampling of plant organs which still must be done manually. Therefore efforts to link CGMs to high-throughput phenotyping data remain limited.

Regardless of their complexity, plant growth models are built from simple units with defined parametric forms. These parametric functions can capture general patterns, but their predetermined nature can lead to a degree of inflexibility. Exponential growth, for instance, can be described by the differential equation:

$$\frac{dy}{dt} = ry, \quad y(0) = A \tag{3.1}$$

where y is size, A is the initial size, r is the continuous growth rate, and t is time.

CERES-wheat, a CGM for wheat, is comprised of multiple modules which integrate envi-

ronmental information with plant developmental and physiological states (Ritchie 1985). However, each module is composed of a simple parametric function. For example, one module relates leaf area index (LAI) to the fraction of above-canopy radiation intercepted (IPAR) and photosynthetically active radiation (PAR) with the following equation:

$$\frac{IPAR}{PAR} = 1 - e^{-0.85 * LAI} \tag{3.2}$$

Equation 3.2 is a function that relates plant size to a plant's ability to interact with the environment. While powerful in its simplicity, it specifies a rigid and fixed relationship with LAI that may not fully capture this relationship.

One issue with both Equations 3.1 and 3.2 is that while they capture the general form of traits related to size or physiology, they cannot capture cases that deviate from their pre-established form. With the availability of denser data due to HTP technologies, we might find that the relationships in our data (growth rates as a function of time, or $\frac{IPAR}{PAR}$ as a function of LAI) are not as simple or consistent in form as predicted by these these parametric models. Therefore, allowing more flexibility and data-driven dynamics may improve the quality of plant growth models. Our proposal here is to use regression splines to parsimoniously, but flexibly, model the functional forms of the modules that compose plant growth models.

In this paper, we present `SplineRGR` (Spline modeling of Relative Growth Rates), a modeling framework that fits differential equations using splines. This results in a flexible modeling framework that also has interpretable parameters, i.e. relative growth rates at various plant sizes. We apply this framework to the simplest case, a parametric model of growth like the exponential or logistic model. In future cases however, we would like to apply this methodology to improve the individual modules that make up CGMs.

SplineRGR can capture the growth dynamics of more complicated growth curves that deviate from the standard models, and only requires data on size throughout time unlike CGMs. Additionally, compared to other non-parametric models, SplineRGR describes effects on relative growth rate (RGR) rather than on size, which can be more parsimonious.

In our study, we focus on quantitative trait loci (QTL) mapping as a case study of the applicability of SplineRGR. However, our results are universal to any treatment effect, including differences in genotype, disease severity, nutrient availability, environmental differences, etc.

In our simulations, we show that when there are QTL that cause changes in RGR during plant development, using SplineRGR to estimate growth rate functions can have higher power to detect QTL than points along the size vs. time curves. We also compare different statistical methods in their ability to capture the genetic architecture of plant size in three mapping populations with publicly available data - including an independently generated data set consisting of *Arabidopsis* recombinant inbred lines (RIL). Our results highlight the importance of using the right model for the data set, as a less accurate model can have lower power and can lead to an imprecise interpretation of QTL effects. SplineRGR can also lead to novel results due to its use of RGR rather than size, which might not be possible with parametric methods.

3.3 Materials and Methods

SplineRGR: RGR and SGR

In SplineRGR, we approximate the standardized relative growth rate (SGR) with smooth functions. SGR is the instantaneous relative growth rate (RGR) expressed as a function of current

size (i.e. RGR(y), where y is current size).

SGR provides several advantages over other growth and size metrics. Compared to the absolute growth rate, using SGR can allow for comparisons of RGR at a similar size, even if the plants are at different points of development. This leads to a more standardized comparison. Compared to size as a function of time, the differential equation for SGR is based on RGR as a function of size. This is more conceptually relevant because plant growth depends on the number and size of its shoots and roots, rather than depending strictly on time. Furthermore, SGR is not affected by initial size. Differences in initial size that result in size or growth differences in later development can obscure relationships between growth and physiology (Turnbull et al. 2008; Rose et al. 2009; Rees et al. 2010; Pommerening and Muszta 2015). This can result in spurious correlations between traits (Turnbull et al. 2008; Paul-Victor et al. 2010).

SplineRGR: Overview

We model the instantaneous rate of growth of a plant ($\frac{dy}{dt}$) as proportional to its current size ($\frac{dy}{dt} = f(y) * y$), with the coefficient of proportionality described by the function f() which itself is a function of the current size. f(y) is the instantaneous RGR at size y, and we allow f(y) to change with size. Rather than enforcing a specific form for f(y), we allow a flexible form using penalized B-splines which constrains changes in f(y) to occur smoothly as a plant grows, as shown:

$$\frac{1}{y}\frac{dy}{dt} = \Sigma b_j(y)\beta_j$$

$$y(0) = y_0$$

Here, $b_j(y)$ are B-spline basis functions and β_j are the unknown parameters. Because the equations cannot be solved analytically, we use numerical integration to solve for $y(t)$, and numerical optimization methods to find optimal values for $\hat{\beta}_j$. We choose $\hat{\beta}_j$ by maximizing an objective function that sums the squared differences between the logarithm-transformed predicted values and the observed data, and adds a penalty on the squared differences between consecutive β_j to control the smoothness of the spline. We use the logarithm transform for the input size data in the objective function because the variance of the data increases as time progresses and this tends to equalize the dispersion along the growth curves.

We implemented `SplineRGR` using the R programming language (R Core Team 2021). Differential equations are solved with the *deSolve* package, and spline basis functions are calculated using functions from the *splines* package (Soetaert et al. 2010). We used a total of three internal knots and an intercept for our analyses, unless otherwise stated. We used non-uniform knot spacing, placing the internal knots closer to the left (smaller) boundary knot than the right boundary knot because most data points are collected on smaller plants. Non-uniform knots have also been shown to be optimal in certain conditions (Likhachev 2021), and this choice resulted in fewer convergence issues in our applications.

Initial values

We selected initial values for the β_j parameters and initial size by fitting exponential or logistic growth models to the data, and calculating RGR values at the knot locations based on the analytical forms of these models. Because the RGR differential equation for an exponential model is constant ($\frac{1}{y}\frac{dy}{dt} = r$, where r is the continuous growth rate), the values for the RGR

parameters are the same (i.e. parameter #1 = parameter #2 = parameter #3, etc.). Consequently, we fitted an exponential model to the data to estimate the continuous growth rate (r), and then used r as the initial value for the RGR parameters. Similarly, we also fitted a logistic model to estimate initial RGR values for plants with logistic growth. In this case, the relationship between RGR and size follows a line with a negative slope, which we calculated to estimate initial values.

Simple simulations for intuitive understanding of SplineRGR

We used simulations to study the ability of `SplineRGR` to capture a wide range of plant growth patterns. These simulations included: 1) testing the ability of `SplineRGR` to accurately estimate RGR from a variety of growth curves, 2) evaluating how effects on initial size can affect size through time and RGR, 3) comparing different timings of RGR effects, and 4) examining how smooth fluctuations in RGR can impact the size vs. time and RGR vs. size curves. These simulations are biologically motivated, and we related effects on RGR to a variety of biological mechanisms: different timings for the reproductive switch, alterations in carbon allocation, changes in photosynthetic efficiency, etc. While we cannot pinpoint the exact mechanisms that affect RGR, `SplineRGR` can estimate deviations from a smooth growth/RGR curve due to any biological mechanism, resulting in a more flexible model than other parametric models.

For simulation set 1), we simulated three different growth curves: exponential, Gompertz, and logistic to represent the different types of growth exhibited by plants. The relative growth rate is constant for exponential growth, but Gompertz and logistic models exhibit decreasing relative growth rates, with asymmetric and symmetric decreases, respectively. The equation

for the exponential growth curve is shown in Equation 3.1.

The equation for the Gompertz growth curve is:

$$y = Ae^{-be^{-rt}} \tag{3.3}$$

where y is size, A is the asymptote, e is Euler's number, b is the displacement along the x-axis, r is the continuous growth rate, and t is time.

Lastly, the equation for the logistic growth curve is:

$$y = \frac{A}{1 + be^{-rt}} \tag{3.4}$$

where y is size, A is the asymptote, $\frac{A}{1+b}$ is the value of y when t = 0, e is Euler's number, r is the continuous growth rate, and t is time.

Each growth curve was simulated for 15 equal time steps. All growth curves were simulated between times 0 and 30, except the exponential growth curve (simulated between times 0 and 20). We used the RGR differential equation for each growth curve to calculate the analytical RGR. We then compared the estimated RGR (using `SplineRGR`) to the analytical RGR.

When simulating the growth curves for simulation set 2), we used only exponential growth curves (Equation 3.1). We generated three growth curves with different initial sizes (0.125 - 0.5) but with the same growth rate (0.5). Each growth curve was simulated for 100 equal time steps between times 0 and 20. We used `SplineRGR` to estimate the RGR of these growth curves to compare the RGR vs. size and size vs. time curves.

We compared the timing of RGR effects in simulation set 3) and how they affect plant growth. We simulated two sets of plants with exponential growth curves, where one plant had a higher RGR value around a particular size than the other plant. We used RGR parameters

to simulate the size vs. time curves rather than Equation 3.1. We used a value of 0.3 for all RGR parameters for the first plant, and used the same value except around one size for the second plant. Around this particular size, we used a larger RGR value of 0.6. We used 7 RGR parameters in total, with the second plant having a higher RGR value (0.6) at the second, third, and fourth parameters in the three different simulation sets. These RGR effects occur when the plant is small, medium, and large, respectively. We used these simulations to understand how RGR effect timings influence the overall growth curve.

Finally, we manipulated RGR vs. size curves to analyze downstream effects in the size vs. time curves in simulation set 4). We wanted to see how the degree of fluctuations (i.e. waviness) in the RGR vs. size curves could influence variation in the time vs. size curves. Like in simulation set 3), we simulated exponentially growing plants using RGR parameters. We simulated smooth fluctuations in RGR by controlling RGR parameter autocorrelation in two cases: $\sim$0% correlation and $\sim$75% correlation. The differences in autocorrelation results in variation in the RGR vs. size curves between plants. To induce autocorrelation, we included independent and shared sources of variation (based on a normal distribution) across the RGR parameters. We scaled both sources accordingly in order to adjust the correlation to 0% and 75%. We then visually examined how fluctuations in the RGR vs. size curves can lead to higher variance in the size vs. time curves. When we refer to RGR parameter correlation, we refer to the average correlation; we realize that the random draws from a normal distribution can result in slight deviations in the observed correlations.

Ground-truth of SplineRGR

To show that `SplineRGR` produces unbiased estimates of RGR parameter values, we simulated growth curves with known parameter values and compared them to `SplineRGR` estimated values. These simulated growth curves were based on Equation 3.1, as this was the simplest case and because the real data sets used in this study consisted mainly of plants with approximately exponential growth (Figure A.14).

For the initial size, we used random numbers drawn from a normal distribution (μ = 0 and σ^2 = 1). To ensure that the values were positive and non-zero, we took the absolute value of the random number and added a small value (0.1). Similarly, for the growth rate (r), we used random numbers sampled from a normal distribution (μ = 0 and σ^2 = 0.25). We also used the absolute value of r as well. We then took these parameters and simulated growth curves, and added measurement error by multiplying size by Euler's number (e) raised to power of a random number. This was to make sure that size remained positive. The random number was drawn from a normal distribution (μ = 0) with differing variances (σ^2 = 0.01 - 0.25) to simulate distinct amounts of measurement error. We then estimated RGR parameters using `SplineRGR`, and then visually compared plots of known and estimated parameter values. We determined that `SplineRGR` performed well if the estimated values were correlated with the true values, and if the scatter of the estimated values were spread evenly around the true values.

QTL mapping simulation setup

We simulated mapping populations to compare QTL mapping power across a variety of methods. We simulated 100 backcross populations - each consisting of 100 individuals - with known QTL positions and effects. We used the R package *R/qtl* to simulate the backcross populations and the corresponding genetic map (Broman et al. 2003), and generate phenotypes - RGR parameters - for these individuals. We used a mean growth rate (r) of 0.5 (Equation 3.1) to simulate the RGR parameters, with the same initial size (A = 1) among individuals.

We simulated mapping populations with two different genetic architectures: 1) a single QTL affecting RGR and 2) multiple (two) QTL affecting RGR. We used a total of 7 RGR parameters, ordered by the sizes they influence (from least to greatest). For the one-QTL simulation, we simulated a QTL on the second RGR parameter, as that generated the largest differences between plants through time (Figure 3.3). For the two-QTL simulations, we simulated independent QTL on the second and third RGR parameters. We found that the same QTL effect on the second parameter resulted in a larger effect in size than a QTL on the third parameter (Figure 3.3). To make the QTL effects more comparable, we calculated the t-statistic for the QTL effect on the final time point for the second and third parameter separately, and then adjusted the the effect on the third parameter until the t-statistics were similar for each parameter. We determined that the t-statistic on the final time point was about 1.6 times greater when the QTL was on the second RGR parameter than the third parameter, so we multiplied the QTL effect on the third RGR parameter by 1.6. Similar to the simulations in 4) above, we induced autocorrelation in the RGR parameters in our mapping populations ranging from 0% to 75% correlation.

We simulated the growth curves using the RGR parameters, and then mapped QTL using various methods: a single-time point (STP) method, a parametric method, the functional principal component (FPCA) method, and the `SplineRGR` method. For the STP analysis, we used the last time point because that is when we find the largest difference between genotypes. For the parametric method, we estimated the initial size (A_0) and growth rate (r) using Equation 3.1 and then mapped the growth rate. For the FPCA method we used the *funqtl* package to calculate an average LOD (SLOD/SL) across principal components (Kwak et al. 2014, 2016). For the `SplineRGR` method, we estimated RGR parameters, estimated principal components from the parameters, and then used the principal components that explain the most variation ($\sim$ 99%) along with the average LOD method from the *funqtl* package. We combined principal components with `SplineRGR` and the average LOD method because we wanted to simulate having no information on which RGR parameter would be best for mapping. Missing data were replaced by the median values for a particular principal component or time point.

For the QTL mapping analysis, we mapped QTL with the method in Kwak et al. (2016) and Haley-Knott regression using *R/qtl* (scanone function) and *R/qtl2* (scan1 function), respectively (Broman et al. 2003; Kwak et al. 2016; Broman et al. 2019). We calculated a 95% LOD threshold for each scan using 1000 permutations, and estimated 95% credible intervals for any significant QTL. If the true QTL is within the credible interval, then that method has successfully identified the QTL for that simulation. We then calculated QTL mapping power as the number of times the true QTL was detected divided by the number of simulations (i.e. 100).

We also expanded our analyses by simulating plants with logistic growth curves, using parameters used in previous publications (Kwak et al. 2014) (Equation 3.4). We simulated one-QTL mapping populations using the logistic model, and used the logistic model as our

parametric model in subsequent analyses.

Measuring AUC

Using QTL mapping power alone does not tell us how much power changes as the false-positive rate changes. Consequently, we also estimated receiver operating curves (ROC), which measure the true positive rate as a function of the false positive rate. We used the *ROCR* package in R to estimate ROC and area under the curve (AUC) for each method across each simulation (Sing et al. 2005). Higher AUCs suggest a better performing method (i.e. a high true-to-false positive ratio). We denoted true positives as markers within 20 centiMorgans (cM) of the true QTL marker. We also calculated the average AUC and confidence intervals for the AUC across simulations.

