

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Landscape Genetics, in 4D: Exploring The Influence of Spatiotemporal Phenomena on Microevolutionary Dynamics

Permalink

<https://escholarship.org/uc/item/0np5q1qr>

Author

Terasaki Hart, Drew Ellison

Publication Date

2022

Peer reviewed|Thesis/dissertation

Landscape Genetics, in 4D: Exploring The Influence of Spatiotemporal Phenomena on
Microevolutionary Dynamics

By

Drew Ellison Terasaki Hart

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Environmental Science, Policy, and Management

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ian J. Wang, Chair

Professor Nina Maggi Kelly

Professor Todd E. Dawson

Fall 2022

© 2022 Drew Ellison Terasaki Hart

Abstract

Landscape Genetics, in 4D: Exploring The Influence of Spatiotemporal Phenomena on Microevolutionary Dynamics

by

Drew Ellison Terasaki Hart

Doctor of Philosophy in Environmental Science, Policy, and Management

University of California, Berkeley

Professor Ian J. Wang, Chair

A major goal of landscape genomics is to understand how spatiotemporal variability in complex environments influences evolutionary dynamics, and consequently geographic patterns of genomic diversity, in natural populations. Recent computational advances enable this spatiotemporal complexity to be described and analyzed in unprecedented detail. One such advance is the improvement of forward-time landscape genomic simulation, allowing arbitrarily complex evolutionary scenarios that mimic real-world systems to be created and studied *in silico*. In chapter 1, I present Geonomics, a new, user-friendly Python package for performing complex, spatially explicit, landscape genomic simulations on changing landscapes and with full spatial pedigrees. I describe the structure and function of Geonomics in detail, show that its results are consistent with expectations for a variety of validation tests based on classical population genetics, then demonstrate its utility and flexibility with example scenarios featuring polygenic selection, selection on multiple traits, and non-stationary environmental change on realistic landscapes. Taken together, these tests and demonstrations establish Geonomics as a robust platform for population genomic simulations that incorporate complex spatiotemporal dynamics.

In chapter 2, I apply the software developed in chapter 1 to one of the major areas of theoretical and applied interest in landscape genomics: the evolutionary consequences of climate change. I model climate change realistically, as the decoupling of historical environmental gradients that generates novel multivariate selective environments. I then simulate evolutionary responses to that climate change event across a range of genomic architectures, defined as the full factorial crossing of discretized levels of three key architectural components: the number of genes per trait (polygenicity), the recombination rate between neighboring genes (linkage), and the number of distinct genotypes generating identical phenotypes (genotypic redundancy). I use the results to test a series of hypotheses about the influence of polygenicity, linkage, and redundancy on gene flow, maladaptation, and demographic decline. Results show that a commonly assumed mechanism of evolutionary rescue, adaptive gene flow from populations whose current climates approximate future projection, can be less effective than *in situ* adaptation under some architectures, likely because of maladaptive introgression caused by the decoupling of

environmental gradients. I also find that high polygenicity aggravates maladaptation and demographic decline, a concerning result given the likely polygenic nature of many climate-adapted traits, but that higher genotypic redundancy increases adaptive capacity across all scenarios, adding to the growing recognition of its importance. Overall, this chapter shows that genomic architecture, though it is often ignored, can exert large influence over the effectiveness and relative magnitudes of adaptive gene flow and *in situ* adaptation in a spatially distributed population subjected to climate change.

Another major computational advance facilitating the study of spatiotemporal evolutionary dynamics is the advent of massive and distributed geocomputation. In chapter 3, I use this tool set to study, in unprecedented detail, global geographic variability in the seasonality of terrestrial plant productivity – i.e., land surface phenology (LSP). Not only does the geography of LSP convey critical information about environmental controls on plant function and carbon cycling, but it has important implications for evolutionary biogeography: spatial asynchrony in LSP can indicate the potential for spatial asynchrony in reproductive phenology, and thus for increased genetic isolation and divergence between conspecific populations. Thus, whereas chapter 2 provides an example of the spatial nature of an evolving system influencing its temporal dynamics, this chapter provides an example of the less-appreciated inverse situation: the potential for a system’s temporal complexity to influence its evolutionary dynamics. Despite its importance, LSP research lacks mapping methodologies that can characterize the full diversity of terrestrial phenologies, and LSP asynchrony mapping is even less developed. Here, I develop a multivariate, generalized, and robustly-validated LSP mapping methodology, based on simple harmonic regression, then apply it to a 10-year, 0.05° dataset of MODIS near-infrared reflectance of vegetation (NIR_v, a proxy of plant productivity). This produces a global LSP map that reveals surprising diversity, including both regional patterns of heterogeneity that are corroborated by prior research and intercontinental patterns of convergence that recapitulate major bioclimatic and biogeographic gradients. Next, I calculate and present a global map of LSP asynchrony, and use machine learning to explore regional variability in its potential climatic and physiographic drivers. I describe LSP asynchrony hotspots in the world’s five Mediterranean climate regions, where asynchrony appears driven by precipitation asynchrony and spatial variability in vegetation structure, and in tropical montane regions, where minimum temperature asynchrony and precipitation asynchrony appear to be interacting drivers. Lastly, I use an ensemble of regressions within global high-asynchrony regions to demonstrate that phenological asynchrony between climatically similar sites is most frequent at lower latitudes, supporting the notion that phenological asynchrony is most likely to cause allochry and consequent evolutionary divergence in the tropics.

to Mom and Dad, who taught me to walk
to Nana, who taught me to walk barefoot
to Science and Art, who guide my walk through life
to Humor and Love, who lift me up when I trip
to Miko, with whom I do my walking
to BB, whose walk is just beginning

Introduction

In nature, evolving populations consist of individuals that are located and mobile in space, yet the fields of population and evolutionary genetics developed around theory and mathematical models that are either aspatial or, at best, spatially implicit (Hartl and Clark 2007). While this body of theory has been applied quite effectually to real populations, the spatial arrangements and movements of individuals, and the spatial patterns of genetic diversity this generates, can both modulate the outcomes expected under classical theory and generate important phenomena not entirely approachable by aspatial classical theory (Slatkin 1987). This realization led to the birth of the field of landscape genetics, a result of the fusion of population genetics – the study of the genetic composition and changes of populations – and landscape ecology – the study of spatially structured ecological processes and the geographic patterns they generate (Manel *et al.* 2003). Though it is quite young, the field of landscape genetics quickly developed its own body of methodologies and provided a wealth of useful basic and applied findings (Manel and Holderegger 2013). However, much more remains to be understood within the ambit of landscape genetics, and numerous horizons remain.

One such horizon is the ability to scale theory and analysis from one or a few genes of interest (i.e., landscape genetics) to large numbers of polymorphic markers and even whole genomes and epigenomes (i.e., landscape genomics; Storfer *et al.* 2018). Numerous new tools, in rapid and ongoing development over the past two decades, are facilitating this transition. A major area of tool development, of course, is the development of new sequencing technology, allowing for large, random samples of genetic markers and even whole genomes to be produced for sizable numbers of individuals sampled from wild populations. A second area of development, that of new computational tools, is necessarily linked to the first, in that improved computing is necessary to process and analyze next-generation sequencing data. New modeling methods specifically suited to the types and volumes of data generated by modern sequencing are critical here, and for landscape genomics this includes improved simulation resources. The large spatial scales, long time scales, large population sizes, and high complexity of most landscape genetics study systems make *in vivo* experimentation inadequate, and thus make *in silico* experimentation critical.

In chapter 1, I introduce one such simulation tool. ‘Geonomics’ is a Python package I developed to enable the simulation of complex evolutionary scenarios on complex and changing landscapes. Simulations are individual-based, spatially explicit, and scriptable, generating simulated populations of individuals, each with its own simulated genome, useful for theoretical, empirical, and applied research, as well as for education. As a flexible scripting framework embedded in one of the most popular scientific programming languages, Geonomics provides an unprecedented combination of ease of use and complexity. I first describe the need for such a tool and describe the structure and function of Geonomics, then validate it against a suite of tests derived from aspatial and spatially implicit population genetics models, and finally demonstrate its utility through a trio of complex simulation scenarios.

Just as landscape genetics was born from the extension of population genetic theory to account for the spatial context of the study systems – i.e., the landscape – another horizon of the field is to continue this extension to incorporate the temporal aspect of landscapes, and thus to better understand how adaptation occurs in both space and time (Rissler 2016). One particularly important goal is to develop theory that can provide the ability to understand and perhaps even

predict evolutionary outcomes under environmental change (Hoffmann and Sgrò 2011). This is an important gap not only from a theoretical perspective, but also from an applied science perspective: climate change, habitat loss, and other forms of anthropogenic environmental change are driving distributional and evolutionary changes in species and communities, often at unprecedented and increasing rates, and the conservation implications of evolutionary responses to global change remain ambiguous in the absence of a robust foundation of evolutionary theory that can help identify species more and less likely to adapt and persist.

In chapter 2, I use Geonomics, the software I presented in chapter 1, to conduct a theoretical study that offers novel insight in this area. Having already passed peer-review, Geonomics now provides a robustly vetted tool for this work. First, rather than modeling climate change as the uniform shift across geographic space of a species' fixed climatic niche, I instead start with a conceptual evolutionary model in which the environmental gradients subtending a species' local adaptation decouple as climate changes, causing the emergence of novel environments (Williams and Jackson 2007) and consequent shifts in the adaptive landscape experienced by local subpopulations, the magnitudes of which vary across the geographic landscape. I then use simulations within this context to test a series of hypotheses about whether and how genomic architecture influences the nature and outcome of evolution under climate change. Simulation results derive from a suite of simulation scenarios using the full factorial crossing of three genomic architecture components of interest: number of genes per locally-adapted trait ('polygenicity'), the quantity of distinct genotypes that can produce identical trait phenotypes ('genotypic redundancy'), and the rate of recombination between neighboring genes ('linkage strength').

A third horizon of landscape genetics is its integration with biogeography more broadly (Rissler 2016). The paramount importance of the spatial context within landscape genetics arises from the recognition that the structure of the environment – e.g., landscape features and movement barriers, environmental gradients that act as selective forces or that determine the local densities of population – can have a predominant influence over processes or outcomes of interest. Because of this focus on the influence of specific landscape characteristics on specific populations, landscape genetics research usually focuses at local to regional scales. However, given that similar geographic contexts arise in numerous places worldwide, we stand to learn much about the nature and repeatability of evolution, and the processes generating the current geographic distribution of biodiversity by studying whether similar geographic contexts worldwide tend to influence microevolutionary processes in similar ways. This is a rich area of inquiry that points at some of the major questions of evolutionary biology – What landscape conditions might cause populations to diverge and species to emerge? Where are those conditions most likely to occur, and why? How, if at all, do spatial patterns of population divergence relate to patterns of speciation and, in turn, to patterns of species diversity? Given the young age and historically parochial focus of the field of landscape genetics, this is a relatively untrodden yet potentially fruitful area of inquiry.

In chapter 3, I provide a global analysis that offers unprecedented insight in this realm: an exploration of the spatial asynchrony of phenology, a biogeographical pattern that has great potential evolutionary importance but that has been largely overlooked in previous research. Phenology, both in plants and in other taxa, is frequently controlled by annual seasonality, even in many perhumid tropical climates where seasonality may be subtle and thus go unnoticed. At

high latitudes, photoperiod and temperature are the predominant seasonal cues of phenology, and are mostly synchronized across broad geographic areas because of their dependence on the fixed orbital geometry of the earth. However, Martin and colleagues (2009) recognized that the most likely seasonal cues at lower latitudes, precipitation and solar radiation, are dependent on factors endogenous to the earth system – atmospheric airflow and its interaction with terrain, at synoptic- and meso-scales – and thus are more subject to being out of sync at relatively short geographic distances (e.g., in a valley versus at altitude, or on one side of a mountain range versus the other). Martin *et al.* used coarse-resolution meteorological data to demonstrate that this spatial asynchrony of seasons is indeed greatest, on average, across the tropics. They then postulated that such seasonal asynchrony could lead nearby populations of a species to exhibit asynchronous breeding phenologies, such that chance migrants between populations would have a lower chance of successful reproduction on average, causing allochronic isolation, reduced gene flow, increased rates of population divergence, and perhaps even increased rates of speciation, helping to generate the latitudinal biodiversity gradient that has long fascinated and puzzled biogeographers. This Asynchrony of Seasons Hypothesis (ASH), as this hypothesis is termed, nonetheless rests on the crucial and as yet untested assumption that high climatic seasonal asynchrony leads to high phenological asynchrony in the tropics, even between sites with similar climatological means, habitats, and thus species whose populations might experience allochronic isolation. In this chapter, I first use time series analysis of satellite imagery to map the characteristic annual phenology of terrestrial vegetation productivity (i.e., land surface phenology; LSP) worldwide. I validate this map extensively, using both a second remote sensing archive and a global set of eddy-covariance flux tower datasets, then explore and interpret it light of prior regional studies on phenology, climatology, and land cover and prior knowledge of global biogeographic patterns. Next, I use that map to map spatial asynchrony of LSP globally, then provide analyses to first explore the potential climatic and physiographic drivers of LSP asynchrony and then test the central assumption of the ASH: that phenological asynchrony occurs more often, and between more climatically similar sites on average, in the tropics.

References

- Hartl, D.L., and Clark, A.G. (2007). *Principles of Population Genetics: Fourth Edition* (Sunderland, Massachusetts: Sinauer Associates, Inc. Publishers).
- Hoffmann, A. A., & Sgrò, C. M. (2011). Climate change and evolutionary adaptation. *Nature*, 470(7335), 479-485.
- Manel, S., Schwartz, M. K., Luikart, G., & Taberlet, P. (2003). Landscape genetics: combining landscape ecology and population genetics. *Trends in ecology & evolution*, 18(4), 189-197.
- Manel, S., & Holderegger, R. (2013). Ten years of landscape genetics. *Trends in ecology & evolution*, 28(10), 614-621.
- Martin, P., Bonier, F., Moore, I. & Tewksbury, J. (2009). Latitudinal variation in the asynchrony of seasons: implications for higher rates of population differentiation and speciation in the tropics. *Ideas in Ecology and Evolution*, 2. <https://doi.org/10.4033/iee.2009.2.3.n>
- Rissler, L. J. (2016). Union of phylogeography and landscape genetics. *Proceedings of the National Academy of Sciences*, 113(29), 8079-8086.

- Slatkin, M. (1987). Gene flow and the geographic structure of natural populations. *Science*, 236(4803), 787–792.
- Storfer, A., Patton, A., & Fraik, A. K. (2018). Navigating the interface between landscape genetics and landscape genomics. *Frontiers in Genetics*, 9, 68.
- Williams, J. W., & Jackson, S. T. (2007). Novel climates, no-analog communities, and ecological surprises. *Frontiers in Ecology and the Environment*, 5(9), 475–482. doi:10.1890/070037

Acknowledgments

The work presented herein is rooted in the influence of so many thoughtful, passionate, and generous scholars at New York University, Hunter College, Stony Brook University, Rutgers University, University of Maryland, and, of course, Cal. It has been shepherded along by my wonderful committee members, without whom it would not have arrived: Dr. Nina Maggi Kelly and Dr. Todd Dawson, whose sage and timely advice strengthened the work at every stage, and Dr. Ian Wang, my PI and chair, whose encouraging patience and earnest advice, since day one, have been the bedrock.

The motivation and inspiration to keep pushing forward came from the great friends and family who somehow continue to tolerate my company:

To Sean: Thank you for telling me to finish my grad school application and then hanging up on me that day; you were the spark.

To Mariko: I don't know how or why you put up with this process for so long. Maybe it's love? Maybe it's that I feed you. At any rate, I sure am glad you did. 今までは、花より団子。これから先は、もっと沢山の花を。

To BB: Thanks for waiting. (... well, you *kinda* waited.) Now that I have more time we'll be sure to name you.

To the other friends and family on whom I have had so little time and energy to spend, and to my poor, sore ass, on which I have spent so much of my time, sitting: I'm sorry, and thank you. No more of this. It's party time!

Lastly, it is important to recognize that these acknowledgments are doomed to fail. How could I ever adequately account for the size and number of my social debts? The artifact I refer to as 'myself' would be nothing without the sculpting influence of the countless others with whom I have shared my moments and thoughts. The best of what I am and have, I owe to their deft and caring hands. The worst, I alone claim, as imperfections in the clay.

Chapter 1

Geonomics: forward-time, spatially explicit, and arbitrarily complex landscape genomic simulations

by

Drew Ellison Terasaki Hart¹

Anusha P. Bishop¹

Ian J. Wang¹

¹ Department of Environmental Science, Policy, and Management, University of California, Berkeley CA, USA

NOTE: This chapter has been published in a refereed journal as: Terasaki Hart, Drew E., Anusha P. Bishop, and Ian J. Wang. "Geonomics: forward-time, spatially explicit, and arbitrarily complex landscape genomic simulations." *Molecular biology and evolution* 38.10 (2021): 4634-4646.). It is my original work, and my co-authors are aware of and in agreement with my use of this work as a dissertation chapter.

Abstract

Understanding the drivers of spatial patterns of genomic diversity has emerged as a major goal of evolutionary genetics. The flexibility of forward-time simulation makes it especially valuable for these efforts, allowing for the simulation of arbitrarily complex scenarios in a way that mimics how real populations evolve. Here, we present Geonomics, a Python package for performing complex, spatially explicit, landscape genomic simulations with full spatial pedigrees that dramatically reduces user workload yet remains customizable and extensible because it is embedded within a popular, general-purpose language. We show that Geonomics results are consistent with expectations for a variety of validation tests based on classic models in population genetics and then demonstrate its utility and flexibility with a trio of more complex simulation scenarios that feature polygenic selection, selection on multiple traits, simulation on complex landscapes, and non-stationary environmental change. We then discuss runtime, which is primarily sensitive to landscape raster size, memory usage, which is primarily sensitive to maximum population size, and other caveats related to the model's methods for approximating recombination and movement. Taken together, our tests and demonstrations show that Geonomics provides an efficient and robust platform for population genomic simulations that capture complex spatial and evolutionary dynamics.

Introduction

Spatial patterns of genomic diversity result from the complex interplay of many underlying ecological and evolutionary processes and are shaped by a wide variety of geographic and environmental factors. Understanding how these patterns develop in natural systems has emerged as a primary goal of modern evolutionary genetics. These systems often occupy complex and potentially changing landscapes and might include populations that are not at demographic equilibrium. They may undergo neutral evolution as well as natural selection, sometimes on multiple traits of variable genetic architecture. The study of complex natural systems is crucial for developing evolutionary and ecological theory (Epperson *et al.* 2010; Barrett *et al.* 2019; Pelletier 2019), understanding the forces governing the evolution and maintenance of genetic diversity (Manel *et al.* 2003; Schoville *et al.* 2012), anticipating ecological futures in the Anthropocene (Bay *et al.* 2018; Capblancq *et al.* 2020), and informing conservation and management (Crossley *et al.* 2017; Lind *et al.* 2017). The complex genomics of many such systems are beyond the reach of analytical population genetics, and their spatial complexity and evolutionary dynamics make them intractable for coalescent simulation (Hoban *et al.* 2012). This hinders not only our understanding of many empirical systems but also our ability to predict their dynamics and, thus, to manage them. Hence, in population and landscape genomics, as in many other fields, forward-time simulation is a crucial tool for dissecting complex study systems.

However, the current suite of forward-time genomic simulators, although numerous, is still of limited utility for such work. Most available software is restricted, either genomically or geospatially, in the complexity it can model. Many programs can model systems of considerable genomic complexity (e.g. simuPOP, Peng and Kimmel 2005; NEMO, Guillaume and Rougemont 2006; QuantiNemo, Neuenschwander *et al.* 2008) yet incorporate only rudimentary spatial components or none at all. Other programs are designed specifically for landscape-genetic simulations (e.g. CDPOP, Landguth and Cushman 2010; CDMetaPOP, Landguth *et al.* 2017; SimAdapt, Rebaudo *et al.* 2013) but are limited in their genomic complexity. For instance, many programs are unable to model simultaneous selection on multiple, polygenic traits. To our knowledge, SLiM (Messer 2013; Haller and Messer 2017; Haller and Messer 2019) is the only package currently capable of simulating scenarios that are sufficiently complex, both genomically and geospatially, to model population genomic patterns emerging under dynamic evolutionary processes (according to a search of the National Cancer Institute's Genetic Simulation Resources website; Peng *et al.* 2013). Its extreme generalizability and complexity allow it to be used for landscape genomics simulation, but it is not explicitly designed for such work. Furthermore, many species are distributed continuously in space, and examining continuous fields of genetic variation can require distinct methods and assumptions (Bradburd and Ralph 2019), yet most population genomic simulation packages are population-based. Such software, SLiM included, requires individuals to be assigned to discrete subpopulations, which can at best be arranged on a high-resolution, regular grid in order to approximate continuously distributed populations.

Here we present Geonomics, a Python package for forward-time, individual-based, continuous-space, population genomic simulations on complex landscapes. Geonomics models are parameterized by way of an informatively annotated parameters file that provides the user a straightforward means of building models of arbitrary complexity while offering reasonable

default settings and ‘off switches’ for parameters and components unrelated to the user’s interests. Models consist of (1) a landscape with one or more environmental layers, each of which can undergo arbitrarily complex environmental change events and (2) one or more species having genomes with realistic architecture and any number of associated phenotypes. Species undergo non-Wright-Fisher evolution in continuous space, with localized mating and mortality, such that species-level phenomena and simulation dynamics are emergent properties of a model’s parameterization. Evolution is comprehensively tracked by way of recently developed data structures that record the complete spatial pedigree (Kelleher *et al.* 2018), providing for the customizable output of rich, three-dimensional datasets in a variety of common formats, including VCF and FASTA for genomic data, GeoTiff for landscape data, and CSV, Shapefile, and GeoJSON for individuals’ non-genomic data (location, environmental values, phenotypes, age, and sex). All of this allows Geonomics to produce realistic landscape genomic results useful for a wide variety of theoretical and empirical purposes.

New Approaches

Model Design: Overview

A Geonomics model consists of two core components: the species and the landscape. The species is composed of a set of individuals and a wide variety of demographic and life-history parameters, including an intrinsic growth rate, mate-search radius, mean number of offspring per mating event, reproductive age, and maximum age, among others. A species can undergo any number of change events, including changes to demographic and life history parameters and various types of population size changes. Each individual in the species has an x,y location, a sex, an age or life-history stage, a set of phenotypes, and a diploid genome consisting of any number of diallelic loci, which can represent either a contiguous haplotype block or a set of distinct loci. Loci can exhibit different types of dominance, and recombination rates can be heterozygous across the genome.

Phenotypic traits are continuous and quantitative, and can be monogenic or multigenic. Each trait is defined by the loci that comprise its genetic basis, the effect sizes of those loci, and a phenotypic selection coefficient, which can be made heterogeneous in both space and time, allowing for spatially complex selection scenarios. While the strength of selection is determined by that coefficient, the force of selection is represented by the environmental raster layer to which the trait is adapting. Loci can have separate mutation rates for three types of mutations: neutral, deleterious, and trait-affecting mutations. Neutral mutations do not affect fitness, and deleterious mutations decrease fitness without affecting simulated phenotypic traits. Trait-affecting mutations, on the other hand, introduce mutations at previously unmutated loci mapped to a trait. This adds to the genetic variation affecting a trait, thus generating phenotypic variance upon which natural selection can operate. Mutation rates can be defined separately for each trait.

The other core component of a Geonomics model, the landscape, is a stack of raster layers. Each layer can be set to serve as one or more of: (1) a resistance raster, which controls individual movement or offspring dispersal, (2) a carrying-capacity raster, which controls population density, and (3) a fitness raster for a trait, which governs natural selection. A key feature of Geonomics is that each layer can undergo any number of arbitrarily complex environmental change events which, as they unfold, influence the dynamics of any species whose carrying capacity, movement and dispersal, or fitness depend on the corresponding layer.

Model Operation: Overview

A Geonomics run begins with a burn-in stage during which individuals move and reproduce, without genomes or selection, until a series of statistical tests is passed. These tests include a time-lagged t-test and an augmented Dickey-Fuller test, which are run as a pair for: (1) the total population size, serving as a test of temporal demographic stability; and (2) both the mean and the standard deviation of timestep-differenced cellwise counts of individuals, serving as a test of spatial demographic stability. This burn-in period results in a stationary spatial distribution of the species on the landscape. Following burn-in, each individual has its genome randomly assigned according to the genomic architecture parameters, such that the main phase of each run begins with no pedigree and, thus, without population structure. Each time step in the main phase is a series of four operations, some optional (Figure 1):

1. **movement** (optional);
2. **mating** (requisite), which includes mate search, mate choice, offspring creation, and offspring dispersal;
3. **mortality**, which is due to density-dependence (requisite) and natural selection (optional);
4. **change events** (optional), including both environmental and demographic changes.

Model Operation: Movement

Movement takes place in continuous space – individuals have x,y coordinates, on either real or simulated landscapes, rather than being arbitrarily restricted to grid cells or bounded populations. Each individual moves along a vector, composed of a distance drawn from a Wald distribution and a direction drawn either from a uniform distribution on the unit circle or from a movement surface – an array of unimodal or multimodal von Mises distributions derived from a landscape layer that serves as a resistance surface (*sensu* McRae 2006; Spear *et al.* 2010). On a unimodal movement surface, each cell is assigned a single von Mises distribution, with mode parameter μ set to the direction of the highest-valued cell in the 8-cell neighborhood. On a multimodal surface, each cell's mixture distribution is a weighted sum of eight such unimodal distributions, one pointing toward the center of each cell in the 8-cell neighborhood and with normalized weights equal to the values of the neighboring cells. This approach to simulating movement generates realistic, anisotropic movement across a heterogeneous landscape (Figure 2) while avoiding time-consuming computational steps, such as repeated searches for minimum-resistance neighboring cells.

Model Operation: Mating

Potential mating pairs are randomly drawn from among all eligible pairs of individuals within the mate-search radius (unless strict nearest-neighbor mating is chosen), with pairing probabilities either uniform or inverse-distance weighted within the mating radius, and with eligibility based on both sex and age. From among those pairs, actual mating-event decisions are Bernoulli distributed, with probability equal to the intrinsic birth rate. Each mating pair produces a number of offspring according to a fixed value or drawn from a Poisson distribution with λ equal to the mean number of offspring. Each parent produces a gamete for each of its offspring, using realistic recombination and Mendelian segregation. Gametes are united to create offspring

individuals, which then disperse to new locations. As with movement, dispersal vectors can be drawn isotropically or anisotropically based on a resistance surface.

Model Operation: Mortality

Mortality is modeled as a Bernoulli process with the probability of an individual death a combination of the probabilities of death due to density-dependence (using a logistic-growth model) and due to natural selection (based on the cumulative fitness for all traits), calculated as:

$$P(d_i) = 1 - \left(1 - P(d_{x,y})\right) \prod_{p=1}^m \omega_{i,p} \quad (1),$$

where $P(d_{x,y})$ is the probability of death due to density-dependence for individual i , m is the number of traits, and $\omega_{i,p}$ is the fitness of individual i for trait p . The probability of density-dependent death at location x,y is calculated as:

$$P(d_{x,y}) = \frac{E[N_{d;x,y}]}{N_{x,y}} = \frac{E[N_{b;x,y}] - \frac{dN_{x,y}}{dt}}{N_{x,y}} \quad (2),$$

where, for location x,y , $E[N_{d;x,y}]$ is the expected number of deaths, $N_{x,y}$ is the population density, $E[N_{b;x,y}]$ is the expected number of births, and $\frac{dN_{x,y}}{dt}$ is the population logistic growth rate. The fitness of individual i for trait p is calculated as:

$$\omega_{i,p} = 1 - \phi_{p;x,y}^{z_{i,p} \gamma_p} \quad (3),$$

where $\phi_{p;x,y}$ is the phenotypic selection coefficient on trait p at location x,y , $e_{p;x,y}$ is the value of the selection layer for trait p at location x,y , $z_{i,p}$ is the phenotype of individual i for trait p , and γ_p defines the curvature of the fitness function for trait p . The phenotype is a result of the additive effects of that individual's genotypes at all underlying loci, and is calculated as:

$$z_{i;p} = \sum_{l=0}^n \alpha_{p,l} g_{i,l} + g_0 \quad (4),$$

where n is the number of loci, $\alpha_{p,l}$ is the effect size of locus l on trait p , $g_{i,l}$ is the genotype at locus l for individual i , and g_0 , the baseline genotype, equals 0 for monogenic traits or 0.5 for polygenic traits.

Model Operation: Change Events

Each demographic or environmental change event unfolds as a series of incremental changes that occur at the ends of scheduled time steps. A demographic change event can be exponential, random, or cyclical, or it can follow an arbitrarily complex, custom trajectory. Each event is parameterized by defining the time steps at which its changes take place and the factor by which the carrying capacity raster is multiplied.

Simple environmental change events are defined by a terminal raster for the final environmental state and a list of time steps at which incremental changes occur (based on cell-wise linear interpolation between the beginning and terminal states). More complex, custom events can be simulated by providing a series of environmental rasters labeled with the time step

at which each will be applied. This option makes it easy to simulate evolution on real-world landscapes undergoing non-linear, spatially heterogeneous environmental change.

Results

Validation

To validate the performance of Geonomics, we ran a series of simulations based on classical population genetic models, covering both neutral and non-neutral evolutionary scenarios. Because the classical models are simpler than the individual-based, spatially explicit, continuous-movement models built by Geonomics, we parameterized the simulations so as to accurately emulate these models while minimizing artefacts (see Validation Testing, Supplementary Information). Our goal was to statistically and heuristically validate Geonomics' full range of functionality and to ensure that it accurately models neutral and non-neutral evolutionary processes.

To verify that Geonomics effectively models neutral evolution, we first examined the average time to fixation for a neutral allele in a finite population using simulations approximating a Wright-Fisher model (Fisher 1923; Wright 1930). We simulated allele-frequency trajectories for 250 independent loci (25 of which are plotted in Figure S1), and we found that that fixation time did, indeed, increase with population size and was proportional to $4N_e$, as expected (Kimura and Ohta 1969), in our simulations (Figure S2). We then tested for changes in the rate of drift surrounding a population bottleneck event by forcing a population to undergo a 70% reduction in size for 50 of 300 timesteps. We found that rates of allele frequency change increased during the bottleneck, then returned to prior levels shortly thereafter (Figure S3). Finally, we quantified the accumulation of genetic structure under a stepping-stone model (Kimura 1953) to certify that genetic covariance decreases with distance (Kimura and Weiss 1964). As expected, we saw that migration rates decreased as a function of inter-island distance (Figure S4), whereas F_{ST} , calculated from both heterozygosity data and genetic variance data, increased as a function of inter-island distance (and, therefore, decreased with pairwise migration rate) and as a function of time (Figures S4 and S5). We also performed a discriminant analysis of principal components (DAPC) using the R package *adegenet* (Jombart *et al.* 2008) to confirm the expected population structure of six island clusters (Figure S6).

To validate the performance of Geonomics for modeling non-neutral evolution, we first performed simulations under a simple scenario of divergent selection between two discrete habitats. As expected, simulations on a landscape evenly divided by two habitat blocks led to local adaptation, producing a significant pattern of phenotype-habitat matching (Figure S7), with mismatches concentrated along the border between habitats. Additionally, over time, the species reached migration-selection equilibrium – the frequencies of the beneficial alleles in each habitat increased up to a stationary level, with that level being positively correlated with the strength of selection (Figure S8). A plot of the mean difference between individuals' phenotypic and environmental values shows a strong decline over model time, with the rate and level of decline increasing as a function of increasing strength of selection (Figure S9). Finally, logistic regressions show no significant relationships between phenotypic and environmental values at the outset (pseudo- R^2 's ≈ 0.0 , p-values > 0.1), but show highly significant relationships at the ends of the simulations ($p < 0.0001$ for all values of ϕ), with the amounts of variation explained

increasing as a function of selection strength (pseudo- $R^2 = 0.327$ for $\phi = 0.01$, 0.376 for $\phi = 0.05$, and 0.406 for $\phi = 0.1$).

We next tested the ability of Geonomics to recreate the genetic structure expected under local adaptation along an environmental cline: monotonic change in the allele frequency of a non-neutral locus across the cline. On a landscape with a symmetric environmental selection gradient, Geonomics again produced the expected spatial pattern of local adaptation (Figure S10), and when we fitted sigmoid *tanh* clines (Szymura and Barton 1986; Porter 2013) for all loci, the locus underlying the monogenic trait was the only one to exhibit clinal variation (Figure S11). In a family of genotype-environment analyses using Bonferroni-corrected, locus-wise logistic regressions, this locus was also the most significantly correlated with the environmental variable based on locus-wise logistic regressions ($p < 0.0001$). A plot of the mean difference between individuals' phenotypic and environmental values shows a strong decline over model time (Figure S12), and logistic regressions show no significant relationship between phenotypic and environmental values at the outset (pseudo- $R^2 = 0$, p -value = 0.812) but a significant relationship at the end of the simulation (pseudo- $R^2 = 0.169$, p -value < 0.0001).

Finally, we verified that Geonomics can effectively model genomic data with physical linkage by simulating a selective sweep, introducing a beneficial mutation at the center of a 101-locus block of otherwise neutral loci. The results exhibited the classic genomic signal of a selective sweep (Przeworski 2002; Kim and Nielsen 2004), with a region of reduced nucleotide diversity surrounding the locus under selection and that region gradually eroding over time (Figure S13). During these simulations, as the beneficial mutant spread through the population the population's mean fitness increased from $1 - \phi$ (where ϕ is the strength of selection) to a saturating value of 1 (Figure S14), confirming that the population dynamics of the selective sweep played out as expected. To further support these results, we also validated Geonomics' method of recombination by examining effective recombination observed in a Geonomics model to those produced by an msprime simulation using the same randomly drawn, heterogeneous recombination map (see Recombination Test, Supplementary Materials); the resulting genome-wide pattern of recombination breakpoint densities recapitulates the one produced by msprime and the true recombination map (Figure S15).

Example Applications: Overview

To demonstrate the broad utility of Geonomics for modeling complex evolutionary scenarios, we performed a series of simulations covering a range of potential applications. These demonstrations highlight scenarios for which Geonomics is particularly well suited, including spatially explicit simulations on highly heterogeneous landscapes, selection on multiple traits with complex genomic architecture, and microevolutionary responses to non-stationary environmental change.

Example 1: Isolation by distance (IBD) and by environment (IBE)

Genetic covariances between individuals or populations are often inversely correlated with linear or resistance-based geographic distance – a pattern known as isolation by distance (IBD; Wright 1943) or isolation by resistance (IBR; McRae 2006; McRae *et al.* 2008) – or with environmental distance – a pattern known as isolation by environment (IBE; Wang and Bradburd 2014). Understanding the landscape factors and population processes generating these patterns

has emerged as a major focus of landscape genetics (Sexton *et al.* 2014; Wang and Bradburd 2014).

To demonstrate how Geonomics can simultaneously generate patterns of IBD and IBE, we built a simulation that uses a heterogeneous resistance layer as a movement surface and models selection for a 10-gene trait on a heterogeneous environmental layer ($\phi = 0.05$). The model features a species with stationary population size (roughly 2450 individuals), experiencing both selection and neutral evolution. The resistance layer consists of a central barrier separating equal-area sides – the barrier has a high resistance to movement, but movement is unconstrained on either side. This layer was also used as the carrying-capacity layer, yielding homogeneous population density on the two sides and zero density within the barrier region. The selection layer consists of two environmental gradients running in opposite directions on either side of the barrier, such that the landscape contains pairs of locations representing a range of combinations of geographic and environmental distances.

To observe the development of population structure, we collected datasets consisting of the genomes for all individuals at timesteps 0 and 1000. We then used principal component analysis (PCA) to calculate pairwise genetic distances between all individuals for each dataset. To visualize population structure, we extracted the first three principal components (PCs) and used them as the red, green, and blue (RGB) color values for mapping individuals on the landscape. To visualize the outcomes of selection, we produced paired maps of the same individuals colored by their phenotypes for the trait under selection (using Geonomics' 'model.plot_phenotype(...)') method), and also created a set of population-structure plots using DAPC. To visualize the time course of the simulation, we plotted the mean phenotype-environment mismatch (i.e. the mean of $|e-z|$, the driving force of selection) and mean fitness. We visualized signals of IBD and IBE in the final dataset using a 3D scatterplot of Euclidean pairwise genetic distance against Euclidean pairwise geographic and environmental distances, colored by pairwise phenotypic distances. We tested the significance of the relationship between genetic distance and environmental distance, controlling for geographic distance, using paired partial Mantel tests with the *vegan* package (version 2.5-6; Oksanen *et al.* 2019) and using multiple matrix regression (MMRR; Wang 2013) in R version 4.0.2 (R Core Team 2020). Finally, because Geonomics models do not use defined landscape-resistance values, we quantified the barrier's increased landscape resistance by tracking all barrier-crossing events, using them to calculate the per-time step crossing rate, then comparing that to the equivalent crossing rate of the same landscape zone in an otherwise identical model that omitted the barrier.

The RGB and phenotype plots of the initial population, with randomly-assigned genomes, showed a clear lack of both spatial structure and local adaptation (Figure 3, top left). However, as expected, spatial structure developed over time, and the species showed signs of local adaptation over the course of the simulation (Figure 3, top right), as well as a corresponding, hierarchical population structure (Figure S15). Average phenotype-environment mismatch decreased and average fitness increased over time (Figure 3, top middle). At the end of the simulation, the species demonstrated significant signals of both IBD (partial Mantel test: $r = 0.560$, $p \leq 0.001$; MMRR: $p \leq 0.001$) and IBE (partial Mantel test: $r = 0.121$, $p \leq 0.001$; MMRR: $p \leq 0.001$; MMRR full model $R^2 = 0.354$), as evidenced by the positive slopes on both horizontal axes of the 3D scatterplot (Figure 3, bottom). The colors of the points in the 3D scatter plot also indicate a clear pattern of increasing phenotypic differences between individuals with increasing

environmental distance (Figure 3, bottom right) but not between individuals separated by increasing geographic distances (Figure 3, bottom left). Finally, the barrier zone had an observed crossing rate of 0.004 individuals per time step in this model, 13 times lower than the rate of 0.052 individuals per time step observed in the barrier-less but otherwise identical model. These results show that Geonomics effectively models IBD and IBE, driven by divergent natural selection, using two simple raster layers. More complex layers could be used to simulate IBD and IBE under a wide range of scenarios, and empirical layers could be used to simulate patterns of spatial genetic variation on real-world landscapes.

Example 2: Simultaneous selection

One of the most powerful features of Geonomics is that it can simulate selection on numerous traits simultaneously, each responding to a separate selection layer. Thus, a simulated species can experience multiple spatial selection regimes. Many natural systems are locally adapted to multiple environmental variables (Fournier-Level *et al.* 2011; Lasky *et al.* 2012; Manel *et al.* 2012), so simulating these scenarios could be broadly valuable for investigating the nature of local adaptation in real environments.