Introduction to different data sets

To complement our simulations, we applied `SplineRGR` to a variety of real data sets. Data Set #1 is a data set that we collected, with $\sim$150 RILs of a Col-0 x Sha mapping population (Simon et al. 2008). Briefly, 6 replicates of each RIL were grown in Conviron growth chambers, and rosette area over time was measured using a system of Canon Powershot S95 cameras that took pictures of plants every hour for $\sim$2-3 weeks. Due to technical issues, we do not have complete hourly or daily data for each plant but around $\sim$80 - 280 data points/pictures per plant. Available pictures were corrected for color and optical distortions, and then merged into orthophotos (An et al. 2016). The orthophotos were then segmented into binary images using the Normalized Green-Red Difference Index (NGRDI), and then the number of pixels

associated with the plant was measured to estimate rosette area. Daily area averages were then used for analysis and mapping. Data Set #2 comprises of a wheat RIL population, with 197 lines derived from the Chinese Spring and Paragon parents (Lyra et al. 2020). The lines were grown in the field, and the Field Scanalyzer phenotyping platform was used to measure canopy height for a total of 22 time point measurements for each plot. Four Arabidopsis RIL populations are used in Data Set #3, which include Bur x Col, Cvi x Col, Bla x Col, and Yo x Col populations (Marchadier et al. 2019). These RILs were grown in well-watered and water deficit conditions, and pictures were taken by Phenoscope, a custom high-throughput phenotyping system. One picture of each plant was taken every day for $\sim$20 days, and projected rosette area was measured by the number of rosette pixels. We mapped QTL using the four different methods explained beforehand, and compared the methods in terms of QTL detected for Data Sets #1 and #2. For Data Set #3, we compared heritability measurements between the individual time points and the estimated RGR parameters.

For Data Set #1, genotype data was provided by Brock et al. (2020) (Brock et al. 2020). For Data Set #2, genotype and phenotype data is publicly available. For Data Set #3, only the phenotype data is publicly available and no genotype data is available. Consequently, we only measured heritability in Data Set #3. For the choice of lambda (penalty of smoothness) in these data sets, we subsetted 10% of the plants and then use 5-fold cross validation with 80% of the time points as the training data. The lambda that resulted in the lowest error (average log difference between observed and predicted) was used for the entire data set. Due to computation time, we did not use a dense grid of lambdas but instead used lambda values ranging from 0.01 to 100 in multiples of 10.

Data availability

Scripts and analyses are available at `https://github.com/jkhta/DEGAM`. The QTL simulations and RGR estimation were run on the FARM cluster at UC Davis. Data Set #2 and #3 are available online (only phenotype data is available for Data Set #3) (Marchadier et al. 2019; Lyra et al. 2020). Genotype data for Data Set #1 is available online (Brock et al. 2020).

3.4 Results

Overview of results section

The results section is organized into two parts: Part One) Evaluating the flexibility and power of `SplineRGR` to describe differences in growth between plants, and Part Two) Applications of `SplineRGR` to real data.

We first illustrate how `SplineRGR` can estimate RGR from various simulations, each with relevant biological underpinnings. These sections mirror the simulation descriptions in **Simple simulations for intuitive understanding of** `SplineRGR` (1) - 4)).

Part One

1): Comparison of a variety of growth curves: size vs. time curve vs. RGR

We first show that `SplineRGR` can accurately represent a variety of classic growth curve patterns. These growth curves differ in size as well as RGR (Figure 3.1A vs. 3.1B). Conceptually,

the dynamics of how RGRs change as a plant grows are driven by mechanisms such as variability in available resources and space. A plant with unlimited resources and space, where there is a constant allocation of new carbon towards further increases in size (RGR is constant across development), would exhibit exponential growth (dark red curve). In contrast, when space or resources are limited, plants may exhibit growth patterns that more closely match logistic and Gompertz growth curves (gray and yellow) which feature decreases in RGR as plants get large and begin to deplete their available resources. This will contribute to lower RGR. In all three cases, `SplineRGR` accurately captured both the size vs. time and RGR vs. size curves (dashed vs. solid line).

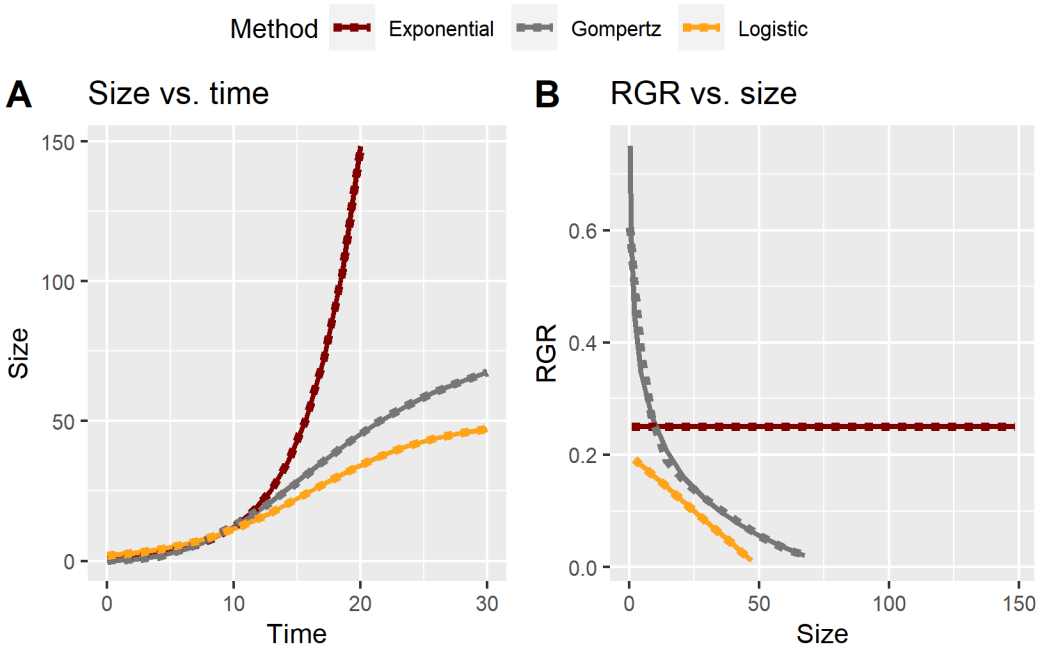

Figure 3.1: Example growth curve fits with `SplineRGR`. A) Example growth curves (exponential, logistic, and Gompertz) of changes in size as a function of time. B) Corresponding RGR vs. size curves. In each panel, the true curves are given solid lines and the approximations using `SplineRGR` are shown with a dashed line.

2): Initial size effects on the size vs. time curve and RGR

Variation in germination time can lead to plants of different sizes and developmental stages at the same time point after planting, even if growth rate dynamics are the same for all plants. To illustrate this point, we show exponential growth curves with the same RGR but with differences in size at time 0 in Figure 3.2A. If we use the size vs. time curves (Figure 3.2A) to test for between-plant differences, we see that differences between plants increase as time increases. Consequently, using later time points to test for plant effects provides higher power (i.e. a larger t-statistic) due to a bigger effect. However, this can disregard the mechanism that led to a difference in final size - an interaction between the different sizes at time 0 and the underlying growth processes.

In Figure 3.2B, we show the time to reach size 10. This is done to capture the differences in germination times between plants, which can also shift the growth curve similar to differences in initial size. However, these mechanisms are different in their interpretation. Size at time 0 can be interpreted as the size of the hypocotyl after germination, while germination time can be interpreted as the time it takes for the hypocotyl to emerge. Having effects on either of these can result in shifts in the growth curve. Overall, we found that plants with higher absolute growth rates at any point in time (dark red) also tended to reach size = 10 at an earlier time point.

Figure 3.2C depicts the RGR vs. size curves for these same three plants. If we use RGR to test for plant effects (Figure 3.2C), we find that RGR is the same across size for all three plants. The only difference between the individuals is the timing at size 10 (Figure 3.2B) when using RGR. In this case, we attribute the differences in size across time only to initial size or

germination time without requiring any differences in the processes controlling growth rate.

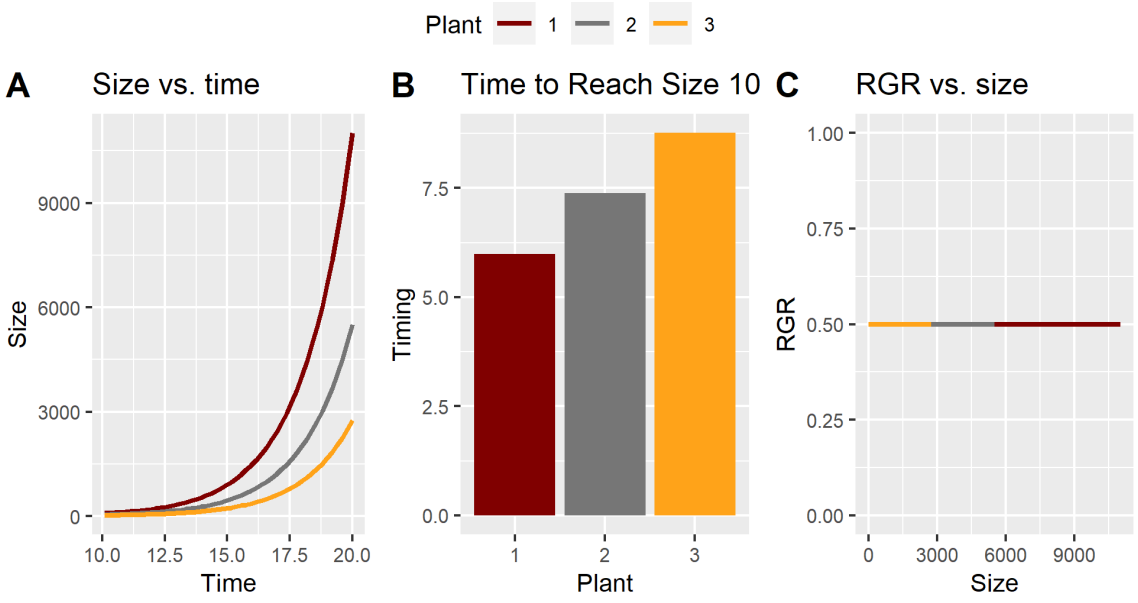

Figure 3.2: Growth curve shifts due to differences in initial size. Panel A shows three exponential growth curves with the same growth rate but shifted in time due to differences in initial size. Panel A is truncated from times 0 - 10 to emphasize the differences between the plants. Panel B shows the time at which each plant reached size 10 (in days). Panel C shows the corresponding RGR vs. size curves for the same three plants. The lines in Panel C are overlapping because the RGR values for all three plants were the same across size, but they ended at different sizes, so we only observed RGR at a reduced range of sizes for plants 2 and 3.

3): The impact of variation in RGR as a function of size

We next perturbed the RGR vs. size curve to understand its effects on the size vs. time curve. To relate this biologically, we can consider these perturbations as changes in photosynthetic efficiency of a plant over development. Photosynthetic efficiency has been shown to change as a plant matures (Bielczynski et al. 2017). A change in photosynthetic efficiency means that the plant will improve or become worse at carbon assimilation through development, which will also result in changes in RGR. In these simulations, we increased the photosynthetic efficiency

of one plant, compared to another plant when the plant is small, medium, or large. In each case, the changes in photosynthetic efficiency are transitory, reverting to the control as the plant grows. In Figure 3.3, when we increased photosynthetic efficiency of a plant when it is small (Figure 3.3A), we found that the individual grows much faster and reaches a larger final size (dark red curve) compared to when the individual has a lower photosynthetic efficiency (gray curve). When we increased the photosynthetic efficiency when a plant is medium-sized, only the later half of the growth curve showed a difference compared to the other individual. Lastly, when we manipulated the photosynthetic efficiency when a plant is large, we saw a minuscule difference in the last part of the growth curve. Overall, we found that differences in RGR among smaller plants have larger effects on final sizes than differences in RGR among larger plants.

4): Effects of RGR parameter correlations on the size vs. time and RGR vs. size curves

We next show how varying correlations on the RGR parameters impacts the size vs. time and RGR vs. size curves in Figure 3.4. The motivation behind these simulations is to show how `SplineRGR` can accommodate fluctuations in the RGR vs. size curve.

In Figure 3.4A and 3.4C, we simulated three plants with 0% and 75% RGR parameter correlations. Generally, we found that as correlation increases the variance of the data across time increases. We also saw the RGR vs. size curves shift accordingly as well. At 0% correlation, the RGR vs. size curves are close to 0.5 (Figure 3.4B), the mean value used in the simulations. For 75% correlation, the RGR vs. size curves have been shifted away

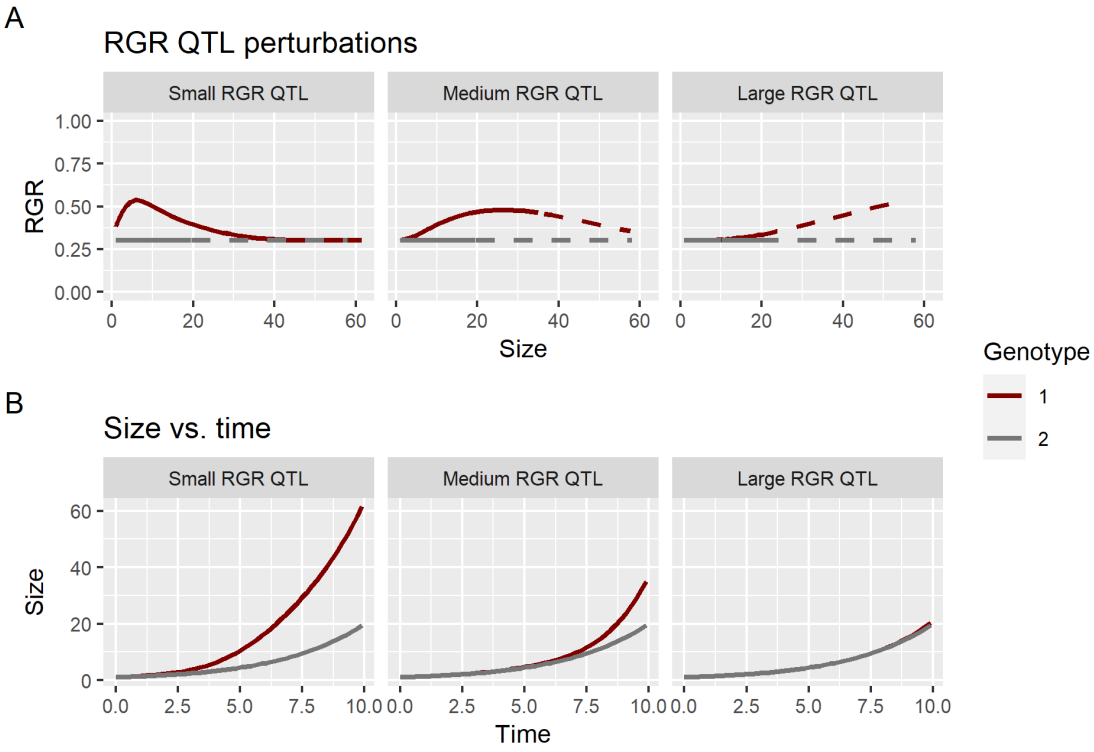

Figure 3.3: Perturbations of RGR vs. size curves have different impacts on growth curves. Panel A illustrates the differences in RGR across size for two different genotypes, with QTL effects on RGR when a plant is small, medium, or large (from left to right). Panel B represents the resulting growth curves when we use the parameters in Panel A for simulating the individuals. The dark red curve and line shows the individual (i.e. genotype) with an increase in RGR relative to the gray curve and line. The solid lines in Panel A show the measured RGR values in the size vs. time curves, while the dashed lines depict the predicted RGR values if the simulations were extended.

from 0.5 (Figure 3.4D). For instance, for individual 2, the entire RGR curve has shifted up to $\sim$0.6 RGR at 75% correlation. Consequently, individual 2 has a larger growth curve at 75% correlation compared to 0% correlation. In contrast, the RGR values have shifted down for individuals 1 and 3, which results in smaller growth curves for both of these individuals at 75% correlation (Figure 3.4C). Generally, we found that as correlation increases, entire RGR curves are affected, usually shifting up or down. This larger variance across the RGR curves resulted

in larger variance in entire growth curves.

These simulations highlight an important aspect of modeling RGR with `SplineRGR`. Because `SplineRGR` is approximated by smooth functions, one hypothesis underlying `SplineRGR` is that RGR should be a smooth function of size rather than having sharp increases or decreases. We demonstrated that `SplineRGR` can reliably estimate smooth shifts in RGR, making `SplineRGR` more flexible than parametric methods that assume a constant line for the RGR vs. size curve.