To demonstrate how Geonomics can model simultaneous selection, we simulated a scenario in which a species undergoes natural selection along two orthogonal environmental gradients, each driving selection for a separate trait ($\phi = 0.05$). Each trait had values ranging from 0 to 1, determined by 10 unlinked loci, all with equal effect sizes. Individuals had a mean movement distance of 0.5 cell widths on a 50 x 50-cell landscape, chosen to limit gene flow and allow for the development of strong spatial structure and, thus, the potential for local adaptation. We let the system evolve for 1000 time steps and then mapped the species on each of the environmental layers, with individuals colored by phenotype in order to visually evaluate whether individual phenotypes matched their environmental backgrounds. The results showed clear patterns of phenotype-environment matching along both independent gradients (Figure 4) that evolved steadily through time (Figure S16; compare to figures S17 and S18, with $\phi = 0$), indicating strong evidence for simultaneous selection across the simulated landscape.

Example 3: Polygenic adaptation to climate change in the Yosemite region

Better understanding evolutionary responses to changing environments is essential for predicting species outcomes and preserving biodiversity under ongoing climate change (Hoffmann and Sgrò 2011; Franks and Hoffmann 2012; Bay *et al.* 2018; Capblancq *et al.* 2020). In many regions, climate shifts are projected to be spatially heterogeneous, including in montane regions where cooler, higher-altitude areas are warming more quickly than warmer, low-altitude regions (Rangwala *et al.* 2013; Mountain Research Initiative EDW Working Group 2015; but see Oyler *et al.* 2015). Of particular interest under these scenarios is the ability of species to adapt to changing local conditions (Franks and Hoffmann 2012).

To demonstrate the utility of Geonomics for studying microevolutionary responses to climate change, we simulated the response to projected climate change of a continuously distributed, locally adapted species, using the sagebrush lizard (*Sceloporus graciosus*) in the topographically complex Yosemite National Park region of California (U.S.A.) as an empirical model. To model climate change, we assembled time series raster stacks of projected mean annual temperature, annual precipitation, and habitat suitability for 19 even time steps from the

present through the year 2100. For present temperature and precipitation, we used PRISM data (Daly *et al.* 2008), calculated as 30-year normals for 1981-2010 at 800 meter resolution. For future years, we used means at a set of 5-year intervals (2015-2100), downscaled to 6 km resolution using the localized constructed analogs downscaling technique (LOCA; Pierce *et al.* 2014), from the Cal-Adapt database (<https://cal-adapt.org/>). We calculated means of both variables from their minima and maxima observed across 32 global climate models, using a conservative representative concentration pathway (RCP 4.5). We developed time series of future temperature and precipitation layers at 800-m resolution by: (1) calculating the raster difference between the first projected year and the current data, aggregated to the projected data's resolution; (2) adding that difference to the current data, such that each cell in the current data received the difference of the coarser, projected cell within which it lay; and (3) repeating that process for all remaining years. All data preparation was done using custom scripts (Supplementary Information) in R (R Core Team 2019).

For the habitat suitability rasters, we constructed a species distribution model (SDM) using the present day temperature and precipitation variables. We downloaded all georeferenced *S. graciosus* occurrence data from the Global Biodiversity Information Facility database (www.gbif.org), using the *gbif* function in the *dismo* R package (Hijmans *et al.* 2017). We clipped the points to California and Nevada, then subsampled the full dataset to remove multiple points within the same raster cells. We generated pseudoabsence data by drawing random points from all cells in the California-Nevada region where the species was not observed (following the recommendations of Barbet-Massin *et al.* 2012). We extracted the current temperature and precipitation data at these points and used them as predictor variables in a binomial generalized linear model (GLM) with a logit link (Elith and Leathwick 2009). We then projected that GLM onto the current and future temperature and precipitation rasters for our study region, producing a time series of predicted habitat suitability.

We generated the simulation's parameters file using the code provided in Code Sample S1, then edited the parameter values therein as needed. To simulate the non-neutral evolution of a polygenic, quantitative trait, we set the trait to be underlain by 100 loci randomly distributed across a genome of 1000 unlinked loci and set a strength of selection of $\phi = 0.5$. We set other life-history and demographic parameters (carrying capacity, age at reproductive maturity, number of offspring per individual, maximum age) to reasonable values based on *S. graciosus* natural history (Stebbins 1948; Tinkle 1973; Rose 1976; Ruth 1978; Tinkle *et al.* 1993; Supplementary Information).

We ran the main phase for 500 time steps without climate change (to develop a pattern of local adaptation), then ran an additional 100 time steps (years) with changing climate (see Code Sample S2). At time steps 500 (before the initiation of climate change events) and 600 (after completion of climate change events), we plotted the current temperature and habitat-suitability landscape layers along with a kriged surface of the current population's phenotypes and a kernel density map of the current population's density, two key emergent properties of the model that should be driven by temperature and habitat suitability, respectively. We then ran the model for an additional 50 timesteps to be able to more clearly visualize the effect of climate change on population size.

The model generated a clear and realistic pattern of adaptation to the spatial temperature gradient in the Yosemite region after the 500 iterations following burn-in, and that pattern

demonstrated a spatial shift in phenotypes that aligns clearly with the spatial shift in temperature under the simulated climate change scenario (Figure 5A, rows 1 and 2; Video 1, Supplementary Information). The model also generated a spatial pattern of population density that clearly aligns with spatial variation in habitat suitability prior to the onset of climate change that likewise shifted as expected in response to the climate change-induced shift in habitat suitability (Figure 5A, rows 3 and 4). We observed demographic changes in response to climate change over the course of the simulation as well. After climate change, mean population size was reduced by roughly 16.8% (from about 255,500 to 212,500 individuals), in line with the 17.9% reduction in the carrying capacity layer derived from the habitat suitability rasters (from 337,089.0 to 276,742.8 individuals, according to sums of the pre- and post-change carrying capacity layers). The population also exhibited sizable fluctuations during the climate change period, with oscillations exceeding 60,000 individuals (roughly 23.5% of the pre-change mean population size; Figure 5B). We interpret this as a result of the stepwise environmental changes comprising the climate change event. Each change causes a shift in the optimum phenotypes of local populations, leading to increased maladaptation and thus increased mortality rates. Subsequent reductions in density-dependent mortality rates because of these reduced population densities, paired with adaptation by natural selection, then reduce overall mortality rates, leading to rebounds in population size, with stochastic movement into and out of local populations, along with other sources of model stochasticity, imposing noise on this oscillatory behavior. Overall, these results show how Geonomics can effectively simulate organismal responses to highly complex environmental scenarios and reveal that these simulations can uncover system behavior that could provide avenues for future investigation.

Discussion

Our validation tests demonstrate that Geonomics simulates molecular evolution in concordance with predictions from theoretical population genetics (Fisher 1923; Wright 1943; Kimura; Szymura and Barton 1983), including dynamics of genetic drift, migration, and selection along clines, and our example applications show that Geonomics is capable of generating accurate and realistic population and landscape genomic datasets under scenarios of varying complexity. Geonomics is embedded in Python (van Rossum 1995; Python Software Foundation 2019), one of the most popular programming languages and one already familiar to many researchers who use bioinformatics. It makes the creation of arbitrarily complex models quick and easy, without even requiring prior Python experience, yet provides advanced users with access to core data structures, enabling broad customization and extension.

Many theoretical questions in population genomics necessitate explicitly spatial study methods, often with full tracking of a population's spatial pedigree (Bradburd and Ralph 2019). Geonomics makes this work more tractable than ever before. Landscape genomics studies draw conclusions about complex, real-world systems, sometimes with direct implications for conservation and management (Epperson *et al.* 2010; Landguth *et al.* 2012). Geonomics not only enables the generation of simulated datasets specific to such study systems, but it will also aid the development and testing of analytical methods in landscape genomics, strengthening our ability to draw accurate and reliable inferences from real-world data.

Runtime and memory

Geonomics models run more slowly and have steeper memory limitations than models written and optimized in compiled languages, such as SLiM (Messer 2013; Haller and Messer 2017; Haller and Messer 2019). However, for users whose scenarios are well served by the design and affordances of Geonomics, what is sacrificed in runtime will be made up for in flexibility, customizability, and ease of use. With a reasonably powerful computer and for moderately sized models, most users should not find runtime or memory a major limitation. Indeed, our first two example applications were run on a laptop computer with 8 GB of RAM and an Intel® Core™ i5-8250U 3.4 GHz quad-core processor. Each run took an average of 271 s (about 0.27 s per time step) for the IBD-IBE model and 144 s (roughly 0.14 s per time step) for the simultaneous selection model. Because the polygenic adaptation example has much higher complexity, approximating the high population density of a small vertebrate, we ran it on a regular-memory node of a computing cluster (the *savio3* partition of UC Berkeley's Savio system) with 96 GB RAM and 2.1 GHz Skylake processors in order to handle the larger memory requirement. This model took considerably longer to run (approximately 7.35 hr, running at about 32 s per time step after time for upfront computation of a series of changing movement surfaces) and had a peak memory usage of 25.433 GB, but even highly complex scenarios like this remain tractable on reasonable research timelines.

Given the complexity of Geonomics and the number of parameters a user can modify, numerous parameters and parameter combinations can influence a model's average runtime. We provide a basic runtime analysis (Figure 6), run on the same 8Gb, quad-core laptop as the examples above. This analysis highlights some basic parameters that are likely to influence a model's average runtime per time step, including the mean population size (as determined by an array of local carrying capacities), the number of offspring per mating event, the size of the landscape, and the number of non-neutral loci in the genome. The effect of landscape size predominates, as runtime scales superlinearly with this parameter. The number of non-neutral loci actually has only slight effects on total runtime, and because neutral genetic data is stored in a set of tables rather than redundantly for each individual (Kelleher *et al.* 2018) runtime is even less sensitive to the number of neutral loci in the genome, meaning that Geonomics can efficiently simulate genome-scale datasets if provided adequate memory to store the set of recombination pathways that is calculated at the outset. Finally, after the upfront cost of computing recombination pathways and movement surfaces, runtime scales roughly linearly with the number of time steps, barring large demographic changes. That means that the moderately complex scenario in our simultaneous selection example could complete 1,000,000 time steps in approximately 40 hours, and even the highly complex scenario in our polygenic adaptation example with more than 200,000 individuals could run through 10,000 generations in about 3.7 days. Hence, Geonomics could even prove useful for research at deeper timescales, for example in phylogeography or geogenomics (Baker *et al.* 2014).

Caveats

Geonomics uses two unconventional approximations to make complex models tractable within reasonable compute time in an interpreted language. The first is the approximation used to model heterogeneous recombination. Enacting recombination between all neighboring loci each time a gamete is produced would require an extremely large and time-consuming number of

random draws. To avoid this, when a model is first created, Geonomics generates and saves, as binary arrays, a large collection of recombination ‘paths.’ The number of paths used is set by the user, and directly determines the minimum recombination rate difference that can be modeled. Each path is just a genome-length array that switches between 0 and 1 at each inter-locus position where a recombination event should occur. The path can then quickly be used to subset an individual’s genome, producing a gamete. As a model runs and gametes are continually produced, these paths are repeatedly shuffled and drawn through, like a deck of cards during multiple rounds of a game. This approach, which we have validated using *msprime* (see Recombination Test, Supplementary Information), can lead to memory limitations for models with a large number of paths and a long genome, because the data structure containing the paths is essentially a two-dimensional binary array whose size is the product of these values. To avoid these problems, for genomic architectures with homogeneous recombination rates, Geonomics provides the option to use an alternative recombination mechanism that simulates recombination on the fly for each new gamete but does so at a cost of increased average runtime per time step.

The second is the approximation used to model the circular distributions from which movement directions are drawn. Conceptually, a movement or dispersal surface is an $x \times y$ array of Von Mises distributions. In practice, each distribution on that surface is represented by a column of angular directions (an ‘approximation column’) drawn, at the time the model is built, from the true, continuous distribution. During a model run, to draw a movement direction from a cell, a random value is sampled from that cell’s approximation column. This increases computational efficiency by avoiding large numbers of calls to random number generators during runtime. The accuracy of these approximation columns is a function of their length, which is set by the user. This length will usually not be so constrained that it significantly impacts the accuracy of the approximation, but such a constraint could arise if the movement or dispersal surface undergoes environmental change. In this case, the movement surfaces corresponding to each step of the change event will be generated and stored when the model is first created, and the series of arrays produced could exhaust memory if the landscape is very large and has many environmental change steps. A solution to this problem would feature some combination of decreasing the temporal resolution of the environmental change event, decreasing the landscape size, or decreasing the approximation column length. In all cases, users may check the accuracy of modeled movement by using built-in functions that visualize the composition and behavior of movement and dispersal surfaces.

Installation

The simplest way to get started with Geonomics is to install it via *pip*. Geonomics uses common, well-established Python packages as required dependencies — Numpy (Harris *et al.* 2020), Matplotlib (Hunter 2007), Pandas (McKinney *et al.* 2010), Shapely (Gillies *et al.* 2007), Scipy (Virtanen *et al.* 2020), Scikit-learn (Pedregosa *et al.* 2011), Statsmodels (Seabold and Perktold 2010), Rasterio (Gillies *et al.* 2019), Bitarray (Schnell *et al.* 2011), *msprime* (Kelleher *et al.* 2016), and *tskit* (Kelleher *et al.* 2018) — and offers optional integration of neutral landscape models through the *NLMpy* package (Etherington *et al.* 2015). The source code is publicly available on GitHub (<https://github.com/drewhart/geonomics>), where it is actively maintained and developed.

Conclusions

Geonomics is a Python package designed to make building and running complex landscape genomic models quick and simple. At the same time, it provides a flexible scripting framework that allows advanced users to customize and extend its functionality. We believe Geonomics will prove highly useful for theoretical, empirical, methodological, and applied research in population and landscape genomics, molecular ecology, global change biology, and conservation.

Acknowledgments

We thank G. Bradburd, T. Dawson, N. Fefferman, M. Fitzpatrick, J. Frederick, N. Graham, M. Kelly, M. McElroy, R. Nielsen, T. O’Connor, D. O’Sullivan, C. Small, E. Westeen, N. Whiteman, G. Wogan, and M. Yuan for feedback and guidance on various iterations of the model that became Geonomics. We also thank I. R. Wang, B. Xu, J. Zhao for contributing to the Python code in its early stages and Berkeley Research Computing for access to the Savio computing cluster. Lastly, we thank M. Terasaki, C. Hart, G. Hart, J. Hart, and M. Tylka for supporting and encouraging a lifetime of curiosity about nature, and C. Adderly, K. Cobain, C. Debussy, L. Hill, S. Jeffes, N. Jones, G. R. Sessions, and T. Yorke for their continual company during many long hours of programming. This work was supported by a Berkeley Fellowship (to DETH) and NSF Award #1845682 (to IJW).

Author Contributions

DETH conceived and wrote the software, validations, and example applications, with contributions from IJW. DETH co-wrote the manuscript and performed analyses. APB co-wrote the manuscript and contributed analyses. IJW co-wrote the manuscript.

Data Availability

Geonomics is installable from the Python Package Index (<https://pypi.org/project/geonomics/>) using pip, and the Geonomics source code is available in the Geonomics GitHub repository (<https://github.com/drewhart/geonomics>). Parameterizations for all simulations run for this study are stored in their respective directories in the source code repository, and other associated data (e.g. raster layers) are available in the ./data directory of a local Geonomics installation. Ancillary code used for the preparation of simulations and execution of analyses for this paper live in a separate GitHub repository (https://github.com/drewhart/geonomics_methods_paper_ancillary_code). Validations tests can be run using the code provided in the respective directories in the source code repository. Demo simulations can be run by loading the Geonomics package into Python and then calling the run_demo function (e.g. ``import geonomics as gnx; gnx.run_demo("IBD IBE")``). Using those materials, readers can reproduce stochastically varying but qualitatively equivalent results as those presented in this paper.

Figures

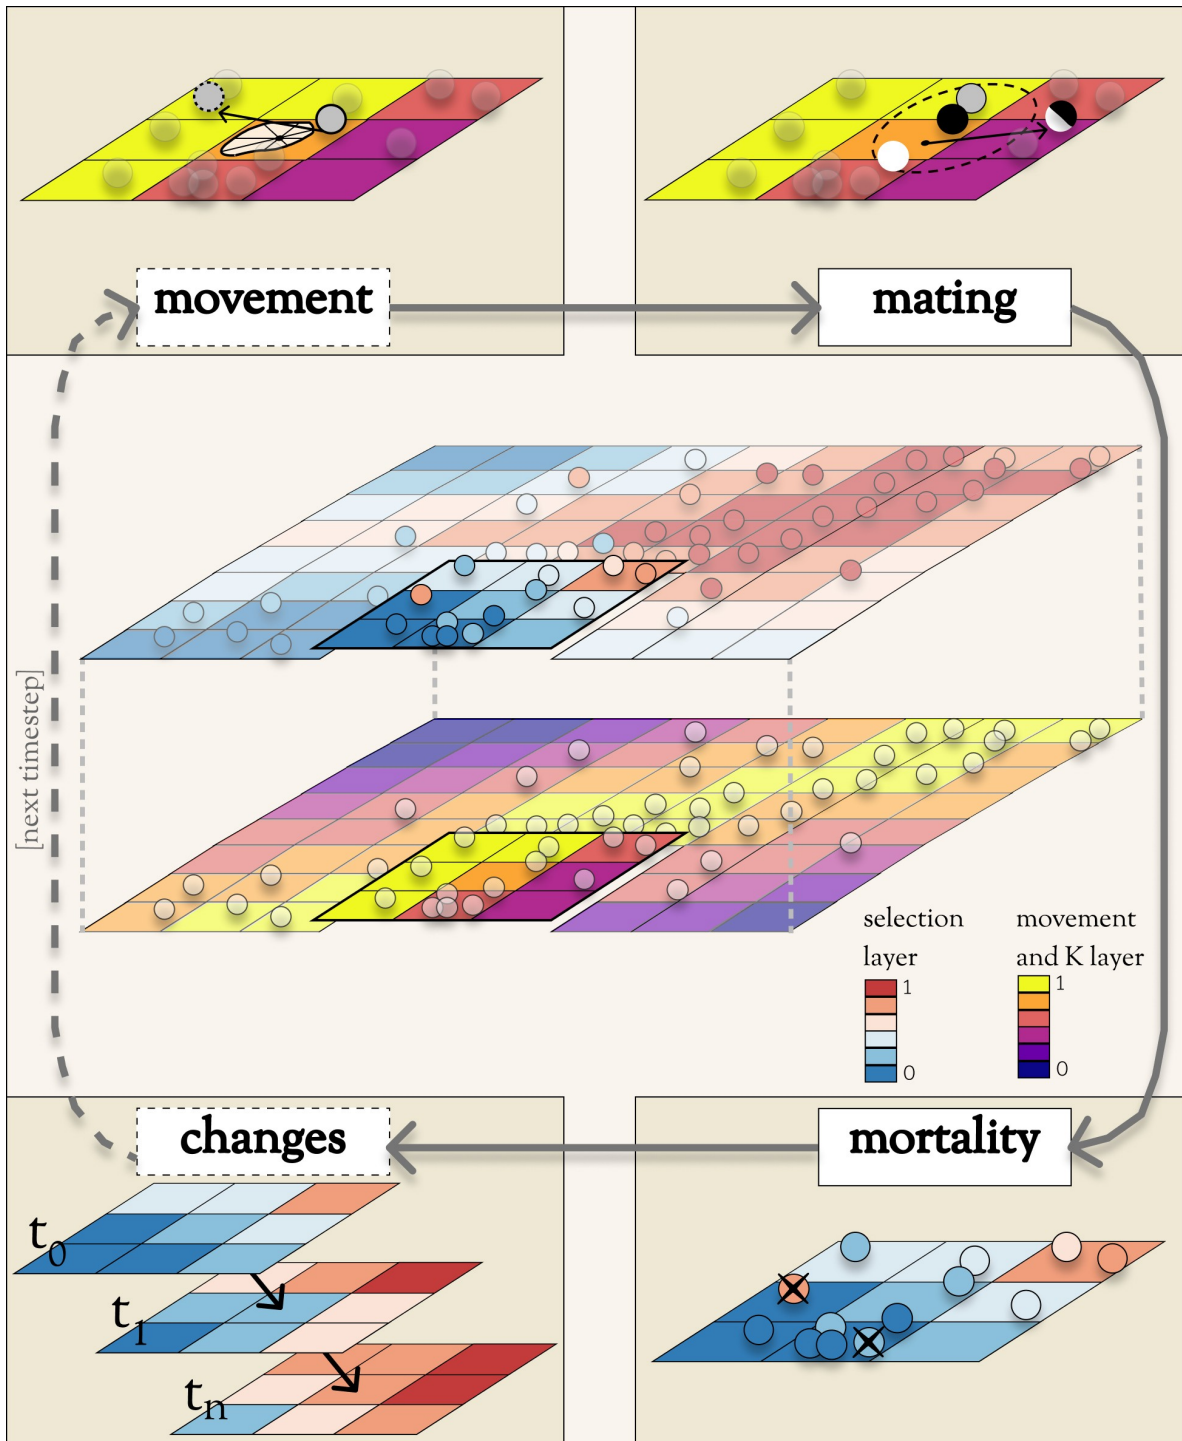


Figure 1: Operations during the main phase of a Geomatics model run. In the center is a species on a multi-layer landscape that includes a selection layer (above) and a layer for movement and carrying capacity (below). Surrounding the landscape is a flow-diagram of the major operations

during a time step. Operations in dashed boxes are optional. During the movement stages (top-left), individuals move along movement vectors drawn from various distribution options (shown is an example of a cell-specific von Mises mixture distribution). During the mating stage (top-right), each mating individual (black circle) randomly chooses a mate (white circle) from all potential mates within its mating radius (dashed circle). The resulting offspring (half-black, half-white circle) disperses from its parents' midpoint along a randomly drawn dispersal vector. During the mortality stage (bottom-right), deaths are modeled as a Bernoulli process, with the probability of mortality a product of density-dependence and selection on all traits. During the changes stage (bottom-left), environmental and demographic change events, which can be represented by a series of change rasters corresponding to scheduled time steps, take place.

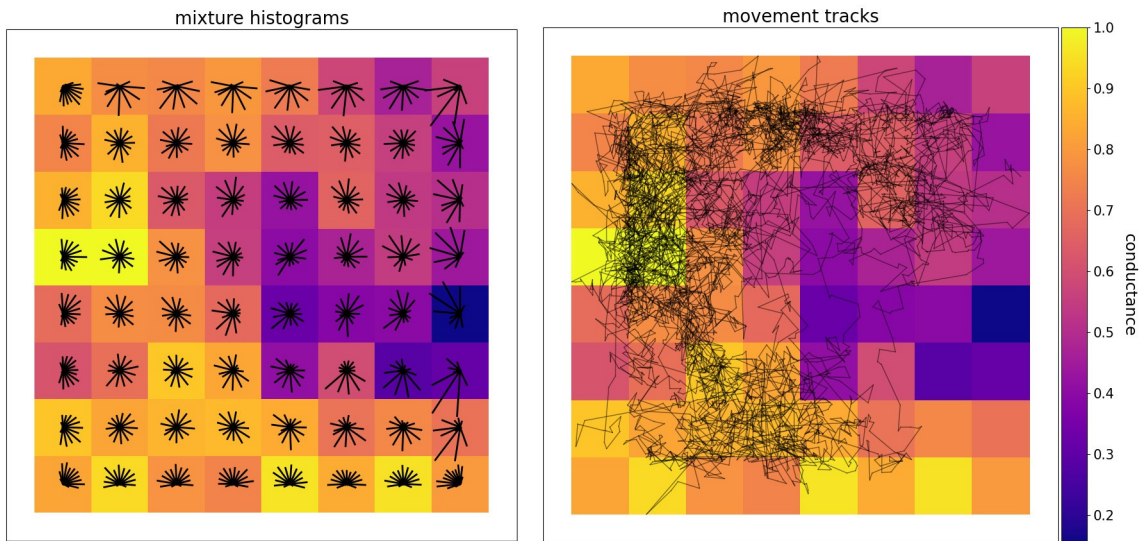


Figure 2: A raster layer representing a movement surface with example movement histograms for each cell (left) and a movement track for a sample individual (right). The circularized histograms represent the movement directions that could be drawn from the von Mises mixture distribution approximations within each cell. Longer bars in a histogram indicate higher probability of movement in their direction. The movement track, plotted with the *gnx.help.param_help.plot_movement* function in Geomatics, is 5000 steps long. Both preferential movement toward higher-suitability regions of the landscape (i.e. cells closer to 1 in value) and occasional long-distance movements between relatively isolated portions of the landscape are evident.

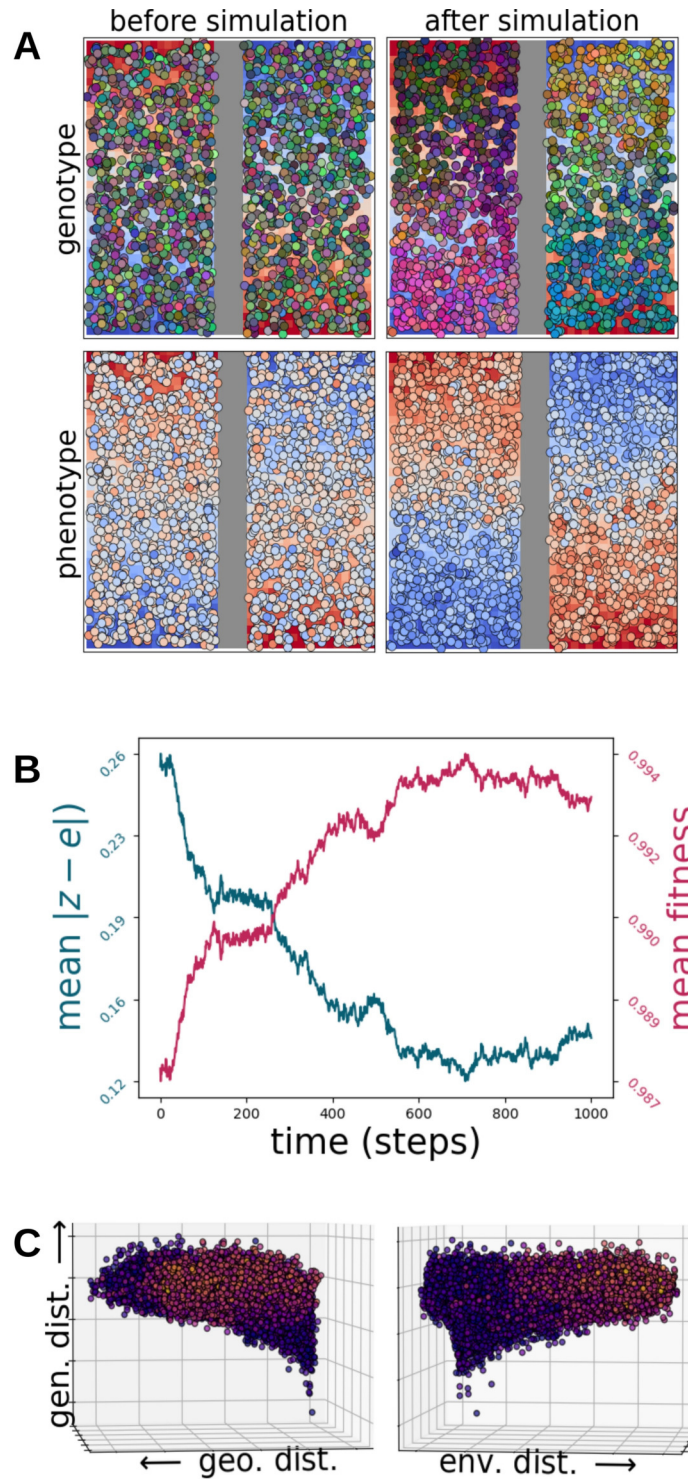


Figure 3: Results of simulations for the isolation by distance (IBD) and isolation by environment (IBE) example application, in which a species evolved on a landscape with a barrier layer that served as the movement surface (displayed as a vertical gray band down the landscape) and an

environmental layer that served as the selective surface for a 10-locus trait (displayed as the red to blue gradient on the landscape). Panel A shows the population before the simulation (left column) and after it (right column), colored by genetic distance (top row), with colors derived from scores on the first three PCs of a genetic PCA used to assign RGB values, and by phenotype (bottom row). The most-fit individuals are those whose phenotypic colors perfectly match the cells on which they are located. Panel B displays the time courses of the mean difference between individuals' phenotypes and their environmental values (blue) and of mean fitness values (red). Panel C shows two views of a 3D scatter plot of pairwise genetic distance as a function of Euclidean geographic distance (left) and Euclidean environmental distance (right), with points colored by phenotypic distance.

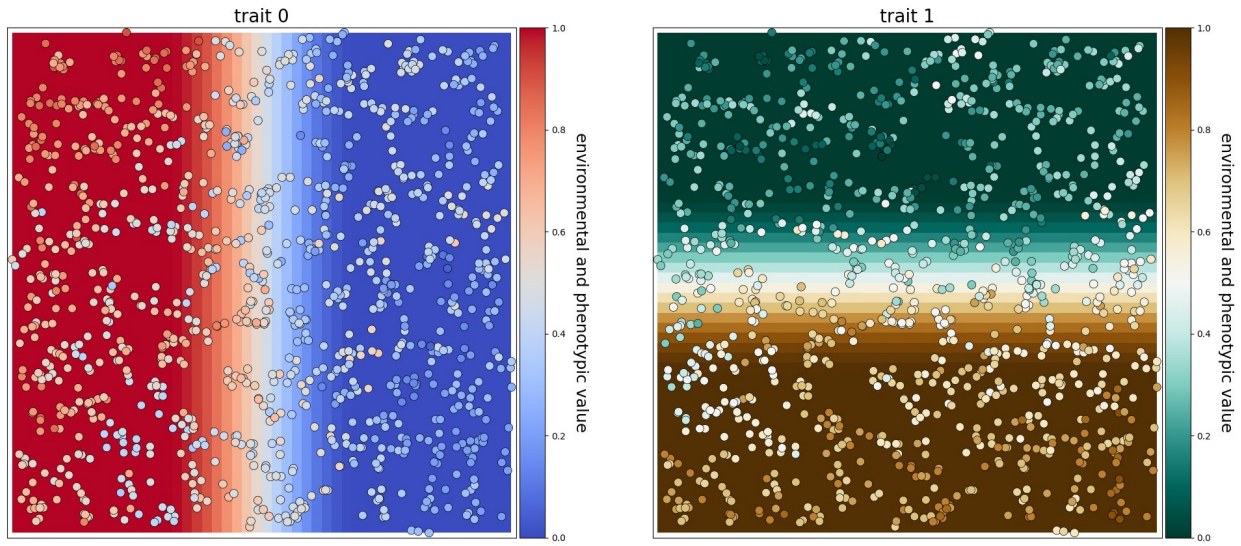


Figure 4: Results of simultaneous selection on two traits with spatially distinct selective regimes. Each trait is controlled by 10 unlinked loci and has a selection coefficient of $\phi = 0.05$. Individuals are colored by phenotype for the trait under selection on each layer.

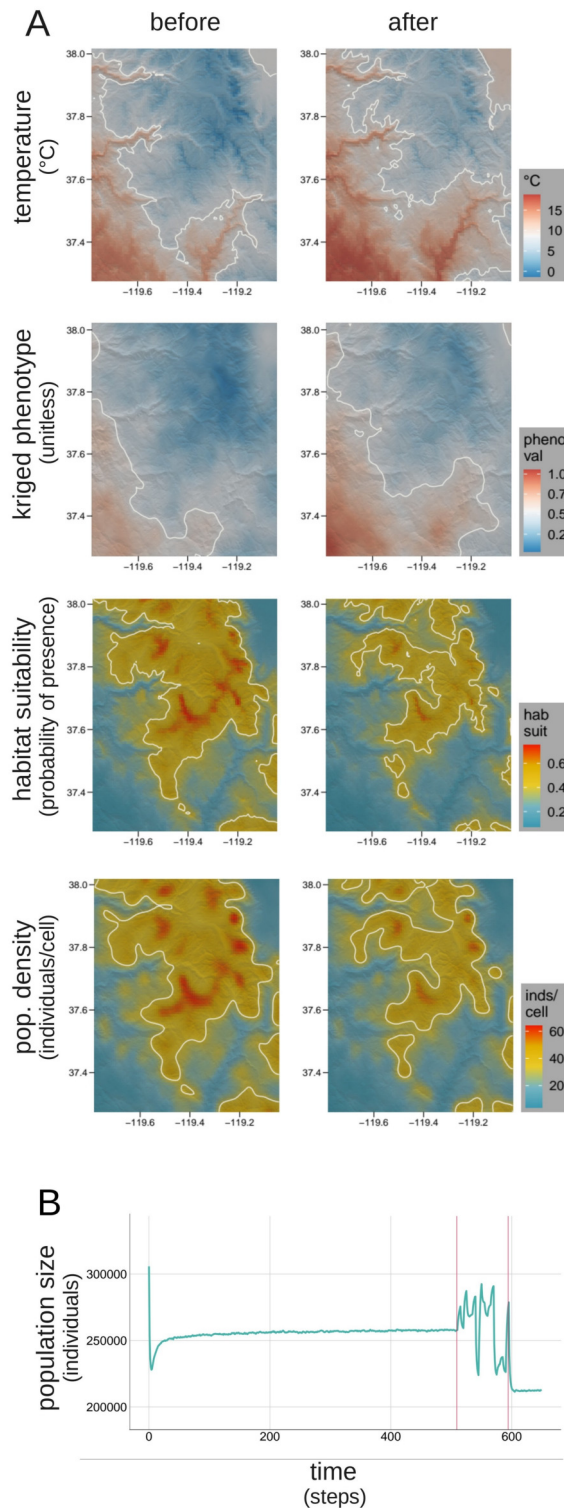


Figure 5: Polygenic adaptation to climate change in the Yosemite region. Panel A shows hillshade plot comparisons of key variables (mean temperature, phenotype, habitat suitability, and population density) before and after the simulated climate change event. The mean of each

variable through time is used to draw mid-value contours on each map (white lines) to help visualize spatial change. As expected, the spatiotemporal shift in temperature (first row) drives a spatially corresponding shift in phenotypes, visualized as a surface kriged from all phenotypic values (second row), and the shift in habitat suitability (third row) likewise drives a corresponding shift in population density, visualized using a 2D kernel density estimator (fourth row). Panel B shows the time course of population size. The early drop in population size results from the onset of natural selection after completion of the unplotted burn-in portion of the model. The oscillations and ultimate reduction at the end of the simulation are a result of the climate change event, which occurs during the period bracketed by vertical red lines.




Figure 6: Average runtime as a function of four major parameters: the number of non-neutral loci (n_loci), the carrying-capacity constant (K) that determines mean population size, landscape size (dim), and the λ parameter of the Poisson distribution from which the number of offspring in each mating event is drawn ($lambda$). Runtime increases with landscape dimension both because of functions whose runtimes scale with the landscape size directly and because of functions whose runtime scales with total population size. The difference between the lines for K and dim can be taken as an indication of the runtime cost of landscape dimension above and beyond population-size effects, which predominates because it is superlinear.

References

- Barbet-Massin M, Jiguet F, Albert CH, Thuiller W. 2012. Selecting pseudo-absences for species distribution models: how, where and how many? *Methods Ecol Evol.* 3:327–338.
- Barrett RDH, Laurent S, Mallarino R, Pfeifer SP, Xu CCY, Foll M, Wakamatsu K, Duke-Cohan JS, Jensen JD, Hoekstra HE. 2019. Linking a mutation to survival in wild mice. *Science* 363:499–504.
- Bay RA, Harrigan RJ, Underwood VL, Gibbs HL, Smith TB, Ruegg K. 2018. Genomic signals of selection predict climate-driven population declines in a migratory bird. *Science* 359:83–86.
- Bradburd GS, Ralph PL. 2019. Spatial Population Genetics: It's About Time. *Ann Rev Ecol Evol S.* 50:427-449.
- Capblancq T, Fitzpatrick MC, Bay RA, Exposito-Alonso M, Keller SR. 2020. Genomic Prediction of (Mal)Adaptation Across Current and Future Climatic Landscapes. *Ann Rev Ecol Evol S.* 51:245-269.
- Crossley MS, Chen YH, Groves RL, Schoville SD. 2017. Landscape genomics of Colorado potato beetle provides evidence of polygenic adaptation to insecticides. *Mol Ecol.* 26:6284–6300.
- Daly C, Halbleib M, Smith JI, Gibson WP, Doggett MK, Taylor GH, Curtis J, Pasteris PP. 2008. Physiographically sensitive mapping of climatological temperature and precipitation across the conterminous United States. *Int J Climatol.* 28:2031–2064.
- Epperson BK, McRae BH, Scribner K, Cushman SA, Rosenberg MS, Fortin MJ, James PMA, Murphy M, Manel S, Legendre P, and Dale MRT. 2010. Utility of computer simulations in landscape genetics. *Mol Ecol.* 19:3549–3564.
- Etherington TR, Holland EP, O'Sullivan D. 2015. NLMpy: a python software package for the creation of neutral landscape models within a general numerical framework. *Methods Ecol Evol.* 6: 164-168.
- Fisher RA. 1923. XXI.—On the dominance ratio. *Proc R Soc Edinb.* 42:321–341.
- Fournier-Level A, Korte A, Cooper MD, Nordborg M, Schmitt J, Wilczek AM. 2011. A Map of Local Adaptation in *Arabidopsis thaliana*. *Science* 334:86–89.
- Franks SJ, Hoffmann AA. 2012. Genetics of Climate Change Adaptation. *Annu Rev Genet.* 46:185–208.
- Gillies S, Shapely contributors. 2007. Shapely: manipulation and analysis of geometric objects. URL: <https://github.com/Toblerity/Shapely>
- Gillies S, Rasterio contributors. 2019. Rasterio: geospatial raster I/O for Python programmers. URL: <https://github.com/mapbox/rasterio>
- Guillaume F, Rougemont J. 2006. Nemo: an evolutionary and population genetics programming framework. *Bioinformatics* 22:2556–2557.
- Haller BC, Messer PW. 2017. SLiM 2: Flexible, Interactive Forward Genetic Simulations. *Mol Biol Evol.* 34:230–240.
- Haller BC, Messer PW. 2019. SLiM 3: Forward Genetic Simulations Beyond the Wright–Fisher Model. Hernandez R, editor. *Mol Biol Evol.* 36:632–637.
- Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, et al. 2020. Array programming with NumPy. *Nature* 585: 357–362.