Bias and accuracy of Spline RGR using simulations

In order to test the performance of `SplineRGR`, we show that `SplineRGR` can accurately estimate RGR in simulations with varying amounts of measurement error in Figure 3.5. Panels A, E, and I show example growth curves with different amounts of measurement error, while the other panels show comparisons of known RGR values to estimated values. Across different measurement error levels (Rows 1-3), the true values appear positively correlated with the estimated values (correlations of parameters in Figure 3.5 > 0.89). The estimated values are evenly spread around the 1:1 line (black line) and appear unbiased in Row 1. As measurement error increases (Row 2 and Row 3), the spread of the estimated values around the 1:1 line also increases. However, the estimated RGR values still appear unbiased. These results suggest that `SplineRGR` can accurately estimate unbiased RGR values from simulated data.

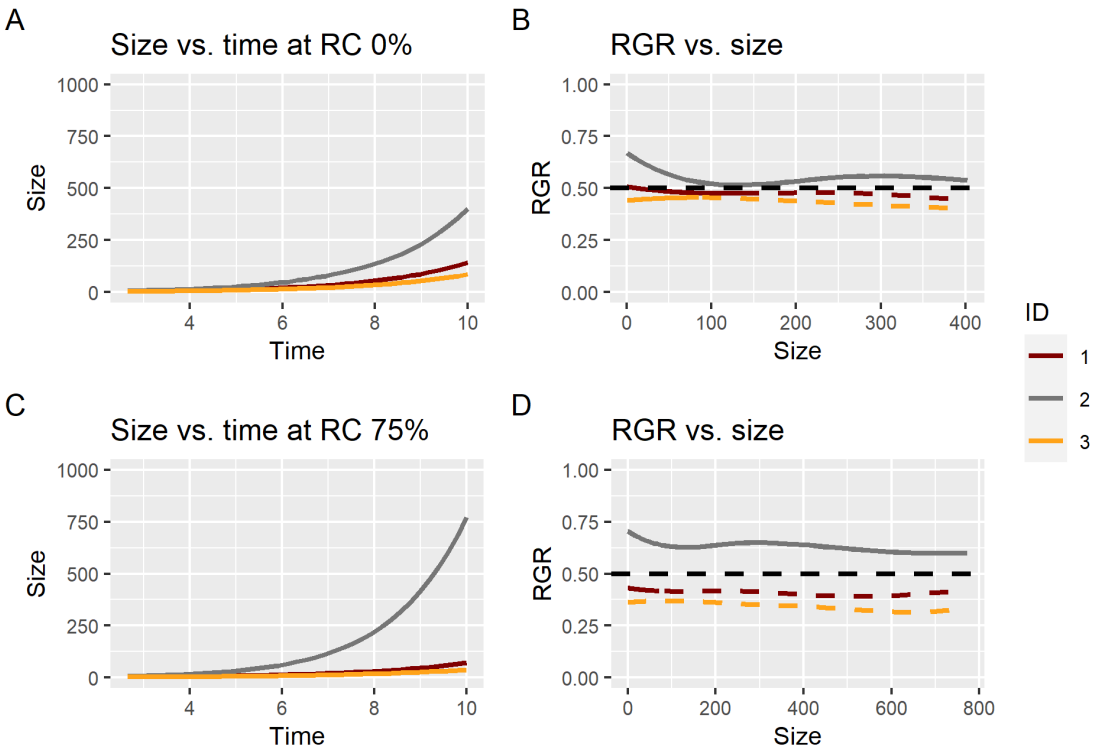

Figure 3.4: RGR parameter correlation effects on growth curves and RGR vs. size curves. We show growth curves (Panels A and C) and RGR vs. size curves (Panels B and D) for two different RGR parameter correlations: 0% and 75% (RC 0% and RC 75%). The growth curves in Panels A and C are truncated and only show the size vs. time curves from time 2.5 to emphasize the differences between the plants. The dashed line at 0.5 in Panels B and D portray the mean RGR value. The solid lines in Panels B and D show the estimated RGR values, while the dashed lines depict the predicted RGR values if the simulations were extended.

QTL mapping power of various methods

Figure A.15 shows QTL mapping power of the different methods in simulated data. In one-QTL simulations, the methods have similar power (∼30% - 50%) for lower RGR parameter correlations (0% and 25%). The STP method has the lowest power, with a max power of ∼40% at 0% correlation. However, as correlation increases (50% or 75%), `SplineRGR` increases in power (up to ∼80% at 75% correlation) while other methods decrease in power (down to

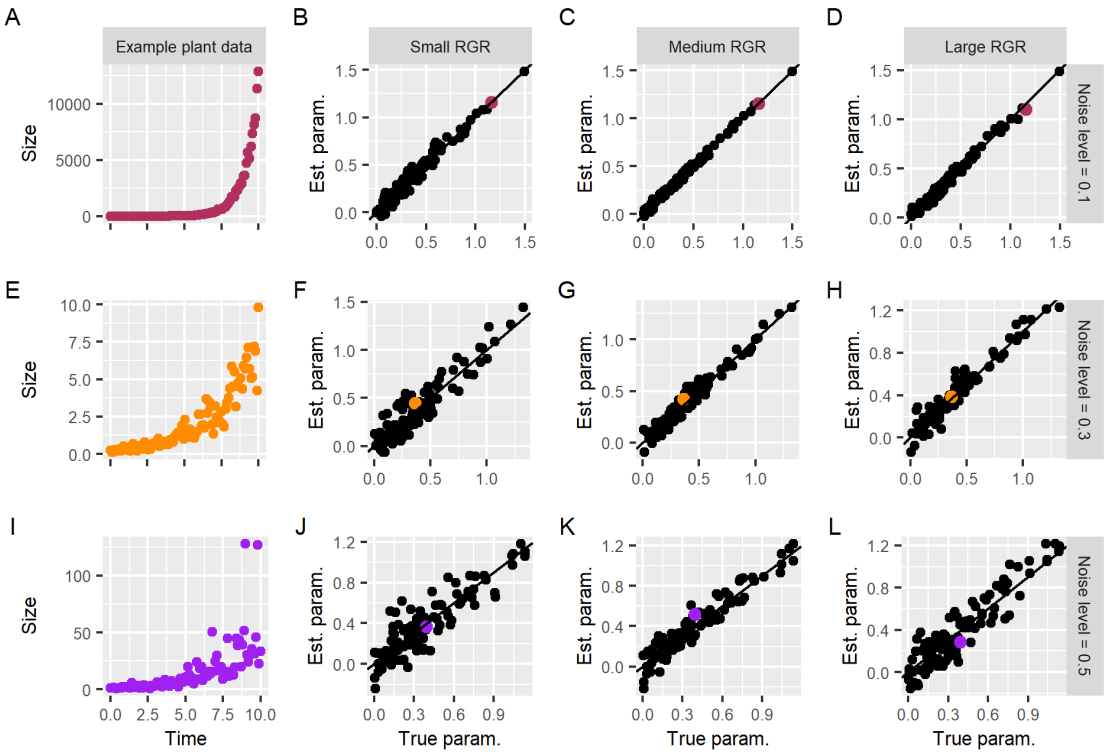

Figure 3.5: Estimated vs. true RGR parameter values from simulations using exponential growth curves. Subsequent rows depict increasing measurement error applied to the time-series data. The first column shows growth curves with differing amounts of measurement error applied, with the least amount of measurement error applied to (A) and the most applied to (I). The latter three columns depict comparisons between the true values of chosen RGR parameters to their estimated values of numerous plants. The first, third, and fifth RGR parameters were chosen (out of 7) to be shown, and correspond to when a plant is small, medium, or large, respectively. The solid lines in the latter three columns is a 1:1 line, starting at the origin. Plants in the first column are colored, and their respective estimated values are shown and colored accordingly in the latter figure columns.

$\sim$10% at 75% correlation). We found similar results for the two-QTL simulations. For the first QTL in the two-QTL simulations, we found similar power among the methods at lower correlation ($\sim$10% - 40% power for 0 - 25% correlation). Again, `SplineRGR` outperforms other methods as correlation increases. `SplineRGR` also has the highest power to detect the second QTL for all correlations in the two-QTL simulations. The gap between methods increases the

fastest for the two-QTL simulations, and SplineRGR caps at 100% power at 75% correlation. Overall, SplineRGR has similar power to detect QTL compared to other methods at lower correlations but has greatly increased power at higher correlations.

These results seem unintuitive: as RGR parameter correlation increases, the RGR vs. size curves should better resemble an exponential function (i.e. a flat line) than at lower RGR parameter correlations. This then suggests that the exponential method should improve in AUC as correlation increases. However, the fact that we do not see this may be explained in part by the differences between genotypes in our one-QTL mapping simulations. In the one-QTL mapping simulations, we can distinguish two genotypes with different RGR vs. size functional forms: the first genotype has an increased value for the second RGR parameter and $\sim$0.5 for the other parameters (i.e. not quite exponential growth). The second genotype has $\sim$0.5 across all RGR parameters giving a flat line in the RGR vs. size space. Therefore the exponential method can model the second genotype well but not the first genotype.

AUC and ROC results

In the one-QTL simulations, SplineRGR performed better than other methods in terms of AUC (Figures 3.6 and A.16). Figure 3.6 depicts AUC results at the highest QTL effect across different levels of RGR parameter correlation. The methods are comparable at lower correlation (0%), but there are larger increases in AUC for SplineRGR as correlation increases. The AUC for SplineRGR increases to $\sim$90% at 75% correlation, while AUC decreases for other methods. We also found similar results in the two-QTL simulations (Figure A.17).

SplineRGR also has higher AUC for most correlations in the logistic simulations (Figure

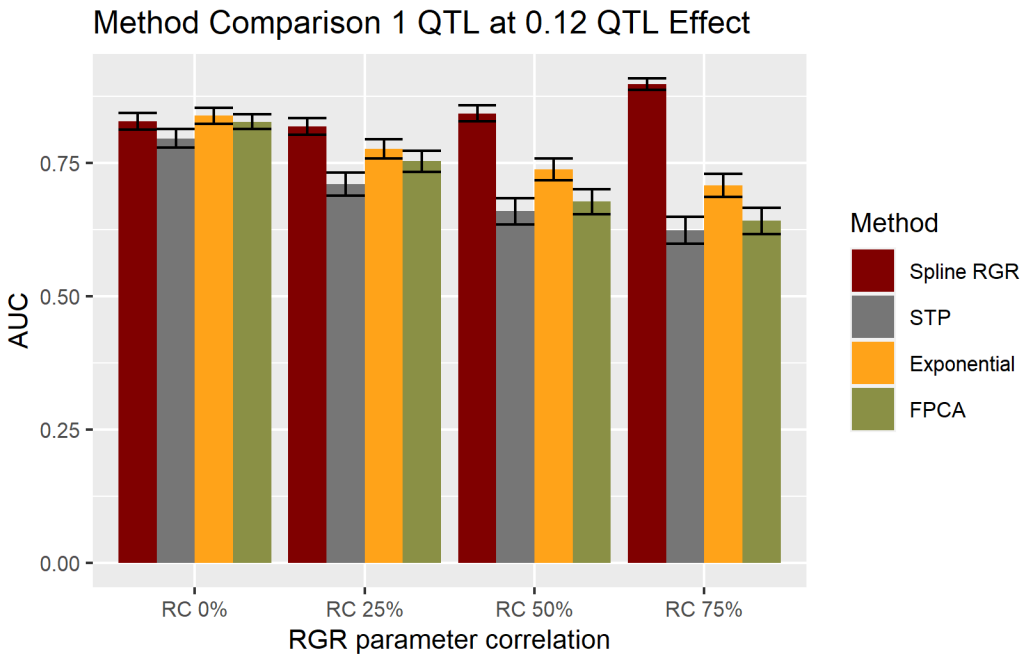

Figure 3.6: AUC at 0.12 QTL effect for the one-QTL simulations (exponential). The mean AUC across the different simulations for QTL effect 0.12 are shown as barplots with error bars. The methods shown are the `SplineRGR` method (`SplineRGR`, dark red), the STP method (STP, gray), the exponential/parametric method (Exponential, yellow), and FPCA method (FPCA, olive). From left to right there is increasing correlation in the RGR parameters, from 0% to 75% (RC 0% - RC 75%).

A.18). AUC increases for most methods as correlation increases, except the parametric method. However, the parametric method has higher AUC than `SplineRGR` when correlations are low (< $\sim$50%). There is a $\sim$0.1 AUC difference between `SplineRGR` and the parametric method when correlation is 0%, and the gap in AUC between the parametric method and `SplineRGR` decreases at 50% correlation. `SplineRGR` has higher AUC than the parametric method above 50% correlation.

Using individual time points may lead to the wrong conclusions

In our simulations, we used the STP method to represent methods that analyze time points individually. In this case, we used the last time point for the STP method because the largest differences between genotype are observed at the last time point (Figure 3.2). However, the last time point might also have the highest variation among plants of the same genotype, at least in exponential growth or when there is no cap to size. In other words, other time points might have higher QTL detection power due to the trade-off between QTL effect size and within-genotype variation. To better understand this, we estimated the t-statistic associated with genotype differences across time. We found that an intermediate time point ($\sim$time 5.0) is associated with the largest t-statistic rather than the last time point (Figure 3.7C).

We next compared the timing of the largest t-statistic to the timing of the largest difference in RGR. Does analyzing time points individually lead to the correct QTL effect timing when genetic variants alter RGR? We compared the timings of the largest RGR effect (Figure 3.7D) to the largest t-statistic (Figure 3.7C). The peak of Figure 3.7C (time $\sim$ 5.0) is later than the peak of Figure 3.7D (time $\sim$ 3.25). This indicates that if we analyze the time points individually, then we would conclude that the QTL occurs later than the true effect. This suggests that using individual time points to detect QTL - when there are genetic variants that affect RGR - can lead to the wrong inference of the timing of QTL effects. This result is not surprising, as effects on RGR should have a delay in effect on the size vs. time curve.

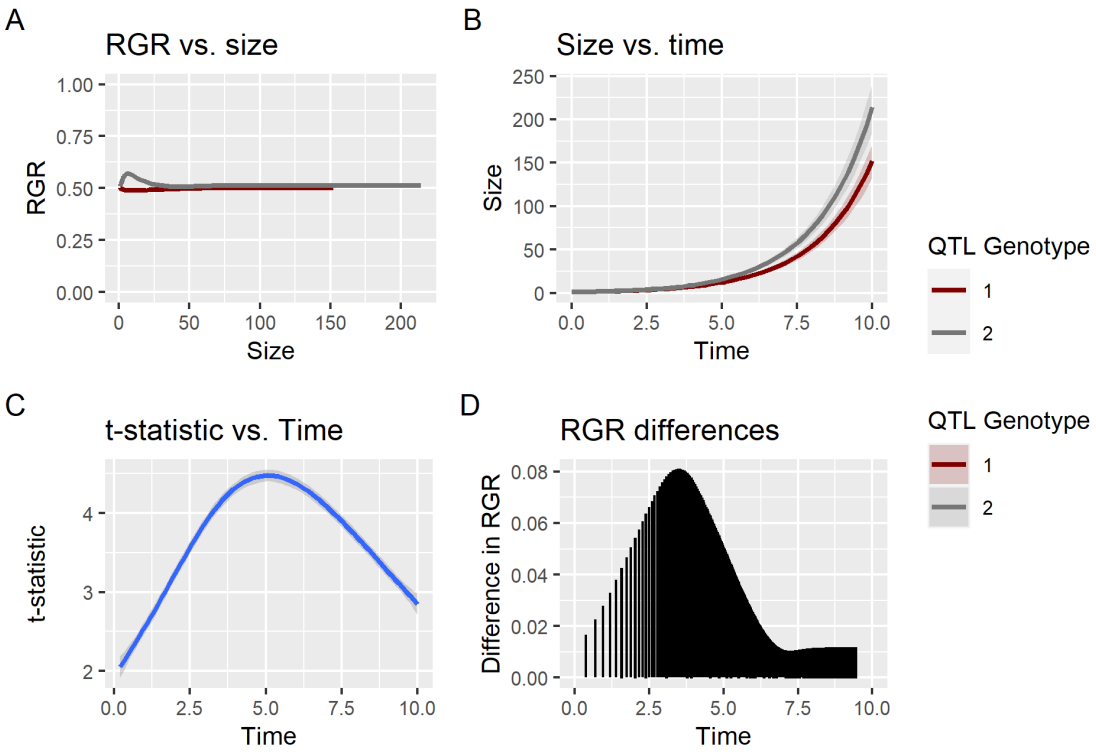

Figure 3.7: Power of STP method to detect QTL across time compared to timings of RGR effects. Panel A shows the RGR values for a one-QTL simulation (QTL on RGR when a plant is small), with the mean RGR across size for a backcross population of 100 for the two different QTL genotypes. Genotype 2 (gray) has higher average RGR when a plant is small compared to Genotype 1 (dark red). Panel B depicts the integrated growth curves using the simulated RGR parameters, with the mean curve and its confidence interval shown in the solid colored lines. In Panel C, we estimated the power of the STP method in terms of the t-statistic across time in our simulations. In Panel D, we show the differences in the RGR values across time between the genotypes, rather than across size.