- Hartl, D.L., and Clark, A.G. (2007). *Principles of Population Genetics: Fourth Edition* (Sunderland, Massachusetts: Sinauer Associates, Inc. Publishers).
- Hijmans RJ, Phillips S, Leathwick J, Elith J. 2017. *dismo*: Species Distribution Modeling. R package version 1.1-4. URL: <https://CRAN.R-project.org/package=dismo>
- Hoban S, Bertorelle G, Gaggiotti OE. 2012. Computer simulations: tools for population and evolutionary genetics. *Nat Rev Genet.* 13:110–122.
- Hoffmann AA, Sgrò CM. 2011. Climate change and evolutionary adaptation. *Nature* 470:479–485.
- Hunter JD. 2007. Matplotlib: A 2D Graphics Environment. *Compu Sci Eng.* 9(3): 90-95.
- Jombart T, Devillard S, Balloux F. 2010. Discriminant analysis of principal components: a new method for the analysis of genetically structured populations. *BMC Genetics* 11(1): 94. DOI:1186/1471-2156-11-94
- Jones E, Oliphant E, Peterson P, *et al.* (2001). SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>.
- Kelleher J, Etheridge AM and McVean G. 2016. Efficient Coalescent Simulation and Genealogical Analysis for Large Sample Sizes. *PLoS Comput Biol.* 12(5): e1004842. DOI:10.1371/journal.pcbi.1004842
- Kelleher J, Thornton KR, Ashander J, Ralph PL. 2018. Efficient pedigree recording for fast population genetics simulation. *PLoS Comput Biol.* 14(11): e1006581. DOI: 10.1371/journal.pcbi.1006581
- Landguth EL, Bearlin A, Day CC, Dunham J. 2017. CDMetaPOP: an individual-based, eco-evolutionary model for spatially explicit simulation of landscape demogenetics. *Methods Ecol Evol.* 8:4–11.
- Landguth EL, Cushman SA. 2010. cdpop: A spatially explicit cost distance population genetics program. *Mol Ecol Resour.* 10:156–161.
- Landguth EL, Cushman SA, Johnson NA. 2012. Simulating natural selection in landscape genetics. *Mol Ecol Resour.* 12:363–368.
- Lasky JR, Des Marais DL, McKAY JK, Richards JH, Juenger TE, Keitt TH. 2012. Characterizing genomic variation of *Arabidopsis thaliana*

- Mountain Research Initiative EDW Working Group. 2015. Elevation-dependent warming in mountain regions of the world. *Nat Clim Change* 5:424–430.
- Neuenschwander S, Hospital F, Guillaume F, Goudet J. 2008. quantiNemo: an individual-based program to simulate quantitative traits with explicit genetic architecture in a dynamic metapopulation. *Bioinformatics* 24:1552–1553.
- Oksanen J, Guillaume Blanchet F, Friendly M, Kindt R, Legendre P, McGlenn D, Minchin PR, O'Hara RB, Simpson GL, Solymos P, Stevens MHH, Szoecs E, Wagner H. 2019. vegan: Community Ecology Package. R package version 2.5-6. <https://CRAN.R-project.org/package=vegan>
- Oyler JW, Dobrowski SZ, Ballantyne AP, Klene AE, Running SW. 2015. Artificial amplification of warming trends across the mountains of the western United States. *Geophys Res Lett.* 42:153–161.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, et al. 2011. Scikit-learn: Machine Learning in Python. *J Mach Learn Res.* 12: 2825-2830.
- Pelletier F. 2019. Testing evolutionary predictions in wild mice. *Science* 363:452–453.
- Peng B, Kimmel M. 2005. simuPOP: a forward-time population genetics simulation environment. *Bioinformatics* 21:3686–3687.
- Peng B, Chen HS, Mechanic LE, Racine B, Clarke J, Clarke L, Gillanders E, Feuer EJ. Genetic Simulation Resources: a website for the registration and discovery of genetic data simulators. *Bioinformatics.* 2013 Apr 15;29(8):1101-2. DOI:10.1093/bioinformatics/btt094.
- Pierce DW, Cayan DR, Thrasher BL. 2014. Statistical Downscaling Using Localized Constructed Analogs (LOCA). *J Hydrometeorol.* 15:2558–2585.
- Porter A. 2013. ClineFit v. 2.0, User's Manual. URL: <http://people.umass.edu/aporter/ClineFit/ClineFit%20Manual.pdf>
- Python Software Foundation. 2019. Python Language Reference, version 3.7. Available at <http://www.python.org>.
- R Core Team. 2020. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL: <https://www.R-project.org/>.
- Rangwala I, Sinsky E, Miller JR. 2013. Amplified warming projections for high altitude regions of the northern hemisphere mid-latitudes from CMIP5 models. *Environ Res Lett.* 8:024040.
- Rebaudo F, Le Rouzic A, Dupas S, Silvain J-F, Harry M, Dangles O. 2013. SimAdapt: an individual-based genetic model for simulating landscape management impacts on populations. Spencer M, editor. *Methods Ecol Evol.* 4:595–600.
- Rose, Barbara R. 1976. Habitat and prey selection of *Sceloporus occidentalis* and *Sceloporus graciosus*. *Ecology.* 57, 3:531-541.
- Ruth, SB. 1978. A comparison of the demography and female reproduction in sympatric western fence lizards (*Sceloporus occidentalis*) and sagebrush lizards (*Sceloporus graciosus*) on Mount Diablo, California. Ph.D dissertation: University of California Berkeley. URL: <https://elibrary.ru/item.asp?id=7216038>.
- Schoville SD, Bonin A, François O, Lobreaux S, Melodelima C, Manel S. (2012). Adaptive genetic variation on the landscape: methods and cases. *Ann Rev Ecol Evol S.* 43(1), 23–

43. DOI:10.1146/annurev-ecolsys-110411-160248
- Schnell I, Bitarray contributors. c2021. bitarray: efficient arrays of booleans. URL: <https://github.com/ilanschnell/bitarray>
- Seabold S, Perktold J. 2010. statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*.
- Sexton JP, Hangartner SB, Hoffmann AA. 2014. Genetic isolation by environment or distance: which pattern of gene flow is most common? *Evolution* 68-1:1-15.
DOI:10.1111/evo.12258
- Spear SF, Balkenhol N, Fortin M-J, Mcrae BH, Scribner K. 2010. Use of resistance surfaces for landscape genetic studies: considerations for parameterization and analysis. *Mol Ecol*. 19:3576–3591.
- Stebbins RC. 1948. Additional observations on home ranges and longevity in the lizard *Sceloporus graciosus*. *Copeia*. 1: 20-22.
- Szymura JM, Barton, NH. 1986. Genetic analysis of a hybrid zone between the fire-bellied toads, *Bombina bombina* and *B. variegata*, near Cracow in Southern Poland. *Evolution* 40: 1141–1159.
- Thakkar S. 2018. Simulate-Correct-ColorBlindness. GitHub Repository: <https://github.com/tsarjak/Simulate-Correct-ColorBlindness>. Commit: 95edbbb6caa75e4869e0b7c12126a6965445c93c.
- Tinkle DW. 1973. A population analysis of the sagebrush lizard, *Sceloporus graciosus* in Southern Utah. *Copeia*. 2:284-296.
- Tinkle DW, Dunham AE, Congdon JD. 1993. Life history and demographic variation in the lizard *Sceloporus graciosus*: a long-term study. *Ecology*. 74, 8:2413-2429.
- van Rossum G. 1995. Python tutorial, Technical Report CS-R9526, Centrum voor Wiskunde en Informatica (CWI). Amsterdam, the Netherlands.
- Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, et al. 2020 SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17(3): 261-272.
DOI:10.1038/s41592-019-0686-2
- Wang IJ. 2013. Examining the full effects of landscape heterogeneity on spatial genetic variation: a multiple matrix regression approach for quantifying geographic and ecological isolation. *Evolution*. 67, 12: 3403–3411
- Wang IJ, Bradburd GS. 2014. Isolation by environment. *Mol Ecol*. 23:5649–5662.
- Wright S. 1930. Evolution in Mendelian populations. *Genetics* 16:97–159.
- Wright S. 1943. Isolation by Distance. *Genetics* 28:114–138.

Supplemental Text

Validation Tests

Full details and reproducible code and parameter files for each of our validation tests are available from the '/tests/validation/' subdirectory of the source code. We discuss the key details and results of these tests below.

Wright-Fisher test: genetic drift

The Wright-Fisher model of genetic drift models a fixed-size haploid population that turns over completely at each timestep (i.e. generation). The population can have any number of independent, biallelic genetic loci. For each locus, each generation's allele frequency is chosen as a binomial random variable, with the number of trials equal to the population size and the probability of success (i.e. of drawing the '1' allele) equal to the previous generation's '1'-allele frequency. The mean persistence time for an allele (i.e. the expected number of generations for which a locus remains segregating) is:

$$t^-(p) = -4N[(1-p)\ln(1-p) + p\ln(p)] \quad (S1),$$

where $2N$ is the number of alleles in the population (such that N can represent the diploid population size) and p is the frequency of an allele at the locus (Fisher, 1923; Hartl and Clark, 2007; Wright, Sewall, 1930).

The Wright-Fisher model is much simpler than the sorts of models for which Geonomics is designed (as are all of the following validation tests)—it is aspatial, panmictic, features fixed population sizes, and models only neutral loci. Thus, we parameterized Geonomics so as to approximate the model as closely as possible. To emulate aspatiality and panmixia, we used a population on a homogeneous landscape, using isotropic movement, with movement and dispersal distributions that broadly encompass the diagonal width of the landscape, and with no mating radius imposed (to allow panmixia) instead of local mating (i.e. with the mating radius set to 'None'). To enforce complete generational turnover, we set the maximum-age parameter to 1 (i.e. 1 timestep). While Geonomics does not maintain constant population size, we maintained the carrying-capacity raster at a constant, uniform value, thus maintaining a stationary mean population size. We simulated 250 independent neutral loci (by setting all inter-locus recombination rates to 0.5), with starting '1'-allele frequencies of 0.5 (although the actual starting frequencies vary slightly around this value because of sampling error when all individuals' genotypes are drawn binomially).

We ran the Wright-Fisher approximation test for three values of the carrying-capacity raster (i.e. three values of 'K_factor'), hence for three mean population sizes (708, 1564, and 2440 individuals). For each mean population size (calculated as the harmonic mean, to account for stochastic fluctuations around the carrying capacity), we compared mean persistence time to that expected by theory, according to equation S1. Figures S1 and S2 show that the results are a close match to the theoretical expectations.

Bottleneck test: population dynamics

Because drift is a stronger evolutionary force in smaller populations, drift accelerates in shrinking populations. If a population undergoes a bottleneck event, the overall effect of drift on the population during that time is expected to be larger than what a constant-size population of

equivalent starting size would experience during that time. Thus, mean fixation time should decrease in a bottlenecked population relative to one of constant size.

As with the Wright-Fisher model, to test the effectiveness of Geonomics for modeling a population bottleneck we used a homogeneous landscape with broad distributions for movement and dispersal and without a mating radius in order to emulate aspatiality and panmixia. To simulate a bottleneck event, we created a custom change event in which the population's carrying-capacity raster is reduced to 30% of its initial value for 50 timesteps (from the 200th to 250th), then returned to its initial value for the remainder of the simulation (through the 300th timestep). These simulations produced a clear signal of accelerated drift during the bottleneck event, with the mean rate of allele-frequency change, calculated in 15-timestep sliding windows, nearly tripling during the period of the bottleneck (Figure S3).

Stepping stone test: population subdivision and genetic differentiation

The stepping-stone model, or one-dimensional island model, is a spatially implicit model. It models a series of subpopulations arranged along a straight line, with migration between all neighboring pairs. As a combined result of divergence by drift and homogenization by effective migration, subpopulations are expected to reach a stationary level of genetic differentiation: migration-drift equilibrium. Theory provides the expected pairwise genetic differentiation between a pair of subpopulations at equilibrium as:

$$F_{ST} = \frac{1}{1 + 4 Nm} \quad (S2),$$

where N is the population size and m is the per-generation migration rate, such that Nm can be interpreted as the per-generation number of migrant individuals (Hartl and Clark, 2007).

To approximate the stepping-stone model, we created a Landscape Layer with a diagonal of six equally spaced, equal-sized islands (1.0-valued cells) embedded in a 'sea' of 0.0-valued cells. We used this layer as the carrying-capacity raster (Figure S4, left). We parameterized dispersal to be very near to parents' midpoints, movement distance to be strongly right-skewed, such that the long-distance movement events leading to migration are uncommon, and the mating radius to a value that makes island populations effectively panmictic but that prohibits mating between individuals on separate islands. Genomes contained 100 neutral loci, and we ran the simulation for 5000 timesteps.

Because Geonomics does not model discrete populations, it does not stipulate migration rates between discrete locations on the landscape. Thus, we manually tracked the number of migration events during each timestep for all possible directional migration events (i.e. for all permutations of island pairs), then used that data to calculate all mean migration rates. With those values, we solved equation S2, then compared the resulting F_{ST} expectations to the observed values (calculated from the simulated data using two common methods; Figure S4, right). We also used Discriminant Analysis of Principal Components (DAPC), performed in the R package *adegenet* (Jombart *et al.* 2008), to visualize population structure.

The results demonstrate that the model approached migration-drift equilibrium, as expected by theory (Figure S5), with all island populations reaching dynamic equilibria around the same mean size. Estimated migration rates and F_{ST} values qualitatively match theoretical expectations: mean migration rate drops off precipitously at greater than one step-distance apart, and genetic differentiation increases to approximate saturation. Values of F_{ST} consistently

undershoot the values expected based on estimated migration rates, however, because subpopulations have yet to approach fixation at most loci (which is the expectation implied by expected F_{ST} values close to 1). DAPC demonstrated that the simulation generated the expected population structure of six distinct clusters, one per island (Figure S6).

Contrasting-habitat test: divergent selection

In a population divided between two, divergent selective environments, if there is standing genetic variation for a biallelic locus controlling the trait adapting to those environments then theory predicts that the two subpopulations will diverge at that locus as each moves toward its respective adaptive peak. The rate at which divergence should occur depends on the relative strengths of two opposing evolutionary forces: natural selection, which causes divergence, and gene flow, which causes homogenization. The rate of allele frequency change in either subpopulation at timestep t is expressed as:

$$\delta q = \frac{-spq[q+h(p-q)]}{1-sq(2hp+q)} + m_i q^{\square} - m_o q$$

(S3),

where p and q are the frequencies of the beneficial and deleterious alleles in the local subpopulation, s is the selection coefficient against the homozygous recessive phenotype, h is the degree of dominance of the recessive allele, m_i and m_o are the migration rates into and out of the subpopulation being analyzed, and q^{\square} is the frequency of the locally deleterious allele in the alternative subpopulation where it is beneficial (Hartl and Clark, 2007).

This model, like the stepping-stone model, is spatially implicit. To approximate this, we created a landscape with two layers. The first was divided into two equal-sized halves, one valued at 0.0, the other at 1.0; this layer was used as the layer driving natural selection. The second was valued uniformly at 1.0; this was used as the carrying-capacity raster (thus setting uniform population density across the landscape and determining, in sum, the overall carrying capacity of the landscape). We created one monogenic trait whose position was randomly chosen within a genomic architecture of 100 unlinked loci. We ran the model for 1000 timesteps for each of three values of the parameter ϕ (identical to s in equation S3): 0.1, 0.05, and 0.01. Given that Geonomics does not directly define a migration rate parameter, we tracked the number of migration events (i.e. individuals crossing the landscape's horizontal midline) during each timestep, then used that data to solve equation S3.

Results depict clear local adaptation to each of the two halves of the landscape, with spillover of opposite phenotypes and resulting heterozygote births occurring along the border between the two habitats (Figure S7). Allele trajectories in each half of the environment follow qualitatively the increasing and saturating trajectories expected by theory, but reach consistently more divergent allele frequencies than expected based on the theoretical calculation (Figure S8). However, these results are an easily understandable artefact of estimating a spatially implicit, population-based model using a spatially explicit, individual-based one—our method of calculating migration rates includes all individuals who cross the habitat boundary, including the large number who only barely cross and who may even quickly migrate back, such that expected allele-frequency trajectories are based on an overestimation of true gene flow and thus serve as lower bounds on the real trajectories. As further validation, a plot of the mean difference between each individual's phenotypic and environmental values shows a strong decline over model time, with the rate and level of decline increasing as a function of increasing strength of

selection (Figure S9). Moreover, logistic regressions show no significant relationships between phenotypic and environmental values at the outset (pseudo- $R^2 \approx 0.0$, p-values > 0.1) but show highly significant relationships at the ends of the simulations (p < 0.0001 for all values of ϕ), with the amounts of variation explained increasing as a function of selection strength (pseudo- $R^2 = 0.327$ for $\phi = 0.01$, 0.376 for $\phi = 0.05$, and 0.406 for $\phi = 0.1$).

Cline test: local adaptation

In a clinal model, a population adapts locally across an environmental gradient, which is characterized by the extremes of its environmental values and its steepness (i.e. the instantaneous rate of environmental change along it). Local adaptation across this gradient will generate a geographic cline in allele frequencies. The clinal pattern is only expected for loci under selection along the cline (and other loci in linkage). Unlinked loci have no long-term clinal expectation (though they could initially be swept along with the selective locus, and any number could continue to show spurious concordant clinal variation). To detect clinal adaptation, we can fit cline curves to the allele-frequency variation across the environmental gradient for all loci, with the expectation that the clines fit to adaptive loci will mirror the gradient. Numerous equations have been used to fit clines, but one of the most common is the sigmoidal *tanh* function:

$$p_x = \frac{1}{2} \left(1 + t \left[\frac{2(x - c)}{w} \right] \right) \quad (S4),$$

where p is the frequency of the reference allele at position x along the cline, c is the centerpoint of the cline (such that $p_{x=c} = 0.5$), and w is the ‘width’, which is defined as $w = \frac{1}{\text{slope}}$ at point c (Porter, 2013).

To implement the cline model in Geomics, we created a landscape with two layers. The first layer was an environmental layer—a symmetrical, non-linear gradient between 0-valued and 1-valued halves (Figure S10). The second was a uniformly valued habitat-quality layer, used to set a uniform population density and thus determine the global carrying capacity. We created a monogenic trait whose locus was randomly placed within a genomic architecture of 100 independent loci. The trait had a ϕ of 0.01, with the gradient layer serving as its selective force. We ran the cline model for 1500 timesteps, then used a numerical optimization function (in Python’s *scipy* package; Jones *et al.* 2001) to fit equation S4 for all loci. We plotted all fitted clines on top of the first landscape layer, with the cline for the single selective locus highlighted. The selective locus consistently and clearly stands out as the only locus with a cline matching the expectation of a monotonic pattern mirroring the environmental gradient and spanning nearly the full range of phenotypic values (Figure S11).

Results clearly show a pattern of clinal adaptation across the landscape—despite isolated patches of maladaptive genotypes potentially resulting from occasional long-distance migration events—with a zone of admixture and phenotypic spillover surrounding the cline’s center (Figure S10). In a Bonferroni-corrected family of locus-wise logistic regression models of environmental value on genotype, the selective locus consistently stands out as the most significant (p-values of roughly 3×10^{-100}). Furthermore, a plot of the mean difference between phenotypic and environmental values shows a strong decline over model time (Figure S12), and logistic regressions show no significant relationship between phenotypic and environmental values at the outset (pseudo- $R^2 = 0$, p-value = 0.370) but a significant relationship at the end of the simulation (pseudo- $R^2 = 0.345$, p-value < 0.0001).

Selective sweep test: genetic hitchhiking

Genomic architecture and linkage add important complexity to models of molecular evolution. The most basic model of selection with linkage is that of a selective sweep: a beneficial mutation occurs in a population, falling on a random genomic background, then rises in frequency because of its selective advantage until it becomes fixed, pulling up the frequency of the surrounding haplotype block in the process. The haplotype block is, nevertheless, subject to recombination, which gradually erodes it symmetrically around the beneficial mutation. Thus, the selective-sweep model predicts that once a beneficial mutation occurs—as long as it is not lost early on by chance—it and the haplotype block around it will rise in frequency, the mutation will eventually fix, potentially with some core block around it, and the rest of the block will erode over time. The haplotype block should be clearly visible in genomic data, where it will manifest as a genomic region of reduced diversity and heterozygosity centered on the mutation.