Part Two

Sample plant pictures from each data set

Figure A.14 shows sample plant data from Data Set #1 - #3. The majority of plant growth curves appear to follow an exponential, rather than logistic, trend. One possible reason for this is that the experiments are not long enough to capture the slowing of growth at later time

points.

Mapping results for Data Set #1 with various methods

In our simulations, `SplineRGR` performed better than other methods especially when RGR parameter correlations are high. To see if `SplineRGR` performs the best in real data, we compared detected QTL among the four methods in Data Set #1 and #2. Figure 3.8 shows the QTL detected among the four methods in Data Set #1. Most methods found similar QTL, with QTL on chromosome 5 being found for all methods. However, `SplineRGR` did not detect a QTL in the middle of chromosome 5 compared to other methods. In contrast, `SplineRGR` detected a QTL on chromosome 2 unlike other methods.

For the two QTL detected using `SplineRGR`, we estimated and compared allelic effects on the size vs. time curve in Figure A.19. We calculated effects on the principal components and the corresponding effects on the RGR parameters. The Sha allele conferred a slight increase in area and RGR compared to the Col-0 allele (Figures A.19A and A.19B).

Mapping results for Data Set #2 with various methods

Similar to results from Data Set #1, multiple methods identified similar QTL for Data Set #2 (Figure 3.9). The parametric and FPCA methods found QTL on chromosome 13 (5A) and 19 (7A). A single QTL on chromosome 13 (5A) was found for the STP method. However, no significant QTL were detected for `SplineRGR`. The FPCA method detected additional QTL on chromosomes 5 (2B) and 14 (5B). The FPCA method seems to be the most powerful method for this data set.

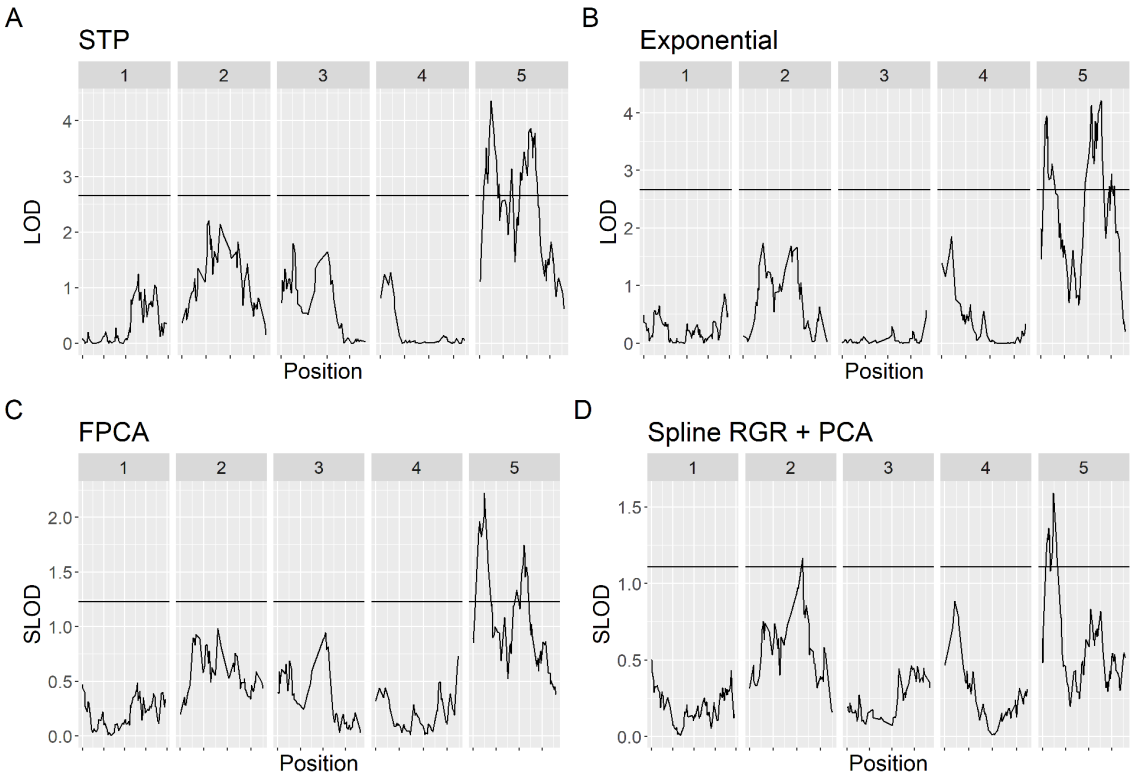

Figure 3.8: Mapping results for Data Set #1. A is showing the STP method using the last time point, B is showing the parametric method (exponential), C is showing the FPCA method, and D is showing `SplineRGR` combined with principal components. The solid horizontal line for each LOD plot depict the 95% LOD threshold using 1000 permutations.

Heritability plots of the `SplineRGR` method across RGR parameters vs. heritability across time points

We compared heritabilities between the individual time points and estimated RGR in Figure 3.10 for Data Set #3. The heritability curves for the individual time points (Figure 3.10A) and the estimated RGR (Figure 3.10B) show different patterns. The heritability stays relatively constant for time points 10 - 25, and then increases in both the well-watered (WW) and water deficit conditions (WD). The maximum heritability is $\sim$0.5 in the WW condition and is $\sim$0.4 in the WD condition. We saw an opposite trend for the estimated RGR. The heritability increases

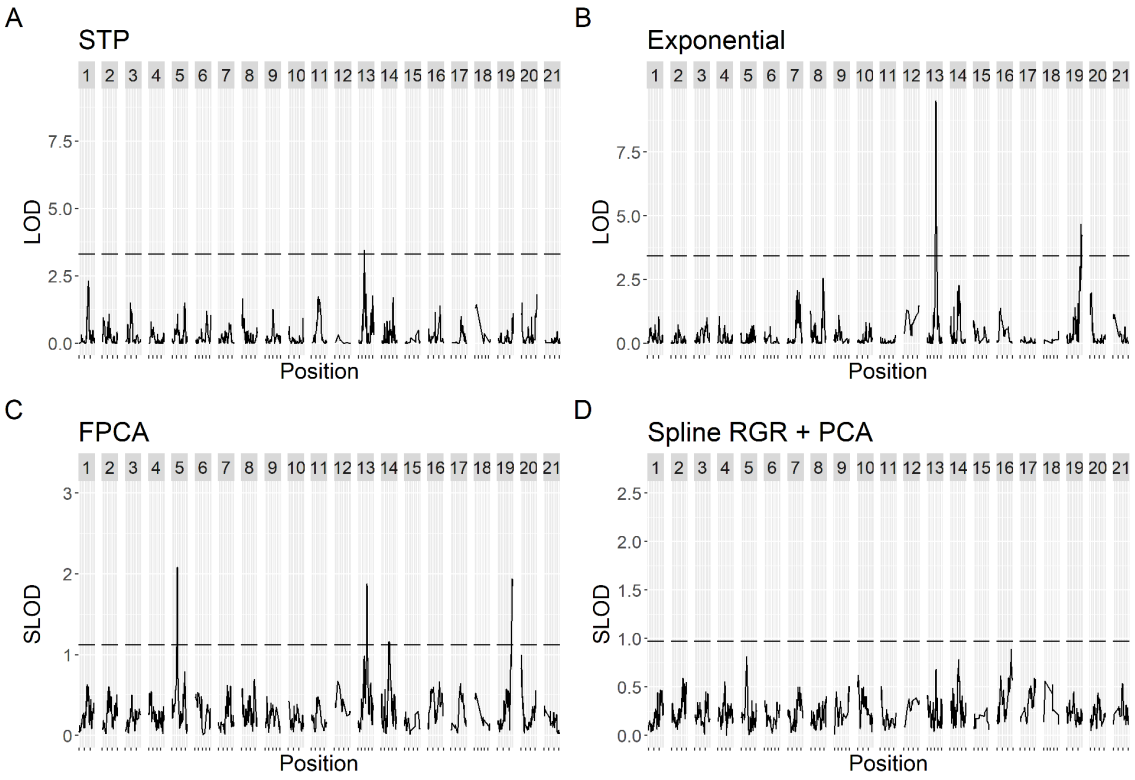

Figure 3.9: Mapping results for Data Set #2. A is showing the STP method using the last time point, B is showing the parametric method (exponential), C is showing the FPCA method, and D is showing `SplineRGR` combined with principal components. The solid horizontal line for each LOD plot depict the 95% LOD threshold using 1000 permutations.

to a maximum early on and then decreases for the later parameters. The maximum RGR heritabilities in the WW and WD conditions are $\sim$0.31 and $\sim$0.29, respectively. These estimates are lower than the maximum heritability for the individual time points. Similar to Figure 3.3, earlier RGR parameters seem to have more importance than later RGR parameters.

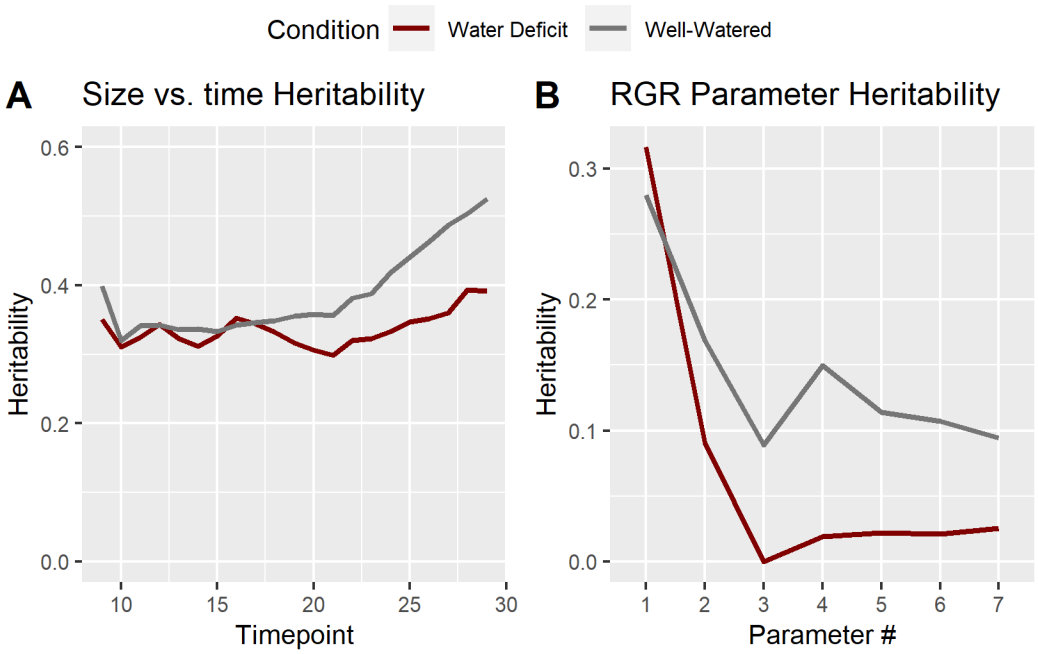

Figure 3.10: Heritability curves for Data Set #3. In Panel A, the x-axis represent the time points, from time point 0 to time point 30; the y-axis shows the heritability estimated from a mixed model. In Panel B, the x-axis shows the parameter number, and the y-axis shows the heritability estimated from a mixed model. The Water Deficit condition is represented by the dark red lines and the Well-Watered condition is shown by the gray lines.

3.5 Discussion

General findings

In this study, we developed `SplineRGR` to estimate relative growth rates (RGR) using differential equations and splines. We compared `SplineRGR` to other methods in simulation conditions and real data. `SplineRGR` has higher power and AUC in numerous situations: across multiple QTL effect sizes, increasing RGR parameter correlations, and with various types of growth curves. In contrast to our simulations, `SplineRGR` did not find the most QTL in Data Sets #1 and #2. This suggests that there are additional factors that contribute to QTL detection, which

we will explore.

Biological relevance of earlier RGR values have larger effects on size

In our simulations, perturbing photosynthetic efficiency when the plant is small has larger effects on size than perturbing when the plant is large. This can be understood through compound interest. With compound interest, the timing of investment and changes in the interest rate can lead to substantial increases or decreases in the principal amount. Our simulations are related to compound interest because we model plant growth proportionally to current size (i.e. an interest rate) and remove the limit to size (exponential growth). Consequently, higher RGRs will result in larger increases in growth, and that increase will have more multiplicative effects the earlier the increase. This will result in larger differences in individual size. Conversely, we found smaller increases in size when effects on RGR occur later because fewer developmental time steps have multiplicative increases in size.

These results might seem less relevant to real organisms, as observed growth curves tend to follow a sigmoidal growth pattern rather than an exponential one (West et al. 2001). However, our results are applicable because the first phase in a sigmoid curve is approximately exponential. Higher initial RGRs can result in larger differences between individuals when growth is still in the (approximately) exponential phase. The difference in earlier RGRs then serves as a mechanism to explain phenotypic variation during these periods of development.

Interpretation of simulation results and RGR

Generally, `SplineRGR` has higher power and AUC than other methods, especially as RGR parameter correlation increases. In the logistic simulations however, the parametric method has higher AUC than `SplineRGR` at lower correlation. The parametric method also has higher power and AUC than the STP and FPCA methods in both exponential and logistic simulations. This contrasts with Kwak et al. (2014) and Kwak et al. (2016) (Kwak et al. 2014, 2016). Kwak et al. (2014) showed that the SLOD method (i.e. average LOD score across time) performed better than the parametric method (logistic) across numerous simulation conditions, and Kwak et al. (2016) showed that the FPCA method tended to perform better than the SLOD method. These results suggest that the FPCA method should perform better than the parametric method, at least in some conditions.

The contrast in results might arise due to the differences in the residual correlation matrix. We induced correlations in the RGR parameters rather than on size as in Kwak et al. (2014) and Kwak et al. (2016). Higher correlations in the RGR parameters can exacerbate variation in the size vs. time curve. This can reduce QTL detection power and AUC for the FPCA method. The types of correlation are also different between studies, as Kwak et al. (2014) and Kwak et al. (2016) explored autocorrelated, equicorrelated, and unstructured residual covariance matrices (Yap et al. 2009). These differences in simulation conditions can result in further distinctions in results between studies.

While `SplineRGR` shows promising results in our simulations, our reasoning might seem circular. Our simulations involved changes in the RGR vs. size curve, which `SplineRGR` is developed to measure. However, our results suggest that using a model that captures

the QTL mechanism can provide higher power. If a QTL affects RGR, then `SplineRGR` can provide higher power than other methods. If a QTL affects photosynthetic rate and carbon assimilation, then models that capture these mechanisms can provide better results. Using a mechanistically relevant model can also accurately decipher the timing of a QTL, as seen in our results (Figure 3.7).

Treatment effects on RGR will have a different interpretation than treatment effects on size. These interpretations are subtle but are more relevant due to their parsimony. Methods that use the size vs. time curve propose that there are treatment effects on size at multiple time points. However, there might be a single treatment effect on a common mechanism that affects multiple time points - i.e. a treatment effect on RGR. Consequently, using RGR as the focal trait in `SplineRGR` is more appealing than other methods due to Occam's razor.

`SplineRGR` tended to perform better with higher RGR parameter correlation, in contrast to other methods. One possible reason for this is the increase in information between RGR parameters as correlation increases. This shared information then increases power and AUC when using RGR combined with PCA. The large amount of variation in the size vs. time curves when RGR parameter correlation is high can explain the decrease in performance in other methods.