To implement the selective sweep model in Geonomics, we again created a model approximating an aspatial, panmictic population (see Wright-Fisher test for details). We created a single, monogenic trait with a ϕ of 0.1. The trait's locus was manually set to position 500, such that it was at the center of the 1001-locus genome. The genome had a homogeneous recombination rate of 0.001 between all neighboring loci. We manually set the starting '1'-allele frequency at this locus to 0.0 but set the trait to be selected upon by a uniform layer of 1 values, such that all individuals began the model equally unfit (i.e. with a fitness value of $1 - \phi = 0.9$). After burn-in, we iteratively chose a random individual, introduced a '1'-mutation in its genome at locus 50, ran the model for 50 timesteps, and checked whether the '1' allele had reached a frequency greater than 0.05 by that time. We iterated until that check was passed, at which point we declared the mutant allele 'established' and continued to run the model until 2500 timesteps after the novel mutation reached fixation. At three timepoints during that model we calculated and recorded genome-wide nucleotide diversity using a sliding-window approach.

We found that Geonomics successfully and realistically simulated the behavior of a selective sweep. The first adaptive mutant that was not immediately lost by drift rose rapidly in frequency, then fixed. The population's mean fitness increased quickly from 0.9 (the universal fitness value before the mutation was introduced) to 1.00 (the universal fitness value after the sweep was complete; Figure S14). The linkage block around the selected locus became a region of depressed nucleotide diversity (Figure S13, top row) and heightened linkage (Figure S13, bottom row) – the classic signature of a selective sweep.

Recombination test

To provide additional validation of Geonomics' recombination model, we compared the effective recombination rates observed in a Geonomics model to those produced by an msprime simulation using the same recombination map. We produced a recombination map by assigning 999 random, interlocus recombination rates to a 1000-length simulated genome. We drew the rates by taking the first 999 values ≤ 0.5 from a random vector drawn from the distribution $\sim \text{Beta}(0.4, 1.3)$, producing a left-skewed distribution that nonetheless sampled the full range of physical linkage values. We ran an msprime model using the msprime.RecombinationMap object created from those values and also ran a Geonomics model using those values as the recombination-rate column in a Geonomics custom genomic architecture file. We then plotted

the true recombination rates and the observed breakpoint densities from both models, binned within even-width genomic windows. The results show that Geonomics' observed breakpoint densities recapitulate the true recombination rates just as closely as do those of msprime (Figure S15).

Applications

Example 3: Polygenic adaptation to climate change in the Yosemite region

To build this simulation, we first generated a template Geonomics parameters file (using the Geonomics function 'gnx.make_parameters_file(...)'), then edited it to best emulate our empirical study system (see Code Sample S1). This created a parameters file for a simulation with: (1) three empirical layers (mean temperature, precipitation, and habitat suitability), two of which (temperature and habitat suitability) have environmental change events; (2) one species, with one trait adapted to mean temperature; and (3) a data collection design. We set life-history parameters to reasonable approximations of *S. graciosus* biology, based on available literature. We set spatial parameters based on the relationship between the resolution of the environmental rasters and the characteristic scales of *S. graciosus* key life-history traits. The resolution of the rasters is 0.00833°, which at latitude 38° is equal to about 730.984 m in the east-west direction and 927.296 m in the north-south direction, giving each cell a total area of roughly 6.78×10^5 m² (67.8 hectares). Using a population density of 208 individuals per hectare (Tinkle, 1973) and the rough estimate that about 10% of the land area covered by our study contains the open habitat favored by *S. graciosus* rather than the more closed habitat favored by the congener *S. occidentalis*, with whom it experiences a large degree of competitive exclusion (Rose 1976), we chose a per-cell carrying capacity (parameter 'K_factor') of 67.8 hectares/cell \times 208 individuals/hectare \times 0.1 proportion of habitat suitable \approx 1410 individuals (which we then further multiplied by 0.1 for computational tractability). We set the reproductive age to 2 years (Tinkle, 1973; Tinkle *et al.* 1993). We left the sex ratio at unity, given controversy in the literature about whether or not it skewed toward females because of lower male survival rates (Tinkle, 1973; but see Tinkle *et al.* 1993). We set the number of births per individual to be a Poisson random variable with $\lambda = 4.464$ individuals/clutch \times 2 clutches/year \times 0.16 survival rate = 1.428, based on an average clutch size of 4.464 across surveyed California populations (Tinkle *et al.* 1993), an average of 2 clutches per year (Tinkle, 1973; Tinkle *et al.* 1993), and an average rate of survival to the first year of 0.16 (Ruth, 1978). We estimated the mean interannual movement distance as 12.457 m (expressed as 0.01704 cell widths), based on an average of all recorded interannual movement events in Stebbins' (1948) study of *S. graciosus* home ranges. We used this as an order-of-magnitude estimate for movement, but increased parameters slightly above this value in order to pair the reduced population density we chose for purposes of computational tractability. Thus, we set a mating radius value of 0.5 cell widths, and parameterized movement and dispersal as '~Lognormal(7×10^{-5} , 0.3)' and '~Lognormal(7.5×10^{-4} , 1),' respectively. In the absence of any known published estimates, we set the population intrinsic growth rate to 0.5. The full code to perform this analysis is available as Code Sample S2 and in the Yosemite demo script, included in the Geonomics package.

Accessibility

For both of the two most common types of color blindness (protanopia and deuteranopia)

we tested the colorblind-friendliness of all of the Matplotlib color palettes used by Geonomics as defaults. For perceptive simulation we used the script provided by Sarjak Thakkar (2018). All color palettes retain interpretability.

Supplemental Figures

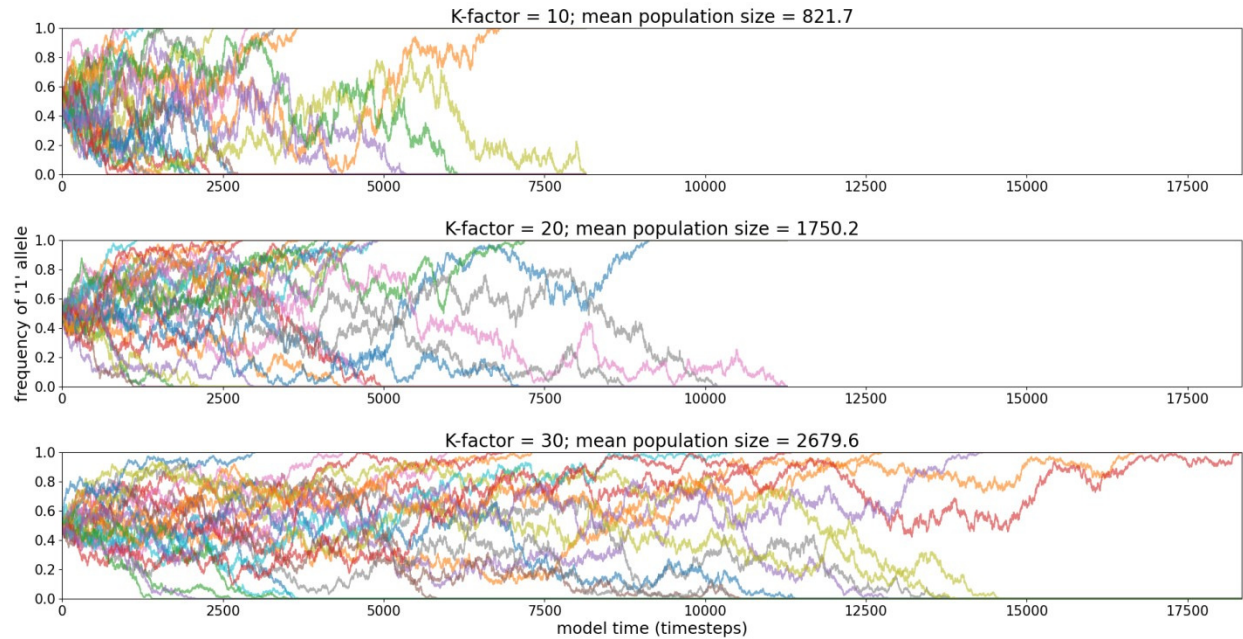


Figure S1: Trajectories for the frequencies of the ‘1’-alleles for 25 of the 250 simulated loci (one line per locus) in a Wright-Fisher model without mutation. We ran simulations for three mean population sizes, as determined by three fixed values of the carrying capacity (‘K_factor’) parameter, until all loci fixed.




Figure S2: Violin plots of mean persistence time distributions across all loci from our Wright-Fisher validation test, shown as a function of harmonic mean population size. Resulting mean persistence times (red dots) are an extremely close match to predictions calculated using Equation S1 (black, horizontal lines).

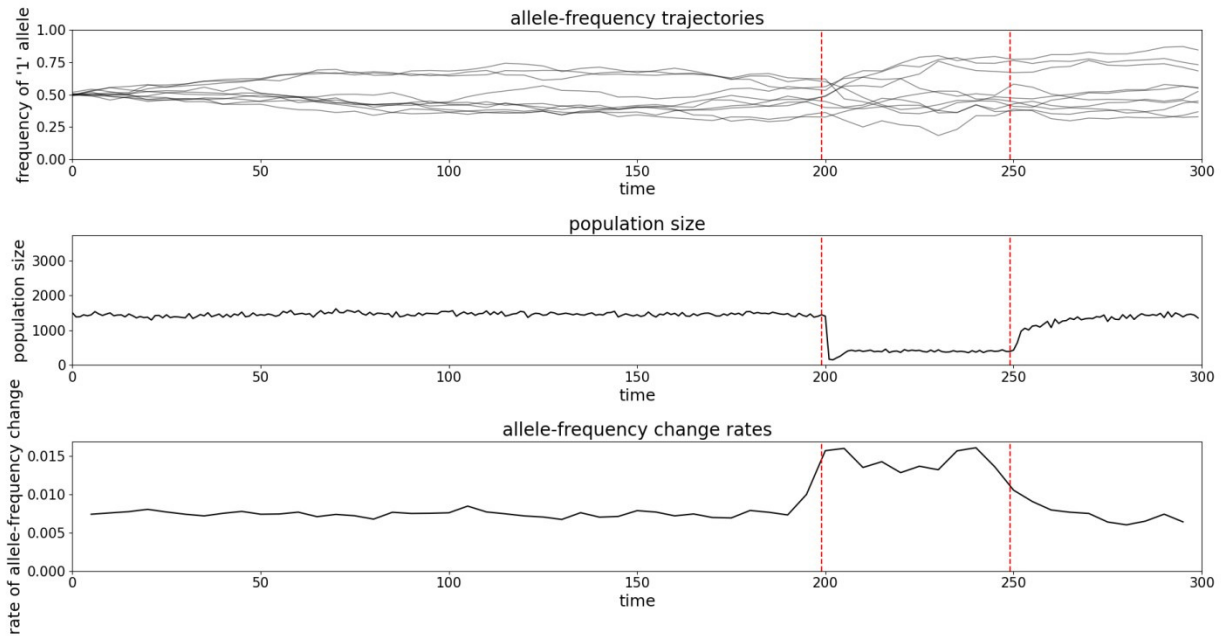


Figure S3: Ten randomly chosen allele frequency trajectories (top), population size (middle), and mean rate of allele frequency change (bottom; calculated for 15-timestep sliding windows) from the bottleneck validation test. We ran the simulation for 300 timesteps with a 50-timestep bottleneck.




Figure S4: Map of six island populations at the end of the simulation for the stepping-stone validation test (left), produced using ‘model.plot’ in Geomics, and plot of pairwise F_{ST} values and inter-island migration rates as functions of inter-island distance (right). R^2 values and p-values result from quadratic regressions of F_{ST} values on inter-island distances and log-log regression of mean migration rates on inter-island distances.




Figure S5: Plot of F_{ST} over model time for the stepping-stone validation test. Each line represents a different island pair, with colors corresponding to increasing inter-island distances (from yellow to green).

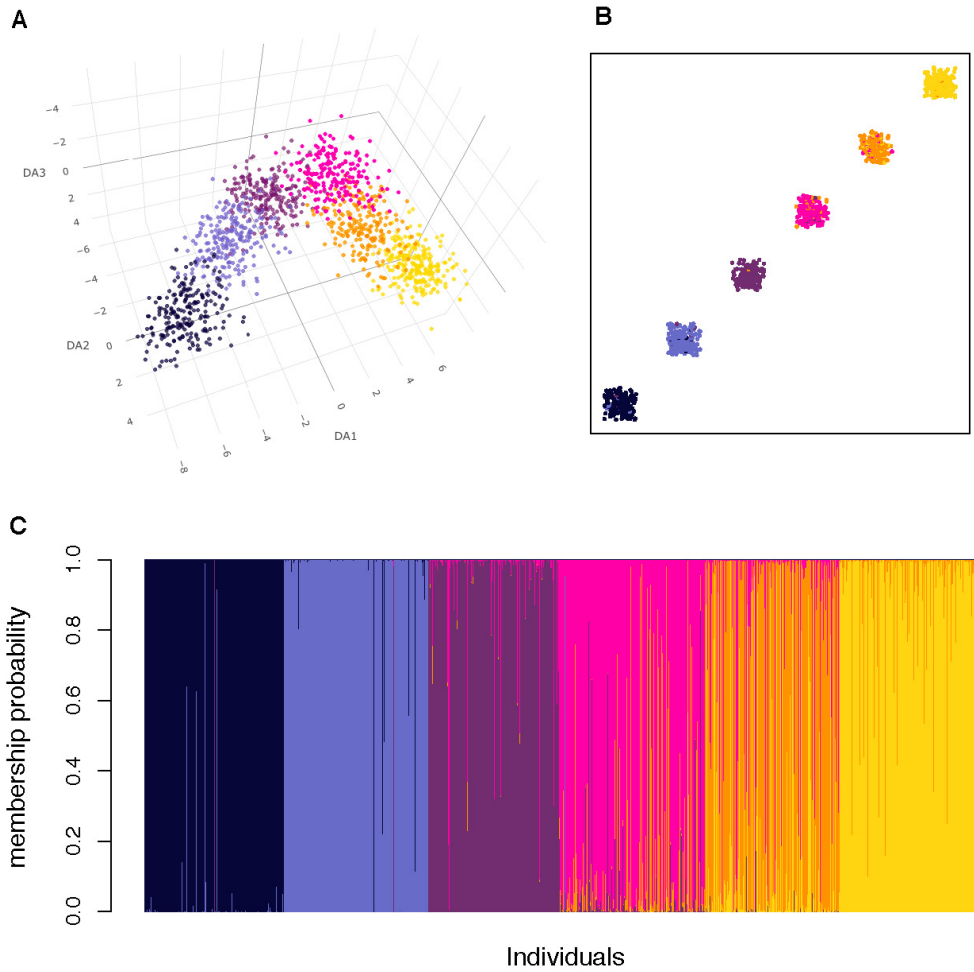


Figure S6: Results of discriminant analysis on principal components (DAPC) from the stepping-stone validation test, including plots of (A) the individual loadings on the first three discriminant axes, (B) individuals at the final time step color coded by population membership assignments, and (C) DAPC membership probabilities for each of the individuals. The optimal number of PCs to retain ($n = 59$) was determined through cross-validation using the ‘xvalDapc’ function in the adegenet R package (Jombart *et al.* 2008).

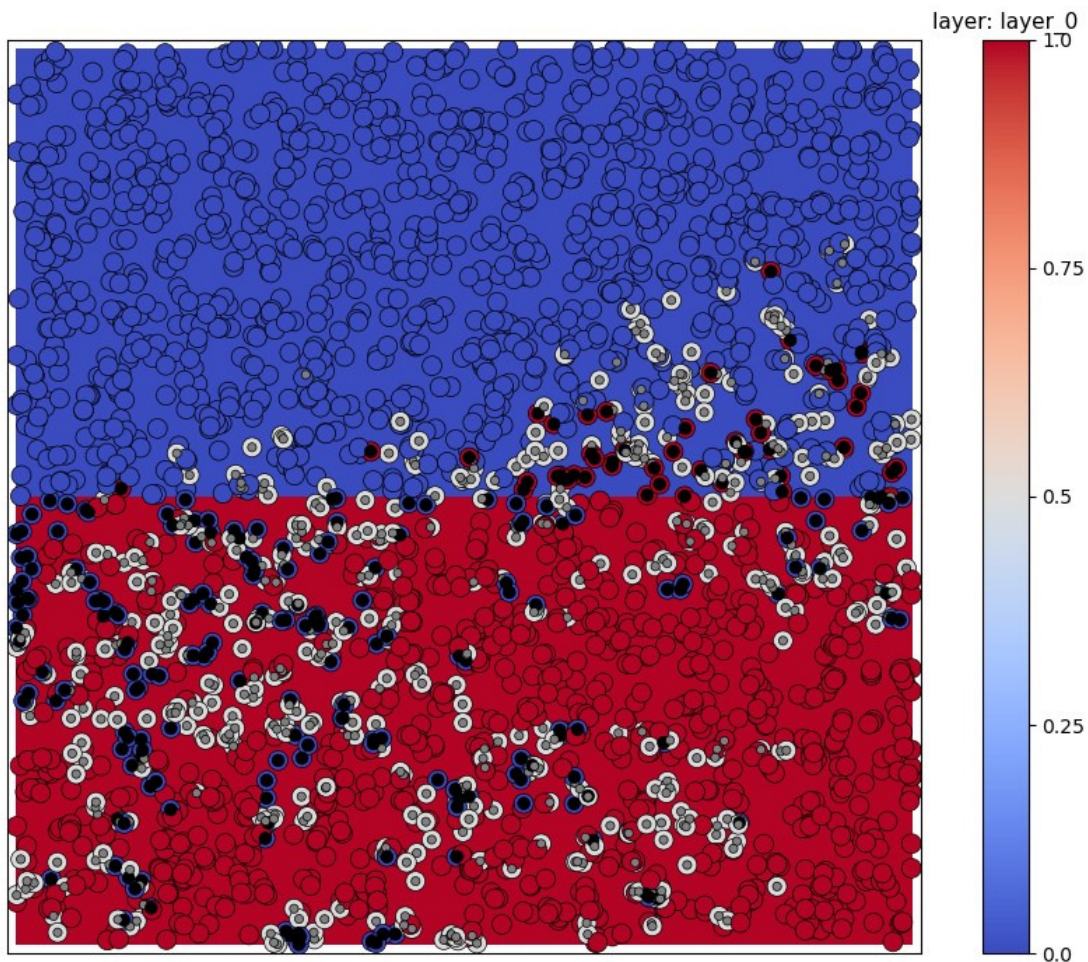


Figure S7: Map of the population after spatially divergent selection at $\phi = 0.10$ in simulations for the divergence validation test, produced using the ‘model.plot_fitness’ function in Geomics. Individuals are plotted on top of the selective landscape layer, which is divided into two halves. Outer circles are colored by phenotype, ranging from dark blue to dark red, representing the optimal phenotypes for each environmental background. Inner circles are color and sized by fitness, such that darker-gray, larger inner circles represent less fit individuals. Stochasticity leads to asymmetry in the structure of the hybrid zone, the nature of which varies from one iteration to the next; at the moment when this figure was produced, more blue alleles were present in the red environment than vice versa.

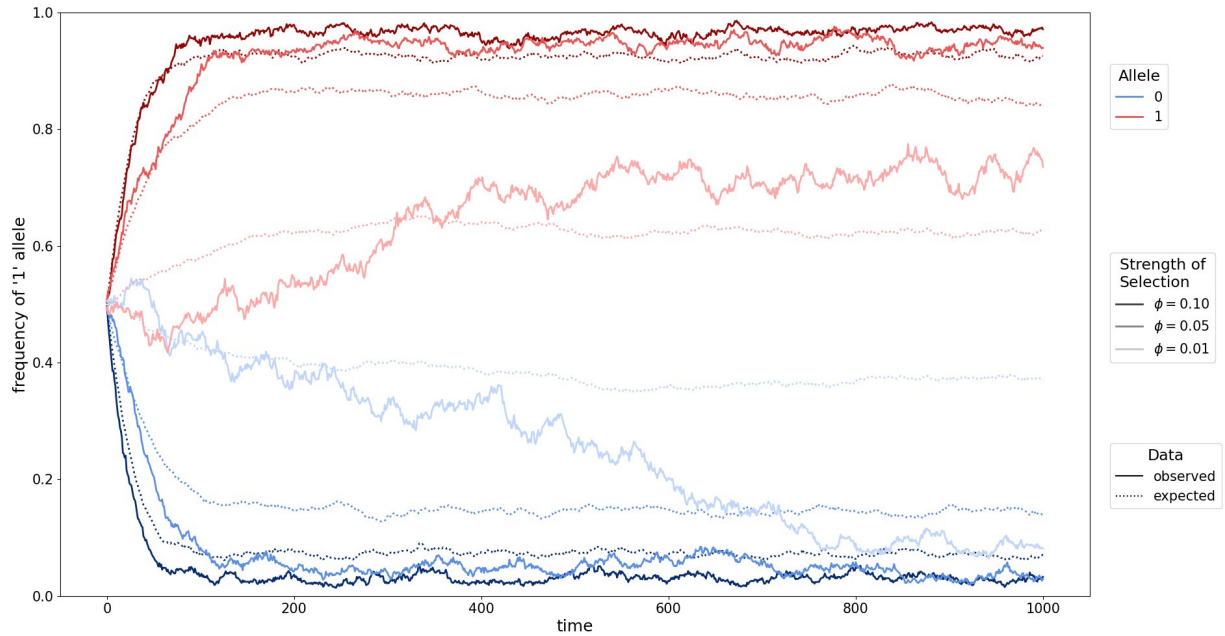


Figure S8: Observed (solid lines) versus expected (dashed lines) allele-frequency trajectories for two contrasting habitats (blue = 0.0-valued; red = 1.0-valued) resulting from divergence test simulations with three selection coefficients: $\phi = 0.01$ (dark), $\phi = 0.05$ (medium), and $\phi = 0.10$ (light).




Figure S9: Plot of the mean difference between each individual’s phenotype and environmental value plotted against time, for divergence test simulations with three different selection coefficients (ϕ). A pattern of background matching, which is indicative of local adaptation, builds up over time. The pattern develops more quickly, and becomes more pronounced, under stronger selection regimes.

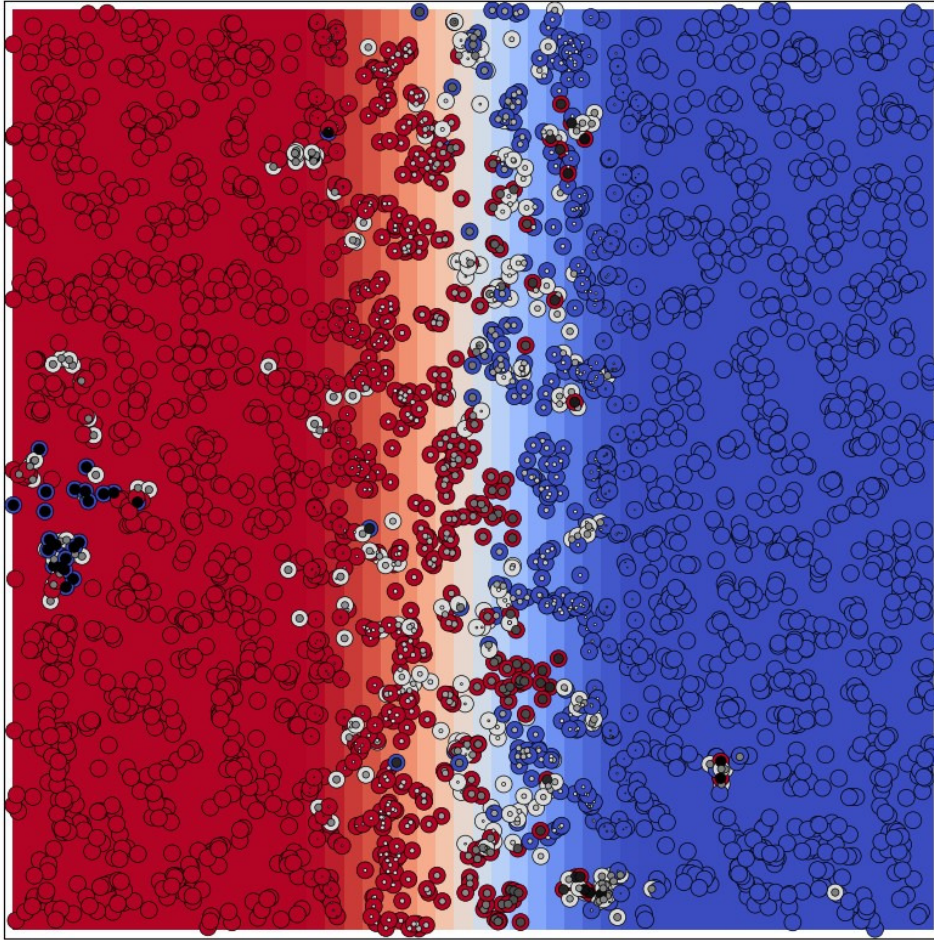


Figure S10: Map of the final generation from the cline test simulation on top of the selective landscape layer, with individuals colored by phenotype (outer circles) and fitness (inner circles), as in Figure S7.

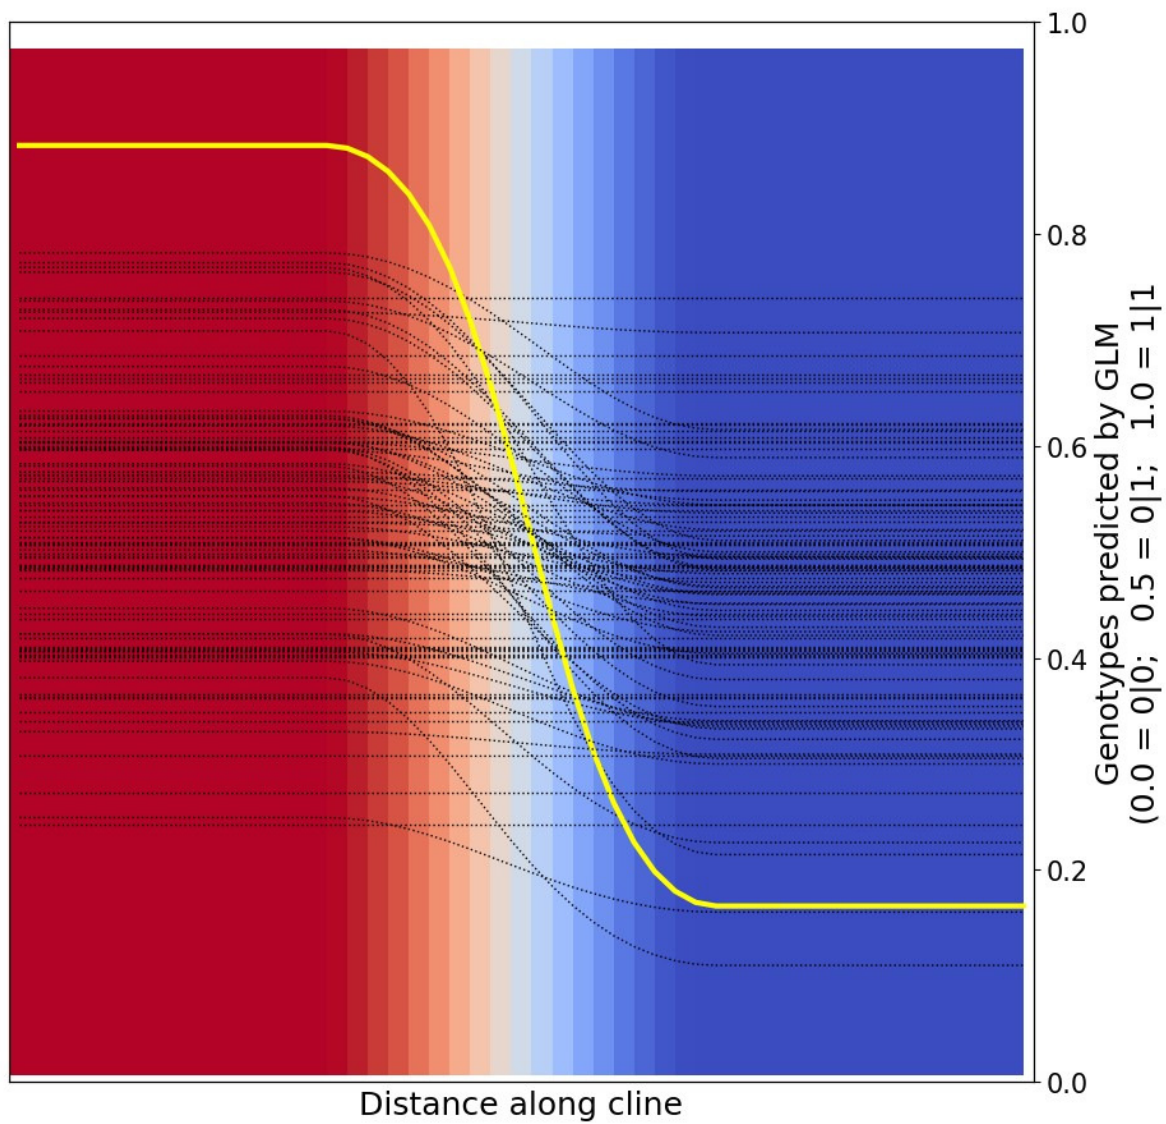


Figure S11: Plot of allele-frequency clines (neutral loci in black, selective locus in bold yellow) against the selective landscape layer (horizontal gradient from red to blue) from the cline test.




Figure S12: Plot of the mean difference between each individual's phenotype and environmental value plotted against time during the cline test simulation. A pattern of background matching builds up over time.

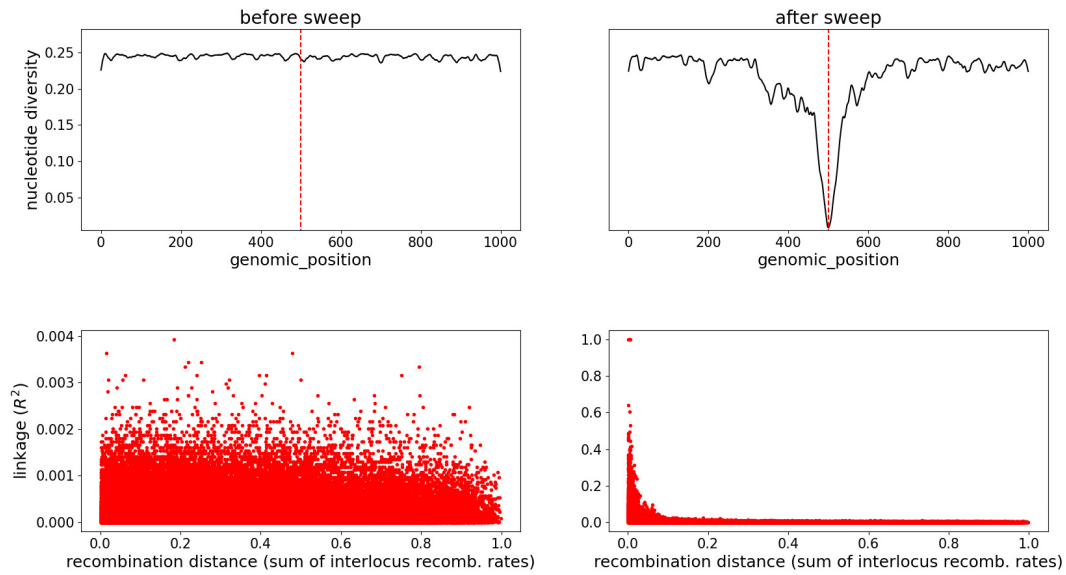


Figure S13: Results of the selective sweep validation test, including nucleotide diversity calculated in 11-locus windows across the genome (top row) and pairwise linkage (R^2) for locus pairs plotted against genetic distance (bottom row). Genetic distance was calculated as the ‘recombination distance’ (the sum of intervening interlocus recombination rates between paired loci).




Figure S14: Mean fitness of the entire population, over the full run of the selective sweep test simulation.




Figure S15: Breakpoint densities, within 50, even-width genomic windows, as calculated from the `tskit.TreeSequence` results of a Geonomics simulation (red) and an equivalent `msprime` simulation (blue). The observed densities clearly recapitulate the true recombination rates expressed in the input recombination map (black dashed line).

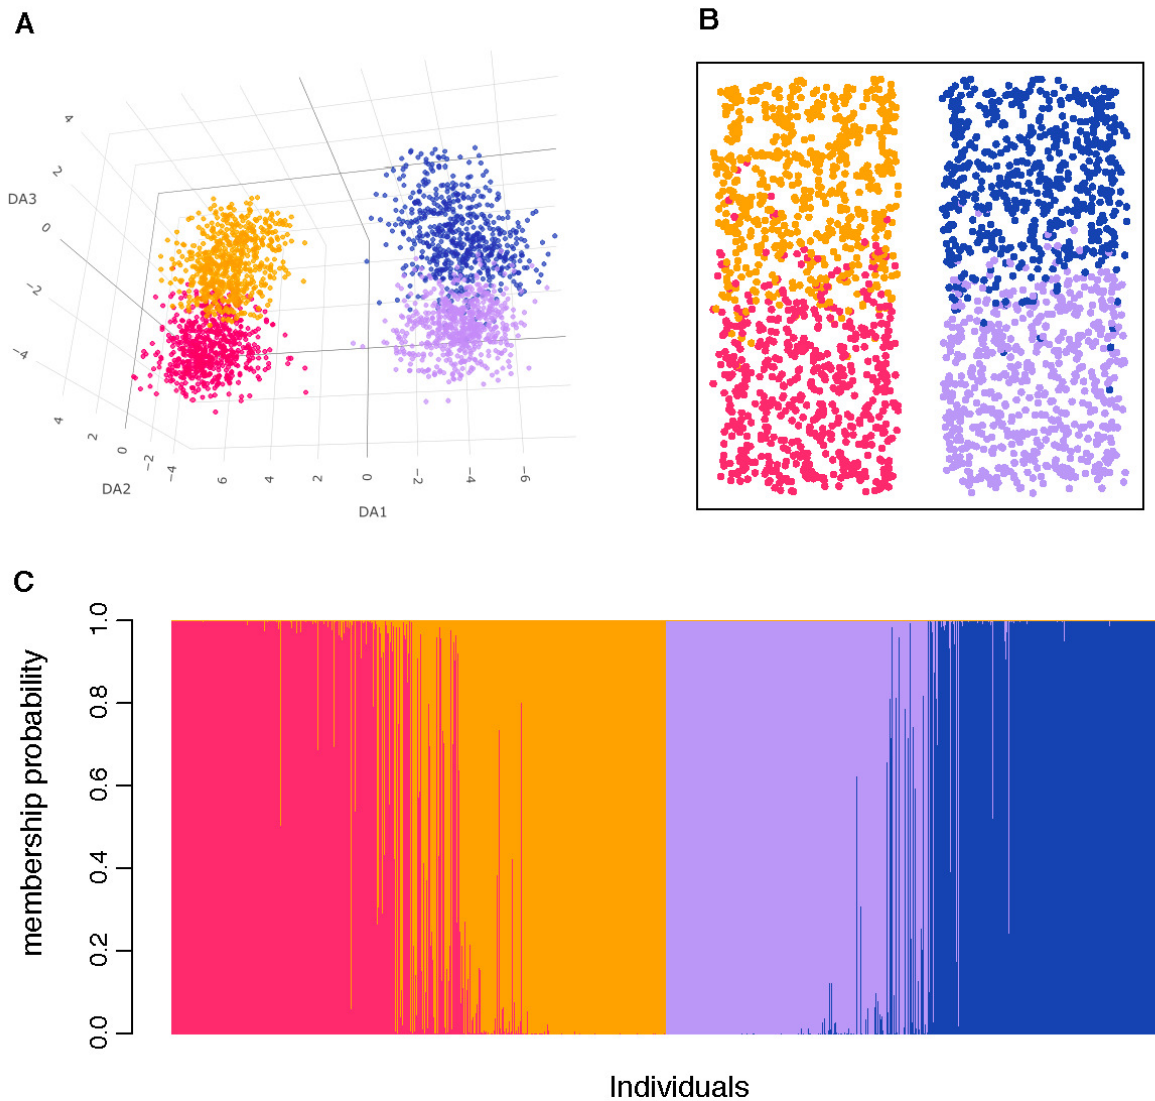


Figure S16: Results from discriminant analysis of principal components (DAPC) of neutral-locus genotypes at the final time step from the simulations for the isolation by distance (IBD) and by environment (IBE) example, conducted using the R package *adegenet* (Jombart *et al.* 2008). The optimal number of PCs to retain ($n = 7$) was determined through cross-validation using the ‘xvalDapc’ function (Jombart *et al.* 2008). A plot of the individuals’ loadings on the first three discriminant axes (A), with each individual colored according to its DAPC-derived population membership assignment (C), recapitulates their spatial arrangement (B). Beyond showing general patterns of IBD, our simulated neutral genetic data clearly match the expected hierarchical population structure: distinct clusters separated by the central barrier, with the subclusters further differentiated along the environmental gradients running in opposite directions on either side.




Figure S17: The mean difference between individuals' phenotypes and environmental values plotted against time, resulting from the simulations for the simultaneous selection example application. Values decrease over time for both traits, reflecting the buildup of a pattern of background matching.

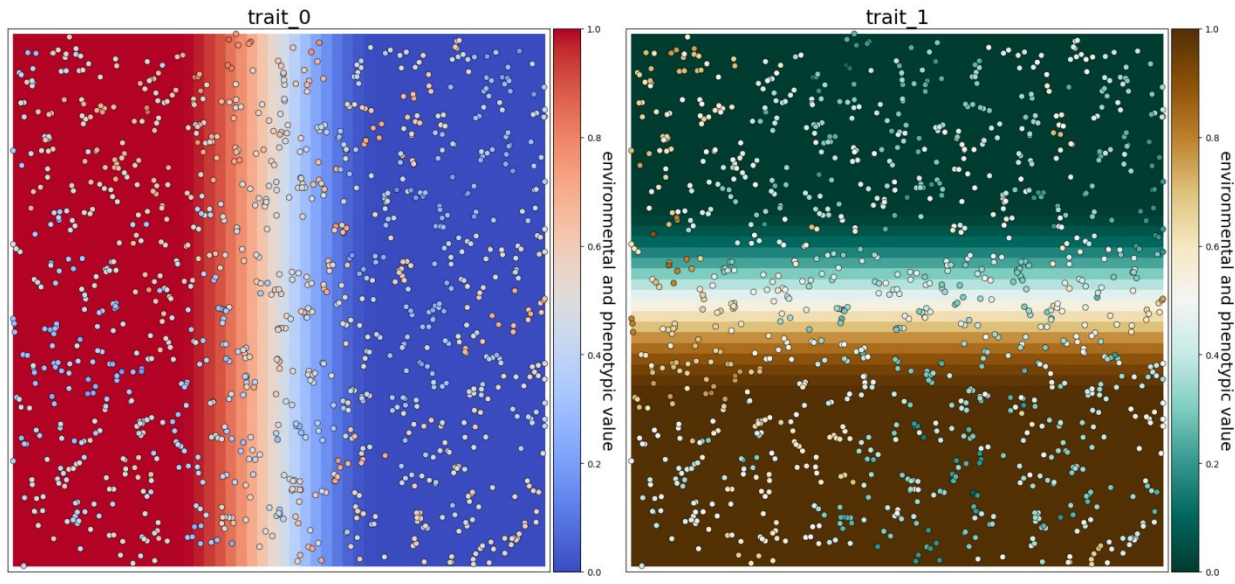


Figure S18: Results of the simultaneous selection simulation when selection is excluded ($\phi = 0$). Individuals are colored by phenotype for the trait corresponding to each layer and show no signal of background matching, as expected.