Different mechanisms of RGR effects on growth curves

This study shows four different mechanisms that attribute variation in plant growth curves to changes in RGR. These mechanisms can result in discrepancies in power to detect treatment effects. Figure 3.2 depicts how changes in size at time 0 can lead to large differences between

the growth curves. Though RGR values were the same between plants, slight changes in size at time 0 (< 1 area differences) led to multiple fold-changes in area between plants at the final time point (>9000 area for Plant 1 and $\sim$3000 for Plant 3). Another related mechanism is changes in germination time or "timing of" events, which can also shift growth curves. In Figure 3.3, we found direct perturbations in RGR can lead to substantial differences in the size vs. time curves. RGR effects have a bigger influence on size when a plant is small than when a plant is large. Finally, in Figure 3.4, shifts in the entire RGR vs. size curve resulted in higher variance in the size vs. time curves. In total, these mechanisms can be understood again by compound interest, whether through changes in the principal amount (i.e. initial size or "timing of" event), changes in the interest rate (i.e. RGR parameter values), how long the interest rate(s) are applied for (i.e. time), and/or any combination of these mechanisms. Our results reveal the numerous processes that increase variation in size vs. time curves, which can decrease power when modeling treatment effects due to an increase in the residuals term. Developing models that parameterize distinct aspects of plant growth can increase power to detect treatment effects by distinguishing these growth processes.

Interpretation of results for real data - Data Set #1

We detected a QTL on top of chromosome 5 using multiple methods for Data Set #1 (Figure 3.8). Previous studies have identified a similar QTL for vegetative traits in the same population (Col-0 x Sha). Simon et al. (2008) detected a QTL for rosette diameter, and Ta et al. (2020) detected a QTL for rosette biomass using a nested association mapping (NAM) population including the Col-0 x Sha population (Simon et al. 2008; Ta et al. 2020). The resemblance be-

tween traits - rosette growth (this study), rosette biomass, and rosette diameter - suggest that a similar QTL might underlie these traits. Results from Ta et al. (2020) suggest that *Flowering Locus C* (*FLC*) is the causal variant underlying this QTL. *FLC* is a flowering repressor that is involved in the vernalization pathway. While our study analyzes RGRs of rosette area and not flowering time, *FLC* has also been shown to influence vegetative traits (Deng et al. 2011). Consequently, *FLC* is a plausible causal variant underlying the QTL in our study because of its relevant functions and support from past research using the same population.

Interpretation of results for real data - Data Set #2

Our results for Data Set #2 corroborate with past work as well. Lyra et al. (2020) detected similar QTL when using the QTL mapping method from Kwak et al. (2014) and Kwak et al. (2016), with QTL on 2B, 5A, 5B, and 7A (Lyra et al. 2020). We detected the same QTL using the FPCA method. However, we only detected one QTL for the last time point (STP method) whereas Lyra et al. (2020) detected multiple QTL for the last time point (228 DAS). This might be due to the differences in methods. We used the scan1 function from the *R/qtl2* package in the STP method, which performs Haley-Knott regression. In Lyra et al. (2020), interval mapping (IM) and composite interval mapping (CIM) were used with the *R/qtl* package.

Lyra et al. (2020) also detected QTL on 2A, 3A, and 6A when analyzing all of the individual time points. However, these QTL were only detected for a single time point, suggesting that they might be false positives. Our results, results from Lyra et al. (2020), and support from past research (Wu et al. 2010; Griffiths et al. 2012; Bentley et al. 2014; Zanke et al. 2014; Gao et al. 2015; Würschum et al. 2015; Tian et al. 2017; Guan et al. 2018), suggest that the FPCA

method has the highest power for this data set.

Contrast between results for simulations and applications of real data

The contrast in results between the simulations and real data might seem conflicting. In the simulations, `SplineRGR` performed the best while FPCA performed better in certain data sets. However, these results emphasize choosing an appropriate model to match the QTL mechanism. In our simulations, we estimated RGR using `SplineRGR` in cases where we manipulated RGR parameters. In real data, unless we have prior information on the causal variants, we have no idea what the QTL mechanism is. There might be multiple QTL that affect mechanisms other than RGR. The FPCA method might have the highest power in various cases because the FPCA method might weigh different mechanisms more evenly than `SplineRGR`, which focuses on RGR.

However, a key difference that distinguishes the FPCA method from `SplineRGR` is mechanistic interpretation. Because the FPCA method uses the size vs. time curve, it is based on absolute growth as a function of time, not as a function of developmental stage or size. If the environment is constant, and growth rate (absolute or relative) is not a function of size, then nothing in the FPCA model can mechanistically explain how growth rate changes with time. In contrast, `SplineRGR` is based on RGR, which relates plant growth to its photosynthetic capacity and carbon output (i.e. size). While `SplineRGR` might have lower power than FPCA in various situations, this difference in interpretation can make `SplineRGR` more appealing in certain circumstances.

Interpretation of results - Data Set #3

Our estimated heritabilities for Data Set #3 are similar to results from Marchadier et al. (2019) (Marchadier et al. 2019). The heritabilities for the last time point in Marchadier et al. (2019) range between 0.23 - 0.67 for the four different populations in the well-watered (WW) conditions, whereas we estimated an overall heritability of $\sim$0.5. Similarly, Marchadier et al. (2019) estimated heritabilities between 0.12 - 0.55 in the water deficit (WD) conditions, while we estimated an overall heritability of $\sim$ 0.4.

In comparison to the stabler heritabilities of the individual time points, the heritabilities for the RGR parameters are higher for earlier parameters and lower for later parameters. This might be related to how earlier RGR parameters having larger effects on size (Figure 3.3).

Trade-offs and conclusion

In total, while `SplineRGR` can flexibly estimate RGR, there are several drawbacks. The method requires adequate sampling (> 10 time points) due to the number of parameters estimated. There is also numerical instability when the initial conditions are random. This can be remedied by using initial values estimated from a parametric model that captures the growth curve well. In our study we used exponential and logistic models to estimate initial values, but additional models can be added, making `SplineRGR` very flexible and accommodating for multiple types of growth curves.

3.6 Appendix

Description of pipeline for pot and rosette segmentation

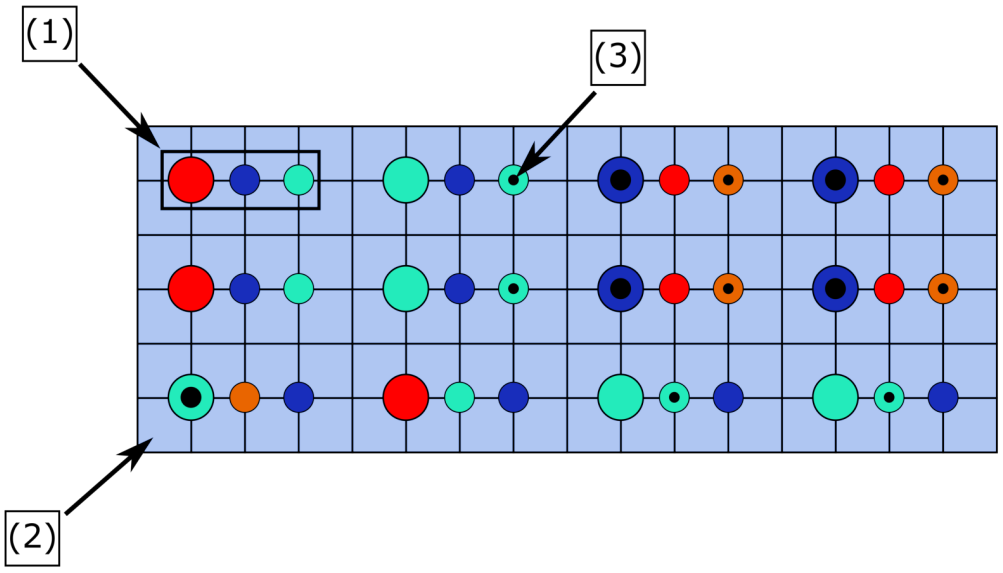

Figure A.11: Diagram depicting the structure of a shelf image file. Each half tray is identified by a dot code identifying it along its horizontal central divider (shown by (1)). In each column of the shelf image there are three half trays. A plant is located at each of the 8 pots within a half tray (shown by (2)). Dot codes are distinguished both by their colors, but also by the presence of a black dot at their center (shown by (3)). The possible dot code colors were pink, dark blue, light blue, red, purple.

To efficiently extract measurements from a high volume of plant images, the OpenCV and PlantCV libraries were utilized to create an automated image processing pipeline (Bradski 2000; Fahlgren et al. 2015a). The plant images utilized in this study consisted of shelf images containing columns of three half trays arranged in a rectangular array. Each half tray was identified by a sequence of colored dots displayed on its central divider. Each half tray contained a set of 8 individual plants arranged in a grid. The position of a particular pot within this grid determined its pot position number. Each dot code was shared by two half trays, such that a

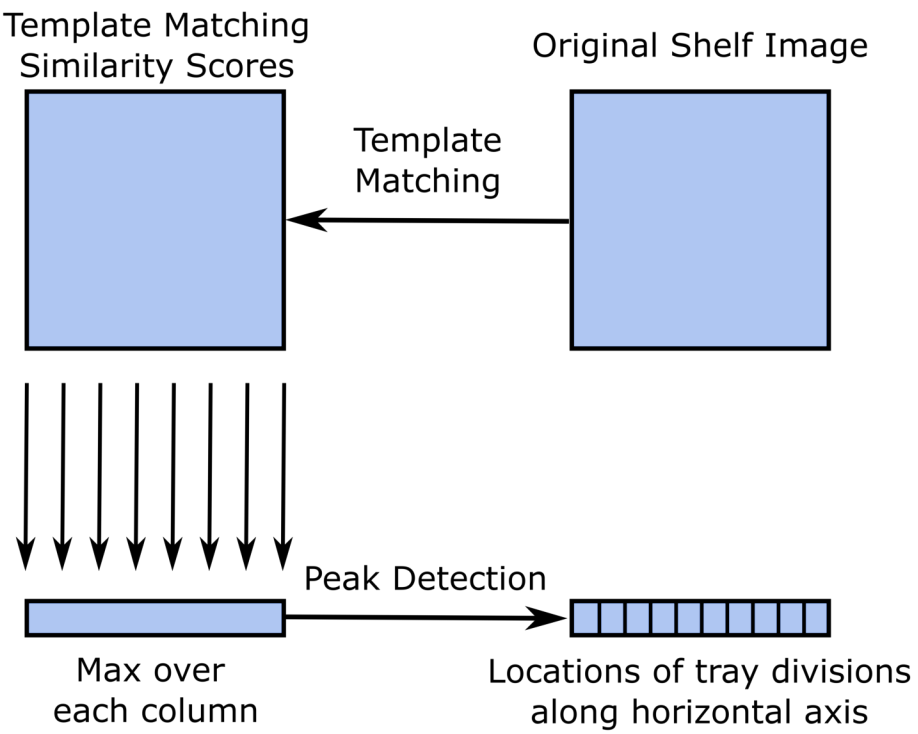

Figure A.12: Finding the horizontal coordinates of the tray divisions by matching a template horizontal divider with the shelf image and finding the peak of similarity along the horizontal axis. A similar approach was used for finding the vertical coordinates of division except the template matching is conducted with a vertical divider template and the max is taken along each row.

given tray contained 16 plants split into two half trays of 8 plants.

The first task was to segment the shelf image files into their constituent half trays. Before this could be done, however, it was necessary to crop the shelf image closely so that there was no blank space surrounding the shelf which could interfere with or offset the segmentation. To automate this cropping procedure, the OpenCV library's contour detection functionality was used to find the contour representing the shelf. The image was then cropped to the bounding rectangle of the shelf contour.

To robustly segment the shelf image into its constituent trays, we utilized template matching

to detect the grey dividers which represented the divisions between the trays. Due to the grid structure of the shelf, the detection of the tray divisions could be conducted separately along the horizontal and vertical axes of the shelf file. A template horizontal divider and vertical divider was selected manually and cropped out from an existing shelf image. The OpenCV template matching functionality was then utilized to find the similarity of the image at each pixel to the template horizontal and vertical dividers. To find the horizontal coordinates for the half tray divisions, a template matching matrix between the shelf image and the template horizontal divider was created. An array containing the maximum matching score along each column was created. The SciPy peak finding algorithm was then used to find the indices of peaks along this maximum array (Virtanen et al. 2020). The peak finder allowed a specification of the minimum distance between peaks which was set at the approximate width of a half tray. The indices returned by the algorithm were the indices along the horizontal axis of the shelf image that matched most closely with the divider template. A similar approach was utilized to find the locations along the vertical axis where tray divisions lay. A template matching result matrix was computed between the shelf image and a template vertical divider. An array containing the maximum template matching result for each row was computed. This was inputted into the SciPy peak finding algorithm to find the vertical indices that most matched the template vertical divider. These vertical indices were taken as the locations along the vertical axis where the half tray divisions lay. After finding the horizontal and vertical coordinates along which to divide the shelf image, the segmentation was performed by slicing the image array along the indices found by the template matching procedure. At this time, the pipeline ensured that each half tray slice fell within certain minimum and maximum dimensions. Any half trays that did not meet these minimum or maximum dimension conditions were separated for manual segmentation.

The next task was to locate and identify the colored dots within each tray image. As the colored dots lay at the center of each tray, a rectangular region of interest was constructed at the center of each tray image. Next, HSI (Hue, Saturation, Intensity) thresholding was conducted to filter out the grey borders of the tray, as well as any soil or plants that lay within the region of interest. Noise removal was conducted to remove any small segments which were not filtered out by the HSI threshold. Contour detection was then conducted, and the three largest contours present in the region of interest were isolated as the three dots. All other contours within the region of interest were discarded and a median blur was applied as additional noise reduction. If fewer than three dot contours were identified by the pipeline, the half tray image was separated from the pipeline for manual processing.

To read the dot codes, it was necessary to determine the color of a dot, as well as whether there was a black dot at the center of the colored dot. The color of each dot was either red, light blue, dark blue, or purple. To assign a dot to one of these color categories, the median hue value within the dot contour was computed and compared to the reference hue value for each of the possible colors. The median was used in this case to lower the susceptibility of the pipeline to any possibly remaining noise. The dot was assigned to the color whose reference hue value was closest to its median hue value. The presence of an inner black dot was determined by checking whether the average intensity (brightness) of pixels in a 20 pixel area at the center of the dot contour fell below a threshold value.

The information contained by the dot code reading was stored as a tuple of strings and booleans of the following format: (colordot1, blackdot1, colordot2, blackdot2, colordot3, blackdot3). In this format, each colordot is a string value containing the name of the color category of the dot, and each blackdot is a boolean value indicating the presence of a central black dot.

Each dot code uniquely corresponded to a tray number. The correspondences between tray number and dot code were shown in a reference dot file. The reference dot file was scanned by the dot code reader to associate the dot code tuples with their corresponding tray numbers. A database associating the dot code tuples and their tray numbers was saved. After each half tray's dot code was read, the dot code tuple was looked up in the database and the half tray was assigned to that tray number. If a dot code tuple was not present within the database, the associated half tray was isolated for manual processing.

After determining the tray identity, each individual plant pot had to be segmented from the rest of the tray so that each plant could be measured individually. This had to be done accurately so that plants were completely contained within their pot, as otherwise there would be inaccuracy in the measurements. The segmentation was conducted by utilizing the locations of the dot centers found during the dot reading process. This is depicted in Figure A.13. The horizontal coordinates of each dot center were used as the positions along the horizontal axis at which to split the half tray image. As the vertical coordinates of the dot centers were nearly the same, the average vertical coordinate of the dot centers was used to compute the division along the vertical axis between pots. Each individual pot image was then associated with its position within the half tray, the tray number it resided in, as well as the time point that the measurement was taken.

Once all the plants pots had been separated and associated with their identifying metadata, they were fed into the measurement phase of the pipeline. Each plant pot image used a HSI threshold to isolate green colored rosettes. Next, contour detection was applied over the image. Contours falling below a threshold area value were discarded as noise and the sum of the areas of all other detected contours was returned as the measurement of rosette area

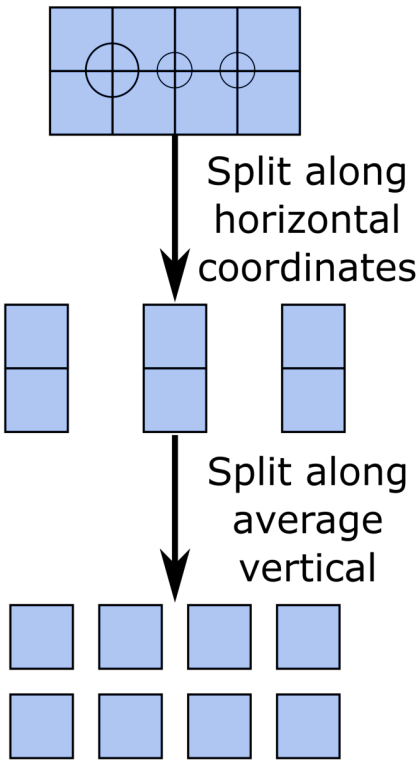

Figure A.13: Splitting a half tray into its constituent pots. The coordinates of the center of each dot are used to identify where to split a half tray along the horizontal and vertical axis. The horizontal coordinates of the split are the horizontal coordinates of each dot center. The vertical coordinate of the split is the average vertical coordinate of all three dot centers.

for that pot. Plant measurements were recorded in an SQLite database, along with the tray number in which they resided, their pot position number and time point that the measurement corresponded to. This database was utilized for further analysis of the data.

Supplementary Figures

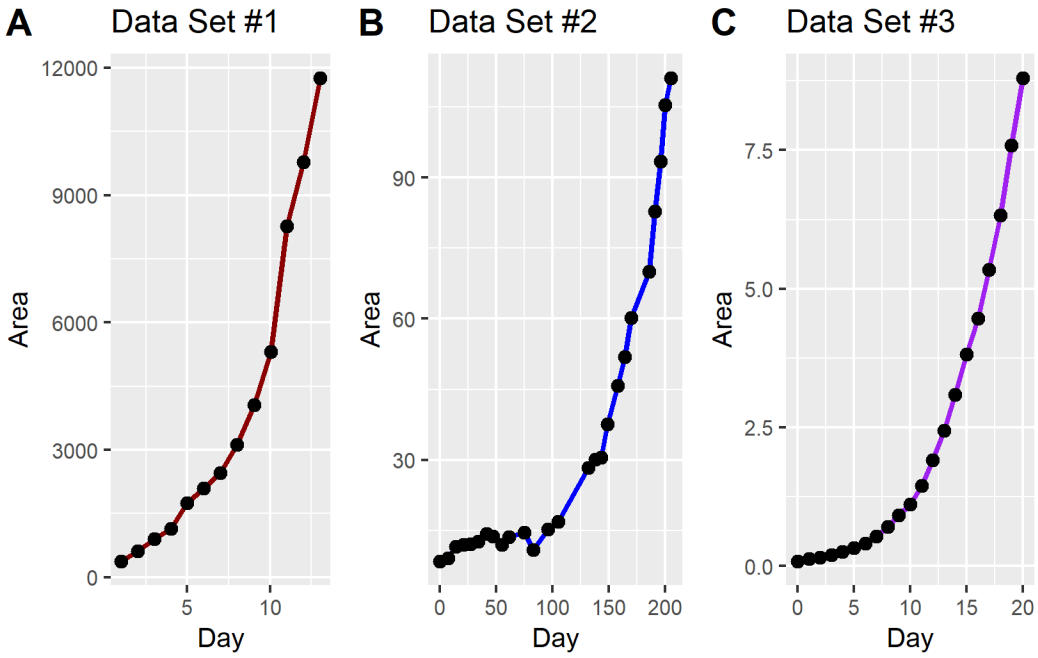

Figure A.14: Raw data of a sample plant for each of the data sets. (A) Sample plant from Data Set #1 (*Arabidopsis*). (B) Sample plant from Data Set #2 (wheat). (C) Sample plant from Data Set #3 (*Arabidopsis*).

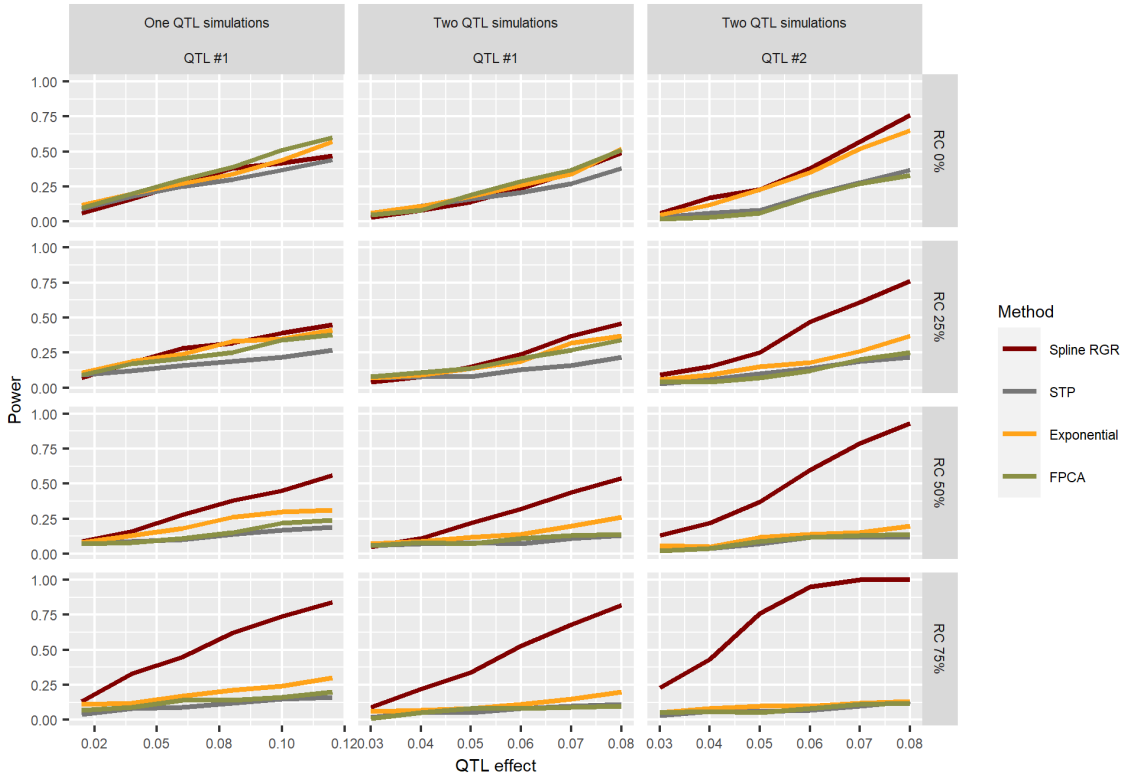

Figure A.15: Power curves of the four different methods across the simulations and RGR parameter correlations. The first column depicts the results for the one-QTL simulations, and the latter two columns show the results for the two-QTL simulations. Each row corresponds to a different RGR parameter correlation (RC), ranging from 0% - 75%. The four different methods are: STP (STP, gray), exponential/parametric (Exponential, yellow), FPCA (FPCA, olive), and our method (`SplineRGR`, dark red). The x-axis depicts QTL effect, with increasing QTL effect as you go from left to right.

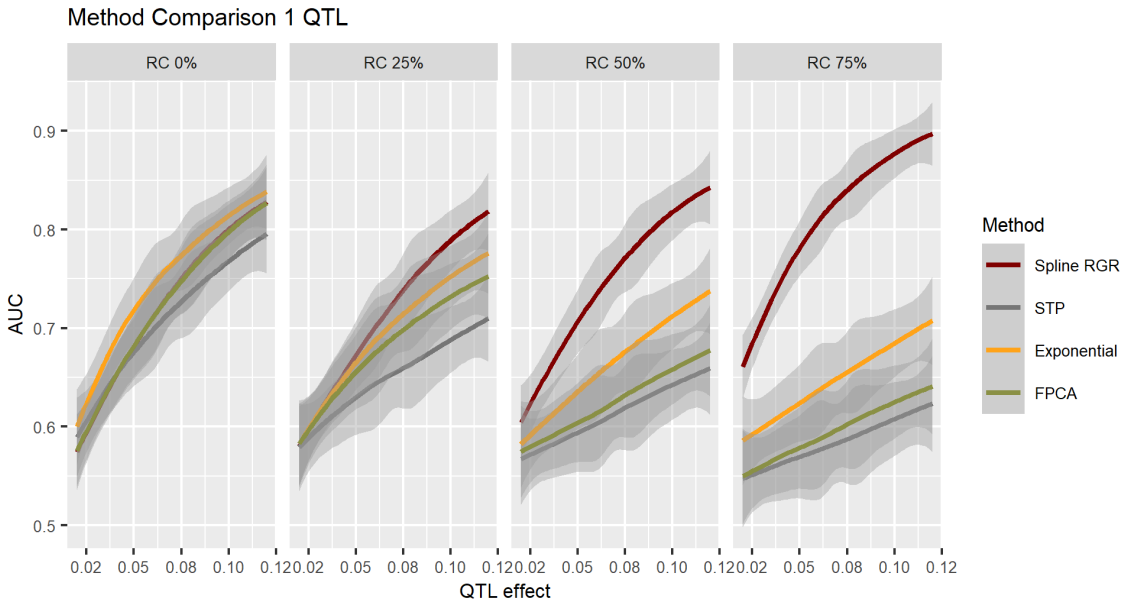

Figure A.16: AUC as a function of QTL effect for the one-QTL simulations (exponential). The mean AUC across the different simulations (QTL effect) are shown, with the grey bands indicating confidence intervals. The methods shown are the STP method (STP, gray), exponential/parametric (Exponential, yellow), FPCA (FPCA, olive), and our method (`SplineRGR`, dark red). Each panel corresponds to a different level of RGR parameter correlation (RC 0% - RC 75%).

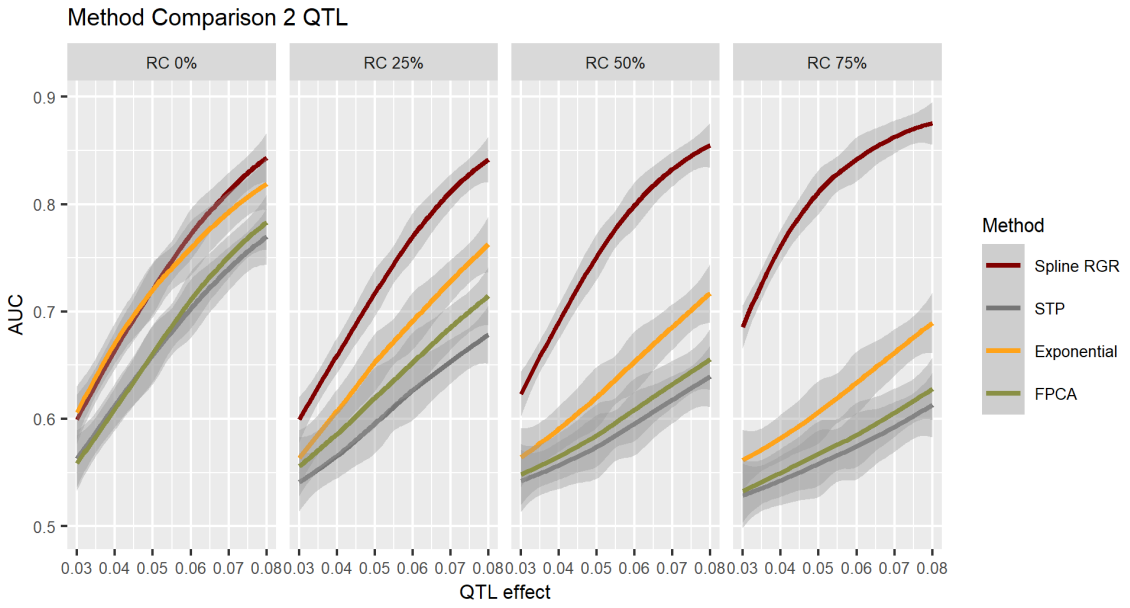

Figure A.17: AUC curves for the two-QTL simulations (exponential). The mean AUC across the different simulations (QTL effect) are shown, with the grey bands indicating confidence intervals. The methods shown are the STP method (STP, gray), exponential/parametric (Exponential, yellow), FPCA (FPCA, olive), and our method (`SplineRGR`, dark red). Each panel corresponds to a different level of RGR parameter correlation (RC 0% - RC 75%).

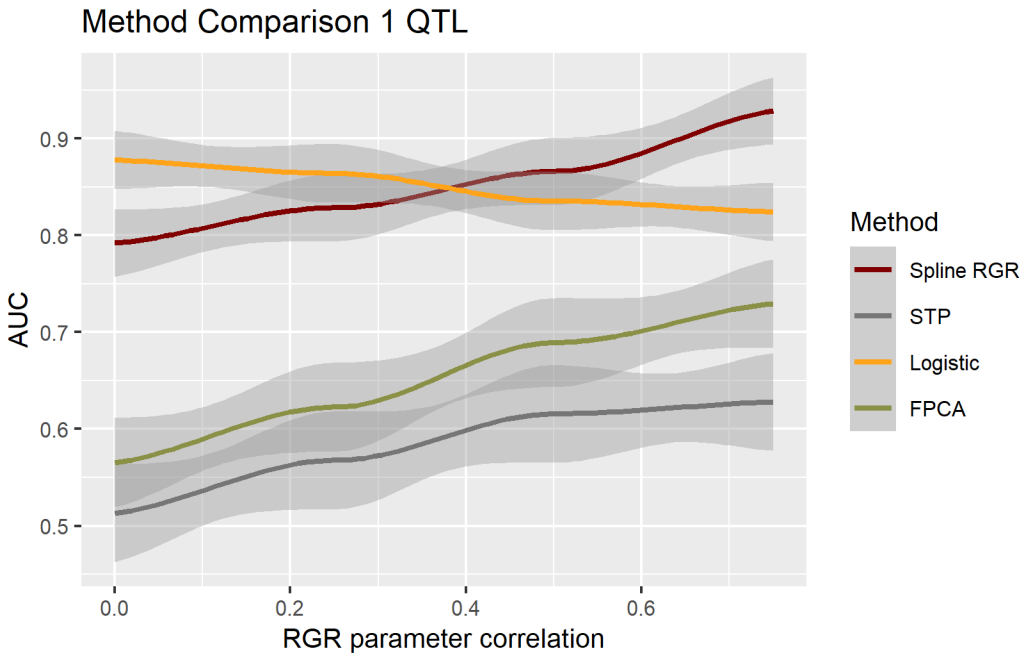

Figure A.18: AUC curves for the one-QTL simulations (logistic). The mean AUC across the different simulations (QTL effect) are shown, with the grey bands indicating confidence intervals. The methods shown are the STP method (STP, gray), logistic/parametric (Logistic, yellow), FPCA (FPCA, olive), and our method (`SplineRGR`, dark red). From left to right is increasing RGR parameter correlation, from 0% to 75%.

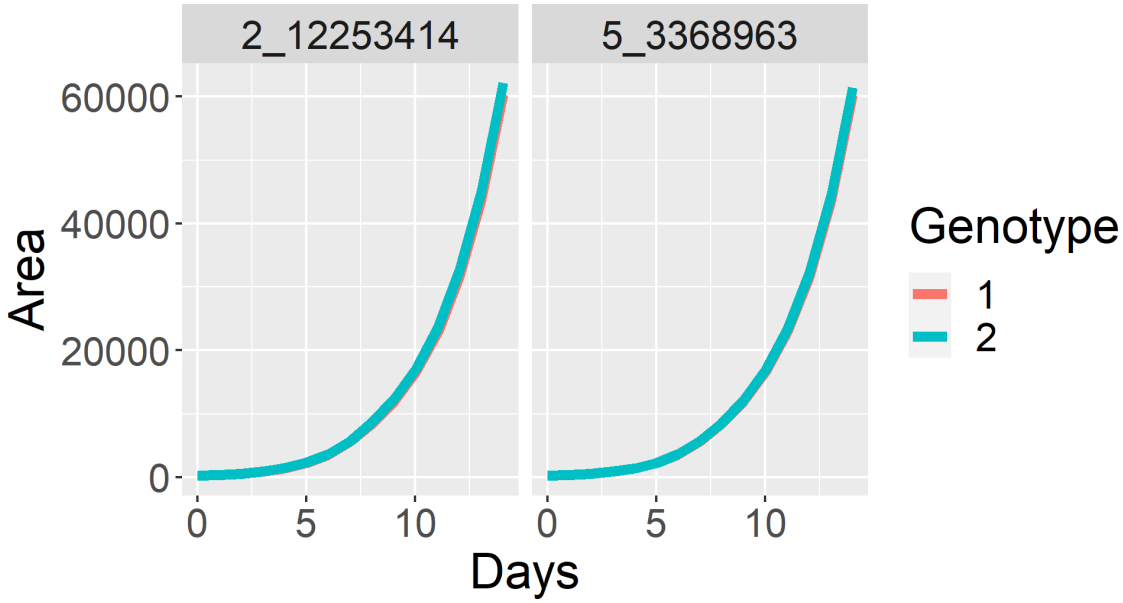

Figure A.19: Estimated QTL effects on the overall growth curve for the two QTL detected in Data Set #1 using `SplineRGR`. Genotype 1 represents the Col-0 allele, while Genotype 2 represents the Sha allele.

Bibliography

Adler, P. B., Salguero-Gomez, R., Compagnoni, A., Hsu, J. S., Ray-Mukherjee, J., Mbeau-Ache, C., and Franco, M. (2014). Functional traits explain variation in plant life history strategies. *Proceedings of the National Academy of Sciences*, 111(2):740–745.

Alokam, S., Chinnappa, C. C., and Reid, D. M. (2002). Red/far-red light mediated stem elongation and anthocyanin accumulation in *Stellaria longipes*: Differential response of alpine and prairie ecotypes. *Canadian Journal of Botany*, 80(1):72–81.

An, N., Palmer, C. M., Baker, R. L., Markelz, R. C., Ta, J., Covington, M. F., Maloof, J. N., Welch, S. M., and Weinig, C. (2016). Plant high-throughput phenotyping using photogrammetry and imaging techniques to measure leaf length and rosette area. *Computers and Electronics in Agriculture*, 127:376–394.

An, N., Welch, S. M., Markelz, R. C., Baker, R. L., Palmer, C. M., Ta, J., Maloof, J. N., and Weinig, C. (2017). Quantifying time-series of leaf morphology using 2D and 3D photogrammetry methods for high-throughput plant phenotyping. *Computers and Electronics in Agriculture*, 135:222–232.

Balasubramanian, S., Schwartz, C., Singh, A., Warthmann, N., Kim, M. C., Maloof, J. N., Loudet, O., Trainer, G. T., Dabi, T., Borevitz, J. O., Chory, J., and Weigel, D. (2009). QTL Mapping in New Arabidopsis thaliana Advanced Intercross-Recombinant Inbred Lines. *PLoS ONE*, 4(2):e4318.

Ballaré, C. L., Scopel, A. L., and Sánchez, R. A. (1997). Foraging for light: Photosensory ecology and agricultural implications. *Plant, Cell and Environment*, 20(6):820–825.

Bentley, A. R., Scutari, M., Gosman, N., Faure, S., Bedford, F., Howell, P., Cockram, J., Rose, G. A., Barber, T., Irigoyen, J., Horsnell, R., Pumfrey, C., Winnie, E., Schacht, J., Beauchêne, K., Praud, S., Greenland, A., Balding, D., and Mackay, I. J. (2014). Applying association mapping and genomic selection to the dissection of key traits in elite European wheat. *Theoretical and Applied Genetics*, 127(12):2619–2633.

Bielczynski, L. W., Łącki, M. K., Hoefnagels, I., Gambin, A., and Croce, R. (2017). Leaf and Plant Age Affects Photosynthetic Performance and Photoprotective Capacity. *Plant Physiology*, 175(4):1634–1648.

Boccalandro, H. E., Ploschuk, E. L., Yanovsky, M. J., Sánchez, R. A., Gatz, C., and Casal, J. J. (2003). Increased Phytochrome B Alleviates Density Effects on Tuber Yield of Field Potato Crops. *Plant Physiology*, 133(4):1539–1546.

Bongers, F. J., Evers, J. B., Anten, N. P. R., and Pierik, R. (2014). From shade avoidance responses to plant performance at vegetation level: using virtual plant modelling as a tool. *New Phytologist*, 204(2):268–272.

Botto, J. F. (2015). Plasticity to simulated shade is associated with altitude in structured popu-

lations of Arabidopsis thaliana. *Plant, Cell and Environment*, 38(7):1321–1332.

Botto, J. F. and Smith, H. (2002). Differential genetic variation in adaptive strategies to a common environmental signal in *Arabidopsis* accessions: phytochrome-mediated shade avoidance. *Plant, Cell and Environment*, 25(1):53–63.

Bradbury, P. J., Zhang, Z., Kroon, D. E., Casstevens, T. M., Ramdoss, Y., and Buckler, E. S. (2007). TASSEL: software for association mapping of complex traits in diverse samples. *Bioinformatics*, 23(19):2633–2635.

Bradski, G. (2000). The OpenCV Library. *Dr. Dobb's Journal of Software Tools*.

Brock, M. T., Rubin, M. J., DellaPenna, D., and Weinig, C. (2020). A Nested Association Mapping Panel in *Arabidopsis thaliana* for Mapping and Characterizing Genetic Architecture. *G3 Genes|Genomes|Genetics*, 10(10):3701–3708.

Broman, K. W., Gatti, D. M., Simecek, P., Furlotte, N. A., Prins, P., Sen, S., Yandell, B. S., and Churchill, G. A. (2019). R/qtl2: Software for Mapping Quantitative Trait Loci with High-Dimensional Data and Multiparent Populations. *Genetics*, 211(2):495–502.

Broman, K. W., Wu, H., Sen, S., and Churchill, G. A. (2003). R/qtl: QTL mapping in experimental crosses. *Bioinformatics*, 19(7):889–890.

Buerkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1):1–28.

Cagnola, J. I., Ploschuk, E., Benech-Arnold, T., Finlayson, S. A., and Casal, J. J. (2012). Stem Transcriptome Reveals Mechanisms to Reduce the Energetic Cost of Shade-Avoidance Responses in Tomato. *Plant Physiology*, 160(2):1110–1119.

Carriedo, L. G., Maloof, J. N., and Brady, S. M. (2016). Molecular control of crop shade avoidance. *Current Opinion in Plant Biology*, 30:151–158.

Casal, J. J. (2012). Shade avoidance. *Arabidopsis Book*, 10(4):e0157.

Casal, J. J. (2013). Canopy Light Signals and Crop Yield in Sickness and in Health. *ISRN Agronomy*, 2013:1–16.

Cheng, C., Krishnakumar, V., Chan, A. P., Thibaud-Nissen, F., Schobel, S., and Town, C. D. (2017). Araport11: a complete reannotation of the *Arabidopsis thaliana* reference genome. *The Plant Journal*, 89(4):789–804.

Cheng, R., Lim, J. E., Samocha, K. E., Sokoloff, G., Abney, M., Skol, A. D., and Palmer, A. A. (2010). Genome-Wide Association Studies and the Problem of Relatedness Among Advanced Intercross Lines and Other Highly Recombinant Populations. *Genetics*, 185(3):1033–1044.

Ciolfi, A., Sessa, G., Sassi, M., Possenti, M., Salvucci, S., Carabelli, M., Morelli, G., and Ruberti, I. (2013). Dynamics of the Shade-Avoidance Response in Arabidopsis. *Plant Physiology*, 163(1):331–353.

Coen, E. S. and Meyerowitz, E. M. (1991). The war of the whorls: genetic interactions controlling flower development. *Nature*, 353(6339):31–37.

Cole, B., Kay, S. A., and Chory, J. (2011). Automated analysis of hypocotyl growth dynamics during shade avoidance in Arabidopsis. *Plant Journal*, 65(6):991–1000.

Coluccio, M. P., Sanchez, S. E., Kasulin, L., Yanovsky, M. J., and Botto, J. F. (2011). Genetic mapping of natural variation in a shade avoidance response: *ELF3* is the candidate gene for a QTL in hypocotyl growth regulation. *Journal of Experimental Botany*, 62(1):167–176.

Dechaine, J. M., Brock, M. T., Iniguez-Luy, F. L., and Weinig, C. (2014). Quantitative trait loci × environment interactions for plant morphology vary over ontogeny in Brassica rapa. *New*

Phytologist, 201(2):657–669.

Demotes-Mainard, S., Péron, T., Corot, A., Bertheloot, J., Le Gourrierec, J., Pelleschi-Travier, S., Crespel, L., Morel, P., Huché-Thélier, L., Boumaza, R., Vian, A., Guérin, V., Leduc, N., and Sakr, S. (2016). Plant responses to red and far-red lights, applications in horticulture. *Environmental and Experimental Botany*, 121:4–21.

Deng, W., Ying, H., Helliwell, C. A., Taylor, J. M., Peacock, W. J., and Dennis, E. S. (2011). FLOWERING LOCUS C (FLC) regulates development pathways throughout the life cycle of *Arabidopsis*. *Proceedings of the National Academy of Sciences of the United States of America*, 108(16):6680–6685.

Diaz-Garcia, L., Covarrubias-Pazaran, G., Schlautman, B., Grygleski, E., and Zalapa, J. (2018). Image-based phenotyping for identification of QTL determining fruit shape and size in American cranberry (*Vaccinium macrocarpon* L.). *PeerJ*, 6:e5461.

Dudley, S. A. and Schmitt, J. (1995). Genetic Differentiation in Morphological Responses to Simulated Foliage Shade between Populations of *Impatiens capensis* from Open and Woodland Sites. *Functional Ecology*, 9(4):655–666.

El-Soda, M., Malosetti, M., Zwaan, B. J., Koornneef, M., and Aarts, M. G. (2014). Genotype × environment interaction QTL mapping in plants: lessons from Arabidopsis. *Trends in Plant Science*, 19(6):390–398.

Fahlgren, N., Feldman, M., Gehan, M., Wilson, M., Shyu, C., Bryant, D., Hill, S., McEntee, C., Warnasooriya, S., Kumar, I., Ficor, T., Turnipseed, S., Gilbert, K., Brutnell, T., Carrington, J., Mockler, T., and Baxter, I. (2015a). A versatile phenotyping system and analytics platform reveals diverse temporal responses to water availability in setaria. *Molecular Plant*, 8(10):1520–1535.

Fahlgren, N., Gehan, M. A., and Baxter, I. (2015b). Lights, camera, action: high-throughput plant phenotyping is ready for a close-up. *Current Opinion in Plant Biology*, 24:93–99.

Filiault, D. L. and Maloof, J. N. (2012). A Genome-Wide Association Study Identifies Variants Underlying the *Arabidopsis thaliana* Shade Avoidance Response. *PLoS Genetics*, 8(3):e1002589.

Fournier-Level, A., Wilczek, A. M., Cooper, M. D., Roe, J. L., Anderson, J., Eaton, D., Moyers, B. T., Petipas, R. H., Schaeffer, R. N., Pieper, B., Reymond, M., Koornneef, M., Welch, S. M., Remington, D. L., and Schmitt, J. (2013). Paths to selection on life history loci in different natural environments across the native range of *Arabidopsis thaliana*. *Molecular Ecology*, 22(13):3552–3566.

Franklin, K. A., Davis, S. J., Stoddart, W. M., Vierstra, R. D., and Whitelam, G. C. (2003a). Mutant Analyses Define Multiple Roles for Phytochrome C in Arabidopsis Photomorphogenesis. *Plant Cell*, 15(9):1981–1989.

Franklin, K. A., Praekelt, U., Stoddart, W. M., Billingham, O. E., Halliday, K. J., and Whitelam, G. C. (2003b). Phytochromes B, D, and E Act Redundantly to Control Multiple Physiological Responses in Arabidopsis. *Plant Physiology*, 131(3):1340–1346.

Franklin, K. A. and Whitelam, G. C. (2005). Phytochromes and shade-avoidance responses in plants. *Annals of Botany*, 96(2):169–175.

Gao, F., Wen, W., Liu, J., Rasheed, A., Yin, G., Xia, X., Wu, X., and He, Z. (2015). Genome-Wide Linkage Mapping of QTL for Yield Components, Plant Height and Yield-Related Physiological Traits in the Chinese Wheat Cross Zhou 8425B/Chinese Spring. *Frontiers in Plant Science*, 6.

Gazzani, S., Gendall, A. R., Lister, C., and Dean, C. (2003). Analysis of the Molecular Basis of Flowering Time Variation in Arabidopsis Accessions. *Plant Physiology*, 132(2):1107–1114.

Gnan, S., Marsh, T., and Kover, P. X. (2017). Inflorescence photosynthetic contribution to fitness releases *Arabidopsis thaliana* plants from trade-off constraints on early flowering. *PLoS ONE*, 12(10):e0185835.

Gratani, L. (2014). Plant Phenotypic Plasticity in Response to Environmental Factors. *Advances in Botany*, 2014:1–17.

Griffiths, S., Simmonds, J., Leverington, M., Wang, Y., Fish, L., Sayers, L., Alibert, L., Orford, S., Wingen, L., and Snape, J. (2012). Meta-QTL analysis of the genetic control of crop height in elite European winter wheat germplasm. *Molecular Breeding*, 29(1):159–171.

Guan, P., Lu, L., Jia, L., Kabir, M. R., Zhang, J., Lan, T., Zhao, Y., Xin, M., Hu, Z., Yao, Y., Ni, Z., Sun, Q., and Peng, H. (2018). Global QTL Analysis Identifies Genomic Regions on Chromosomes 4A and 4B Harboring Stable Loci for Yield-Related Traits Across Different Environments in Wheat (Triticum aestivum L.). *Frontiers in Plant Science*, 9:529.

Haley, C. S. and Knott, S. A. (1992). A simple regression method for mapping quantitative trait loci in line crosses using flanking markers. *Heredity*, 69(4):315–324.

Heuvelink, E., Bakker, M. J., Elings, A., Kaarsemaker, R., and Marcelis, L. F. (2005). Effect of leaf area on tomato yield. *Acta Horticulturae*, 691:43–50.

Honsdorf, N., March, T. J., Berger, B., Tester, M., and Pillen, K. (2014). High-Throughput Phenotyping to Detect Drought Tolerance QTL in Wild Barley Introgression Lines. *PLoS ONE*, 9(5):e97047.

Hu, L. and Bentler, P. M. (1999). Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling: A Multidisciplinary Journal*, 6(1):1–55.

Initiative, T. A. G. (2000). Analysis of the genome sequence of the flowering plant arabidopsis thaliana. *Nature*, 408(6814):796–815.

Jiménez-Gómez, J. M., Wallace, A. D., and Maloof, J. N. (2010). Network Analysis Identifies *ELF3* as a QTL for the Shade Avoidance Response in Arabidopsis. *PLoS Genetics*, 6(9):e1001100.

Kami, C., Lorrain, S., Hornitschek, P., and Fankhauser, C. (2010). Light-regulated plant growth and development. *Current Topics in Developmental Biology*, 91:29–66.

Kebrom, T. H., Burson, B. L., and Finlayson, S. A. (2006). Phytochrome B Represses *Teosinte Branched1* Expression and Induces Sorghum Axillary Bud Outgrowth in Response to Light Signals. *Plant Physiology*, 140(3):1109–1117.

Kendziorski, C. M., Cowley, A. W., Greene, A. S., Salgado, H. C., Jacob, H. J., and Tonellato, P. J. (2002). Mapping Baroreceptor Function to Genome: A Mathematical Modeling Approach. *Genetics*, 160(4):1687–1695.

Kerdaffrec, E., Filiault, D. L., Korte, A., Sasaki, E., Nizhynska, V., Seren, U., and Nordborg, M. (2016). Multiple alleles at a single locus control seed dormancy in Swedish *Arabidopsis*. *eLife*, 5:e22502.

Keuskamp, D. H., Sasidharan, R., and Pierik, R. (2010). Physiological regulation and functional significance of shade avoidance responses to neighbors. *Plant Signaling and Behavior*, 5(6):655–662.

Knoch, D., Abbadi, A., Grandke, F., Meyer, R. C., Samans, B., Werner, C. R., Snowdon, R. J., and Altmann, T. (2020). Strong temporal dynamics of QTL action on plant growth

progression revealed through high-throughput phenotyping in canola. *Plant Biotechnology Journal*, 18:68–82.

Kover, P. X., Valdar, W., Trakalo, J., Scarcelli, N., Ehrenreich, I. M., Purugganan, M. D., Durrant, C., and Mott, R. (2009). A Multiparent Advanced Generation Inter-Cross to Fine-Map Quantitative Traits in Arabidopsis thaliana. *PLoS Genetics*, 5(7):e1000551.

Kwak, I.-Y., Moore, C. R., Spalding, E. P., and Broman, K. W. (2014). A Simple Regression-Based Method to Map Quantitative Trait Loci Underlying Function-Valued Phenotypes. *Genetics*, 197(4):1409–1416.

Kwak, I.-Y., Moore, C. R., Spalding, E. P., and Broman, K. W. (2016). Mapping Quantitative Trait Loci Underlying Function-Valued Traits Using Functional Principal Component Analysis and Multi-Trait Mapping. *G3 Genes|Genomes|Genetics*, page 8.

Lempe, J., Balasubramanian, S., Sureshkumar, S., Singh, A., Schmid, M., and Weigel, D. (2005). Diversity of Flowering Responses in Wild *Arabidopsis thaliana* Strains. *PLoS Genetics*, 1(1):e6.

Li, R., Tsaih, S.-W., Shockley, K., Stylianou, I. M., Wergedal, J., Paigen, B., and Churchill, G. A. (2006). Structural model analysis of multiple quantitative traits. *PLoS genetics*, 2(7):e114.

Likhachev, D. V. (2021). On the optimization of knot allocation for b-spline parameterization of the dielectric function in spectroscopic ellipsometry data analysis. *Journal of Applied Physics*, 129(3):034903.

Lippert, C., Listgarten, J., Liu, Y., Kadie, C. M., Davidson, R. I., and Heckerman, D. (2011). FaST linear mixed models for genome-wide association studies. *Nature Methods*, 8(10):833–835.

Lizaso, J. I., Boote, K. J., Jones, J. W., Porter, C. H., Echarte, L., Westgate, M. E., and Sonohat, G. (2011). CSM-IXIM: A New Maize Simulation Model for DSSAT Version 4.5. *Agronomy Journal*, 103(3):766–779.

Lovell, J. T., Juenger, T. E., Michaels, S. D., Lasky, J. R., Platt, A., Richards, J. H., Yu, X., Easlon, H. M., Sen, S., and Mckay, J. K. (2013). Pleiotropy of *FRIGIDA* enhances the potential for multivariate adaptation. *Proceedings of the Royal Society B: Biological Sciences*, 280(1763):20131043.

Lyra, D. H., Virlet, N., Sadeghi-Tehran, P., Hassall, K. L., Wingen, L. U., Orford, S., Griffiths, S., Hawkesford, M. J., and Slavov, G. T. (2020). Functional QTL mapping and genomic prediction of canopy height in wheat measured using a robotic field phenotyping platform. *Journal of Experimental Botany*, 71(6):1885–1898.

Ma, C.-X., Casella, G., and Wu, R. (2002). Functional Mapping of Quantitative Trait Loci Underlying the Character Process: A Theoretical Framework. *Genetics*, 161(4):1751–1762.

Mackay, T. F. C. (2001). The Genetic Architecture of Quantitative Traits. page 37.

Marchadier, E., Hanemian, M., Tisné, S., Bach, L., Bazakos, C., Gilbault, E., Haddadi, P., Virlouvet, L., and Loudet, O. (2019). The complex genetic architecture of shoot growth natural variation in *Arabidopsis thaliana*. *PLoS Genetics*, 15(4):e1007954.

Martin, W., Rujan, T., Richly, E., Hansen, A., Cornelsen, S., Lins, T., Leister, D., Stoebe, B., Hasegawa, M., and Penny, D. (2002). Evolutionary analysis of Arabidopsis, cyanobacterial, and chloroplast genomes reveals plastid phylogeny and thousands of cyanobacterial genes in the nucleus. *Proceedings of the National Academy of Sciences*, 99(19):12246–12251.

McKay, J. K., Richards, J. H., and Mitchell-Olds, T. (2003). Genetics of drought adaptation in *Arabidopsis thaliana*: I. Pleiotropy contributes to genetic correlations among ecological

traits. *Molecular Ecology*, 12(5):1137–1151.

McKay, J. K., Richards, J. H., Nemali, K. S., Sen, S., Mitchell-Olds, T., Boles, S., Stahl, E. A., Wayne, T., and Juenger, T. E. (2008). Genetics of drought adaptation in *Arabidopsis thaliana* II. QTL analysis of a new mapping population, Kas-1 × Tsu-1. *Evolution*, 62(12):3014–3026.

McNellis, T. W., von Arnim, A. G., Araki, T., Komeda, Y., Misera, S., and Deng, X.-W. (1994). Genetic and Molecular Analysis of an Allelic Series of *cop1* Mutants Suggests Functional Roles for the Multiple Protein Domains. *Plant Cell*, 6(4):487–500.

Meng, L., Li, H., Zhang, L., and Wang, J. (2015). QTL IciMapping: Integrated software for genetic linkage map construction and quantitative trait locus mapping in biparental populations. *The Crop Journal*, 3(3):269–283.

Mitchell-Olds, T. (1996). Genetic Constraints on Life-History Evolution: Quantitative-Trait Loci Influencing Growth and Flowering in *Arabidopsis thaliana*. *Evolution*, 50(1):140–145.

Morgan, D. C. and Smith, H. (1979). A systematic relationship between phytochrome-controlled development and species habitat, for plants grown in simulated natural radiation. *Planta*, 145(3):253–258.

Nozue, K., Devisetty, U. K., Lekkala, S., Mueller-Moulé, P., Bak, A., Casteel, C. L., and Maloof, J. N. (2018). Network Analysis Reveals a Role for Salicylic Acid Pathway Components in Shade Avoidance. *Plant Physiology*, 178(4):1720–1732.

Nozue, K., Tat, A. V., Kumar Devisetty, U., Robinson, M., Mumbach, M. R., Ichihashi, Y., Lekkala, S., and Maloof, J. N. (2015). Shade Avoidance Components and Pathways in Adult Plants Revealed by Phenotypic Profiling. *PLoS Genetics*, 11(4):e1004953.

Paul-Victor, C., Züst, T., Rees, M., Kliebenstein, D. J., and Turnbull, L. A. (2010). A new method for measuring relative growth rate can uncover the costs of defensive compounds in *Arabidopsis thaliana*. *New Phytologist*, 187(4):1102–1111.

Pommerening, A. and Muszta, A. (2015). Methods of modelling relative growth rate. *Forest Ecosystems*, 2(1):5.

R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rees, M., Osborne, C. P., Woodward, F. I., Hulme, S. P., Turnbull, L. A., and Taylor, S. H. (2010). Partitioning the Components of Relative Growth Rate: How Important Is Plant Size Variation? *The American Naturalist*, 176(6):E152–E161.

Ritchie, J. T. (1985). Description and performance of ceres wheat : A user-oriented wheat yield model. *ARS Wheat Yield Project*, pages 159–175.

Robson, P. H. R., McCormac, A. C., Irvine, A. S., and Smith, H. (1996). Genetic engineering of harvest index in tobacco through overexpression of a phytochrome gene. *Nature Biotechnology*, 14:995–998.

Rose, K. E., Atkinson, R. L., Turnbull, L. A., and Rees, M. (2009). The costs and benefits of fast living: The costs and benefits of fast living. *Ecology Letters*, 12(12):1379–1384.

Rossel, Y. (2012). lavaan: An R Package for Structural Equation Modeling. *Journal of Statistical Software*, 48:1–36.

Runcie, D. E. and Crawford, L. (2019). Fast and general-purpose linear mixed models for genome-wide genetics. *PLoS Genetics*, 15(2):e1007978.

Rutkoski, J., Poland, J., Mondal, S., Autrique, E., Pérez, L. G., Crossa, J., Reynolds, M., and Singh, R. (2016). Canopy Temperature and Vegetation Indices from High-Throughput Phenotyping Improve Accuracy of Pedigree and Genomic Selection for Grain Yield in Wheat.

G3 Genes|Genomes|Genetics, 6(9):2799–2808.

Salomé, P. A., Bomblies, K., Laitinen, R. A., Yant, L., Mott, R., and Weigel, D. (2011). Genetic Architecture of Flowering-Time Variation in *Arabidopsis thaliana*. *Genetics*, 188(2):421–433.

Schmid, B. (1992). Phenotypic Variation in Plants. *Evolutionary Trends in Plants*, 6(1):45–60.

Schmitt, J., Dudley, S. A., and Pigliucci, M. (1999). Manipulative Approaches to Testing Adaptive Plasticity: Phytochrome-Mediated Shade-Avoidance Responses in Plants. *The American Naturalist*, 154(S1):S43–S54.

Schmitt, J., Stinchcombe, J. R., Shane Heschel, M., and Huber, H. (2003). The Adaptive Evolution of Plasticity: Phytochrome-Mediated Shade Avoidance Responses. *Integrative and Comparative Biology*, 43(3):459–469.

Schrager-Lavelle, A., Herrera, L. A., and Maloof, J. N. (2016). Tomato phyE Is Required for Shade Avoidance in the Absence of phyB1 and phyB2. *Frontiers in Plant Science*, 7:1275.

Sellaro, R., Pacín, M., and Casal, J. J. (2017). Meta-Analysis of the Transcriptome Reveals a Core Set of Shade-Avoidance Genes in Arabidopsis. *Photochemistry and Photobiology*, 93(3):692–702.

Sheehan, M. J., Kennedy, L. M., Costich, D. E., and Brutnell, T. P. (2007). Subfunctionalization of *PhyB1* and *PhyB2* in the control of seedling and mature plant traits in maize. *Plant Journal*, 49(2):338–353.

Shindo, C., Aranzana, M. J., Lister, C., Baxter, C., Nicholls, C., Nordborg, M., and Dean, C. (2005). Role of *FRIGIDA* and *FLOWERING LOCUS C* in Determining Variation in Flowering Time of Arabidopsis. *Plant Physiology*, 138(2):1163–1173.

Simon, M., Loudet, O., Durand, S., Bérard, A., Brunel, D., Sennesal, F. X., Durand-Tardif, M., Pelletier, G., and Camilleri, C. (2008). Quantitative Trait Loci Mapping in Five New Large Recombinant Inbred Line Populations of *Arabidopsis thaliana* Genotyped With Consensus Single-Nucleotide Polymorphism Markers. *Genetics*, 178(4):2253–2264.

Sing, T., Sander, O., Beerenwinkel, N., and Lengauer, T. (2005). ROCR: visualizing classifier performance in R. *Bioinformatics*, 21(20):3940–3941.

Soetaert, K., Petzoldt, T., and Setzer, R. W. (2010). Solving Differential Equations in *R* : Package **deSolve**. *Journal of Statistical Software*, 33(9).

Somerville, C. R. (2001). An Early Arabidopsis Demonstration. Resolving a Few Issues Concerning Photorespiration. *Plant Physiology*, 125(1):20–24.

Ta, J., Palmer, C., Brock, M., Rubin, M., Weinig, C., Maloof, J., and Runcie, D. (2020). Multiple Loci Control Variation in Plasticity to Foliar Shade Throughout Development in *Arabidopsis thaliana*. *G3 Genes|Genomes|Genetics*, 10(11):4103–4114.

Tian, X., Wen, W., Xie, L., Fu, L., Xu, D., Fu, C., Wang, D., Chen, X., Xia, X., Chen, Q., He, Z., and Cao, S. (2017). Molecular Mapping of Reduced Plant Height Gene Rht24 in Bread Wheat. *Frontiers in Plant Science*, 8:1379.

Turnbull, L. A., Paul-Victor, C., Schmid, B., and Purves, D. W. (2008). GROWTH RATES, SEED SIZE, AND PHYSIOLOGY: DO SMALL-SEEDED SPECIES REALLY GROW FASTER. *Ecology*, 89(5):1352–1363.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa,

F., van Mulbregt, P., and SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.

Volpato, L., Pinto, F., González-Pérez, L., Thompson, I. G., Borém, A., Reynolds, M., Gérard, B., Molero, G., and Rodrigues, F. A. (2021). High Throughput Field Phenotyping for Plant Height Using UAV-Based RGB Imagery in Wheat Breeding Lines: Feasibility and Validation. *Frontiers in Plant Science*, 12:591587.

Weinig, C. (2000). Differing selection in alternative competitive environments: shade-avoidance responses and germination timing. *Evolution*, 54(1):124–136.

Weller, J. L., Schreuder, M. E., Smith, H., Koornneef, M., and Kendrick, R. E. (2000). Physiological interactions of phytochromes A, B1 and B2 in the control of development in tomato. *Plant Journal*, 24(3):345–356.

Werner, J. D., Borevitz, J. O., Henriette Uhlenhaut, N., Ecker, J. R., Chory, J., and Weigel, D. (2005). *FRIGIDA*-Independent Variation in Flowering Time of Natural *Arabidopsis thaliana* Accessions. *Genetics*, 170(3):1197–1207.

West, G. B., Brown, J. H., and Enquist, B. J. (2001). A general model for ontogenetic growth. *Nature*, 413(6856):628–631.

Wille, W., Pipper, C. B., Rosenqvist, E., Andersen, S. B., and Weiner, J. (2017). Reducing shade avoidance responses in a cereal crop. *AoB PLANTS*, 9(5):plx039.

Wollenberg, A. C., Strasser, B., Cerdan, P. D., and Amasino, R. M. (2008). Acceleration of Flowering during Shade Avoidance in Arabidopsis Alters the Balance between *FLOWERING LOCUS C*-Mediated Repression and Photoperiodic Induction of Flowering. *Plant Physiology*, 148(3):1681–1694.

Woodward, A. W. and Bartel, B. (2018). Biology in Bloom: A Primer on the *Arabidopsis thaliana* Model System. *Genetics*, 208(4):1337–1349.

Wu, R. and Lin, M. (2006). Functional mapping — how to map and study the genetic architecture of dynamic complex traits. *Nature Reviews Genetics*, 7(3):229–237.

Wu, X., Wang, Z., Chang, X., and Jing, R. (2010). Genetic dissection of the developmental behaviours of plant height in wheat under diverse water regimes. *Journal of Experimental Botany*, 61(11):2923–2937.

Würschum, T., Langer, S. M., and Longin, C. F. H. (2015). Genetic control of plant height in European winter wheat cultivars. *Theoretical and Applied Genetics*, 128(5):865–874.

Yap, J. S., Fan, J., and Wu, R. (2009). Nonparametric Modeling of Longitudinal Covariance Structure in Functional Mapping of Quantitative Trait Loci. *Biometrics*, 65(4):1068–1077.

Yu, J., Holland, J. B., McMullen, M. D., and Buckler, E. S. (2008). Genetic Design and Statistical Power of Nested Association Mapping in Maize. *Genetics*, 178(1):539–551.

Zanke, C. D., Ling, J., Plieske, J., Kollers, S., Ebmeyer, E., Korzun, V., Argillier, O., Stiewe, G., Hinze, M., Neumann, K., Ganal, M. W., and Röder, M. S. (2014). Whole Genome Association Mapping of Plant Height in Winter Wheat (Triticum aestivum L.). *PLoS ONE*, 9(11):e113287.

Zhang, X., Huang, C., Wu, D., Qiao, F., Li, W., Duan, L., Wang, K., Xiao, Y., Chen, G., Liu, Q., Xiong, L., Yang, W., and Yan, J. (2017). High-Throughput Phenotyping and QTL Mapping Reveals the Genetic Architecture of Maize Plant Growth. *Plant Physiology*, 173(3):1554–1564.

Zhou, X. and Stephens, M. (2012). Genome-wide efficient mixed-model analysis for association studies. *Nature Genetics*, 44(7):821–824.