Figure S19: The mean difference between individuals' phenotypes and environmental values plotted against time, during the simultaneous selection simulation, when selection is excluded ($\phi = 0$). These values show no decreasing trend over time, unlike in the model including selection (Figure S16).

Code Sample S1:

```
>>> gnx.make_parameters_file(filepath='yosemite_params.py',
                             layers=[{'type': 'file', 'change': True},
                                       {'type': 'file', 'change': True},
                                       {'type': 'file', 'change': False}],
                             species=[{'movement': True,
                                       'movement_surface': True,
                                       'genomes': True, 'n_traits': 1}],
                             data=True)
```

Code Sample 2:

```
>>> model = gnz.make_model(filepath='yosemite_params.py')
>>> model.walk(100000, mode = 'burn')
>>> model.walk(500, mode = 'main')
```


Chapter 2

Genomic architecture controls multivariate adaptation to climate change

by

Drew Ellison Terasaki Hart¹

Ian J. Wang¹

¹ Department of Environmental Science, Policy, and Management, University of California,
Berkeley CA, USA

Abstract

As climate change advances environmental gradients decouple, generating novel multivariate environments that stress wild populations. A commonly assumed mechanism of evolutionary rescue is adaptive gene flow from populations whose current climates approximate future projections – but novel multivariate environments have no current analogs by definition, and gene flow from similar climates can cause maladaptive introgression with respect to decoupling environmental variables. Genomic architecture can play an important but often-ignored role in determining the effectiveness and relative magnitudes of adaptive gene flow and *in situ* adaptation. Here, I use Geonomics, the Python package presented in chapter 1, to simulate multivariate evolutionary responses to climate change under scenarios of variable trait polygenicity, linkage, and genotypic redundancy. I test a series of hypotheses about cross-scenario variability in evolutionary outcomes, reporting results that partially corroborate but also complicate or extend previous findings based on univariate and static environments. First, I find that up-gradient gene flow is consistently present, but its magnitude is strongly constrained under lower linkage and higher polygenicity and redundancy, suggesting *in situ* adaptation as a more effective mechanism of evolutionary rescue under these conditions. Second, I find that high polygenicity causes increased maladaptation and acute demographic decline, a concerning result given the likely polygenic nature of many climate-adapted traits. Finally, I show that adaptive capacity increases under higher genotypic redundancy across all scenarios, adding to the growing recognition of its importance and suggesting opportunities for better understanding the climatic vulnerability of real populations.

Introduction and Background

Climate change is one of the foremost threats to biodiversity in the Anthropocene. The ability of species to persist within their current ranges will likely depend largely upon their abilities to locally adapt to climate change through natural selection – a concept frequently referred to as ‘adaptive capacity’ or ‘evolutionary potential’ (Chevin *et al.* 2010; Harrisson *et al.* 2014, Nicotra *et al.* 2015; Vilas *et al.* 2015; Wade *et al.* 2017). Because adaptive *de novo* mutations take a long time to arise, this adaptation will instead most likely be facilitated by the reconfiguration of existing adaptive diversity (Bomblies *et al.* 2022). A common conceptual model underlying this assumption is that of adaptive gene flow tracking a shifting climatic gradient (i.e., in the vector direction of climate velocity; (Loarie *et al.* 2009; Ackerly *et al.* 2010), which would bring beneficial genes into recipient populations from ‘climate-suitable’ populations – i.e., populations whose current climates approximate future local projections (Bellis *et al.* 2020). This model of adaptive gene flow has both theoretical (Aitken and Whitlock 2013; Slatkin 1987; Tigano and Friesen 2016) and empirical (Feder *et al.* 2012; Bell *et al.* 2011) support, but meets resistance under theoretical conditions in which gene flow can be maladaptive (Wang and Bradburd 2014; Lenormand 2002; Slatkin 1987; Haldane 1930; Wright 1931; Felsenstein 1976). In these circumstances, shifting allelic covariance – the *in situ* recombination of standing genetic variation into new, adaptive genotypes – could be a more efficient mechanism of local adaptation to environmental change.

In recent decades, research bridging the fields molecular population genetics and quantitative genetics (Barghi *et al.* 2020; Barton 1999; Pritchard *et al.* 2010a; Pritchard *et al.* 2010b) has revealed that the genomic architecture of a trait is a core determinant of whether and how that trait becomes locally adapted (Yeaman *et al.* 2022). The number of loci underlying a trait (henceforth, ‘polygenicity’), the rate of recombination between trait loci (i.e., linkage), and the number of distinct genotypes that yield identical trait phenotypes (i.e., genotypic redundancy; Yeaman 2022; Láruson *et al.* 2020; Barghi *et al.* 2020) are among the key aspects of genomic architecture that influence adaptation (Barton 1999; Yeaman and Whitlock 2011; Yeaman 2022; Le Corre and Kremer 2012). Previous research suggests that ecologically-important traits can vary from few loci of large effect (Martin and Orgogozo 2013; Rees *et al.* 2020) to many loci of small effect (Boyle *et al.* 2017; Rockman 2012; Savolainen *et al.* 2013; Sella and Barton 2019; Barghi *et al.* 2020), and shows that this variation in polygenicity can determine the rate and nature of local adaptation (Yeaman 2015). Linkage controls the likelihood that adaptive alleles cluster together, essentially forming alleles of larger effect size that are stronger targets of selection and more resistant to swamping gene flow (Yeaman and Whitlock 2011), thereby facilitating local adaptation (Tigano and Friesen 2016). Genotypic redundancy can facilitate local adaptation by allowing the existence of a stable phenotypic cline subtended by a genotypic dynamic equilibrium consisting of continuous and concerted shifts in non-neutral allele frequencies (i.e., ‘transient genomic architectures’ (Barghi *et al.* 2019; Manceau *et al.* 2010; Yeaman 2015).

The influence of genomic architecture on the nature and outcomes of local adaptation to changing environmental gradients has been studied to a limited extent, with nearly exclusive focus on univariate models of the selective environment (but see Schiffers *et al.* (2013)). Yet, such models have limitations for studying adaptation to climate change because species can be simultaneously adapted to multiple, unaligned environmental gradients (Guillaume 2011) that

can shift differentially and thus decouple as climate change advances (Crimmins *et al.* 2011; Daly *et al.* 2010), leading to the emergence of novel multivariate landscapes (Williams and Jackson 2007; Williams *et al.* 2007; Fitzpatrick *et al.* 2018). Thus, it is important to assess how the genomic architectures of multiple traits can combine to determine the nature and outcome of multivariate adaptation under climate change. Gene flow from ‘climate-suitable’ portions of a species’ range is often assumed to be adaptive from the perspective of a single trait adapted to a shifting climatic gradient, but this assumption may be invalid if this gene flow also conveys linked variation for a trait adapted to a second environmental gradient from which the climatic gradient is decoupling – i.e., if the gene flow contains alleles of countervailing adaptive value. Under these circumstances, the genomic architectures of both traits (e.g., polygenicity, linkage, and genotypic redundancy) can determine the relative likelihoods of adaptation by gene flow and *in situ* adaptation by shifting allelic covariance (Aitken and Whitlock 2013; Schiffers *et al.* 2013) – and thus, the overall evolutionary outcome.

Spatially explicit simulation is one of our strongest tools for improving our understanding of the complex dynamics of gene flow and adaptation under climate change (Capblancq *et al.* 2020). In the previous chapter, I developed, validated, and demonstrated the extended utility of a novel Python simulator, Geonomics (Terasaki Hart *et al.* 2021), which I designed with precisely this purpose in mind. Here, I use individual-based, spatially explicit Geonomics simulations to test how genomic architecture influences multivariate adaptation to climate change. We simulate the adaptation of a single population continuously distributed in space to a bivariate environment composed of two horizontal environmental gradients, each exerting selection on a separate trait. In our main models, we then simulate climate change on that landscape by holding one gradient constant while gradually shifting the other gradient horizontally, such that the decoupling environment pushes local fitness peaks toward novel regions of bivariate trait space (Fig. 1). We run 100 pairs of climate change simulations and null (i.e., stable climate) simulations, for each of eighteen scenarios resulting from the full factorial crossing of three key components of genomic architecture: genotypic redundancy, polygenicity, and linkage (the precise values of which are given in Table 1).

We analyze cross-scenario variation in the resulting spatiotemporal patterns of gene flow, population size and density, and phenotypic distributions – all of which are emergent properties of our simulation parameterizations, as explained in our methods and specified in Appendix 1 – to test a series of hypotheses about the influence of genomic architecture on multivariate adaptation under climate change. First, we hypothesize that up-gradient gene flow will be higher under climate change than under a stable climate across all scenarios, but that gene flow contributes least to climate change adaptation when linkage is low and polygenicity is high. This is because we expect gene flow to always have at least some adaptive value, but we also expect low-linkage, high-polygenicity architectures (i.e., ‘dispersed’ architectures; Yeaman 2015) to exhibit quick *in situ* adaptation via shifting allelic covariance among many small-effect alleles, akin to adaptation under transient genomic architectures, facilitating phenotypic shifts in the absence of up-gradient gene flow. Second, we hypothesize that stronger linkage and higher polygenicity will reduce a population’s adaptive capacity under climate change, manifesting as greater reductions in population size and mean fitness and more persistent maladaptation, because both conditions impose longer expected wait times for the emergence of recombinant haplotypes that push phenotypes further from their pre-change fitness peaks. Finally, we

hypothesize that higher genotypic redundancy facilitates adaptation to shifting gradients, much as it does on stable gradients (Barghi *et al.* 2019; Manceau *et al.* 2010; Yeaman 2015), resulting in smaller reductions in population size and mean fitness.

Results

Gene flow

Climate change led to a nearly universal increase in up-gradient gene flow compared to null simulations: difference in gene flow was >0 in all scenarios in Fig. 2, and in our regression results the fitted values' 95% confidence intervals were >0 in all but one (the moderate-polygenicity, low-linkage, high-redundancy scenario, where the confidence interval was just below 0). However, the magnitude of this increase was minimal under some scenarios, and was positively correlated with linkage ($\beta_l = 0.0129 \pm 0.0006$, $P < 2 \times 10^{-16}$) and inversely correlated with polygenicity ($\beta_p = -0.0142 \pm 0.0006$, $P < 2 \times 10^{-16}$), corroborating our first hypothesis. Correspondingly, down-gradient gene flow was universally suppressed under climate change, as expected. Of the three components of genomic architecture that we tested, polygenicity had the most striking effect on the extent to which up-gradient gene flow contributes to adaptation; moderate and high polygenicity scenarios generally had much lower up-gradient gene flow than did low-polygenicity scenarios, with low-redundancy, independent-linkage scenarios being the main exception. Moderate-polygenicity scenarios actually showed the lowest overall increase in up-gradient gene flow, though differences between moderate- and high-redundancy scenarios were minor.

Linkage and polygenicity

As expected, our null simulations showed essentially no changes in mean fitness (Fig. 2) or population size (Fig. S2) – aside from small modeling artefacts that applied to both null and climate-change scenarios – and their phenotypic distributions were stable through time (Fig. 3). The results of our climate change simulations exhibited decreases in population size and mean fitness that are the expected results of increasing maladaptation (Aitken and Whitlock 2013). They also revealed environment-tracking bivariate phenotypic shifts (Fig. 4) in line with the expected direction depicted in the lower half of Fig. 1 but with variable amounts of shortfall below the fitness-maximizing phenotypic shift. Across scenarios, the demographic impacts of climate change increased with increasing linkage (change in fitness: $\beta_l = -0.00183 \pm 0.00010$, $P < 2 \times 10^{-16}$; change in population size: $\beta_l = -33.33 \pm 1.287$, $P < 2 \times 10^{-16}$; maladaptation: $\beta_l = 0.0038 \pm 0.0004$, $P < 2 \times 10^{-16}$). (Maladaptation is defined as the area in bivariate trait space separating the central line of a population's post-change phenotypic distribution from the central line of the distribution that would optimally track the changing environment; see wedges in Figs. 4 and S3, and see Methods for details.) Demographic impacts also showed a signal of overall increase with increasing polygenicity (change in fitness: $\beta_p = -0.00217 \pm 0.00010$, $P < 2 \times 10^{-16}$; change in population size: $\beta_p = -15.07 \pm 1.287$, $P < 2 \times 10^{-16}$; maladaptation: $\beta_p = 0.0097 \pm 0.0004$, $P < 2 \times 10^{-16}$), although the cross-scenario trend was non-monotonic and complex: impacts were smallest at moderate polygenicity; more pronounced at low polygenicity and at high polygenicity when redundancy was high; and extreme at high polygenicity and low redundancy (Fig. 3, Fig. S2). Indeed, demographic decline was so strong at low redundancy and high polygenicity that adaptive capacity was effectively outstripped: demographic decline persisted throughout the

climate change period, with little indication of evolutionary rescue (i.e., stabilization and rebound) occurring until the post-change period (see Fig. 3, Fig. S2). The collapse of adaptive capacity in these scenarios is also visible in the large red areas of phenotypic-shift shortfall in Fig. 4. The low-redundancy, high-polygenicity, strong-linkage scenario had such low adaptive capacity that mean fitness declined by 5.2% on average (from 0.934 to 0.885), mean population size declined by 17.1% on average (from 6326 to 5246 individuals), and the simulated population ceased to occupy the rightmost, fastest-changing portion of the landscape. This is visible as the disappearance of population density in the portion of the post-climate change phenotypic space corresponding to that region of the landscape in Fig. 4, but is more clearly visible in the before-after population density maps (Fig. S4).

Genotypic redundancy

Our high-redundancy scenarios showed consistently smaller demographic impacts of climate change and higher adaptive capacity than their low-redundancy counterparts (change in fitness: $\beta_r = 0.00402 \pm 0.00016$, $P < 2 \times 10^{-16}$; change in population size: $\beta_r = 39.06 \pm 2.101$, $P < 2 \times 10^{-16}$; maladaptation: $\beta_r = -0.0098 \pm 0.0006$, $P < 2 \times 10^{-16}$), supporting our hypothesis that genotypic redundancy can facilitate adaptation to shifting environmental gradients (Fig. 2, Fig. S2). This effect was most pronounced in the high-polygenicity scenarios, which exhibited much milder demographic decline under high redundancy, despite still showing no evidence of demographic rebound until after climate change (Fig. 3). Indeed, increased redundancy put the demographic impacts of these scenarios on par with those of the low-polygenicity scenarios (e.g., compare low- and high-redundancy boxplots in Fig. 3 and in Fig. S2).

Discussion

Current theoretical understanding of evolutionary responses to climate change largely derives from a simplified mechanistic model in which adaptation is universally facilitated by up-gradient gene flow. This model is not only the conceptual basis for research but also the inspiration for some climate-smart approaches to biodiversity management (e.g., assisted gene flow; Aitken and Whitlock (2013)). However, adopting this model as the basis for theoretical and mechanistic research risks overlooking the influence of genomic architecture on multivariate adaptation to environmental change. Starting from a more realistic, multi-trait framework, our simulations demonstrate that up-gradient gene flow does indeed occur under climate change, but that its contribution to adaptation can be constrained by polygenicity and, to a lesser extent, linkage. Given the plausibility of the range of genomic architectures we simulate (Barghi *et al.* 2020; Boyle *et al.* 2017; Rockman 2012; Savolainen *et al.* 2013; Sella and Barton 2019; Bomblies and Peichel 2022), these results raise the compelling possibility that up-gradient gene flow, while unlikely to be net-maladaptive, could have limited ability to support adaptation in many systems, especially those where climate-adapted traits have more dispersed architectures (i.e., architectures composed of many genes of small effect; Yeaman (2015)), providing a higher capacity for *in situ* adaptation to rapid change. This poses an important question for subsequent research: How often are climate-adapted trait architectures dispersed, versus being more concentrated (i.e., composed of few loci of large effect)?

We also show that the genomic architecture of climate-adapted traits can influence the

nature and size of the demographic fluctuations that occur during evolutionary responses to climate change. Our results suggest that neglecting this fact can cause important contingencies to go overlooked, while accounting for it could extend current theory and provide valuable insight. First, our results suggest that strong linkage between non-neutral loci, especially under high polygenicity, can increase maladaptation and demographic decline during climate change. In the most extreme case, evolutionary rescue is absent: high polygenicity and low redundancy combine to drive dramatic and persistent demographic declines, and even cause local extinction when linkage is strong. This was unexpected in light of previous work reporting that dispersed architectures produce stable, resilient phenotypic clines despite transient genotypic composition (Yeaman 2015; Yeaman 2022), and thus that species with such architectures could exhibit rapid local adaptation (Aitken *et al.* 2008) – but may make sense in light of the quick evolutionary rescue observed in these scenarios’ artificial ‘post-change’ periods, suggesting that the pace of environmental change exceeded the pace of adaptation. We did expect evolutionary responses to climate change to be slower in these scenarios, because natural selection is less effective on smaller-effect alleles and because high linkage leads to longer expected wait times for the generation of novel, adaptive recombinants, but we did not expect adaptive capacity to be completely outstripped.

The fact that high genotypic redundancy reduces demographic decline, not just in the aforementioned scenarios but across them all, contributes to the building recognition of the importance of redundancy as a driver of population-genetic outcomes in polygenic systems (Láruson *et al.* 2020; Yeaman 2022). This also presents a possible mechanism to be explored in real-world populations living at species’ colder range edges. Much like the local populations in the rightmost region of our low-redundancy scenarios, these local populations could already be at the edge of the phenotypic space defined by their standing genetic variation; in this case, segregating redundancy (*a la* Láruson *et al.* (2020)) and thus adaptive capacity would be low, so local extinction vulnerability would be substantial. However, species whose cold range edges are predominantly determined by geographic barriers or biotic interactions rather than by bioclimate (Thomas 2010) could feature local populations more similar to our high-redundancy scenarios: segregating redundancy would be higher, so selection would be balancing rather than directional and adaptive capacity would be substantial.

Our findings also help advance the theoretical understanding of local adaptation with recombination. Recombination is generally regarded as disadvantageous in situations of clinal adaptation with gene flow, because it disrupts the association between adaptive loci underlying a single trait, but unstable environments are considered a major exception (Tigano and Friesen 2016). Our results suggest that recombination can also be advantageous under the monotonic environmental change that is characteristic of climate change, particularly when simultaneous adaptation is occurring for multiple traits adapted to decoupling gradients. This likely occurs because recombination allows for more effective *in situ* adaptation by shifting allelic covariance – despite the fact that it still disrupts the associations between loci that would otherwise allow for the development of larger-effect gene clusters. This suggests that *in situ* shifts in allelic covariance provide an alternative to adaptive gene flow as a mechanism for evolutionary rescue, especially in multi-trait systems where gene flow can be adaptive for shifting climatic gradients but maladaptive with respect to other decoupling gradients.

Second, the minimal demographic decline in our moderate-polygenicity models contrasts

with previous work finding that adaptation to a gradient is more effective under either concentrated or dispersed genomic architectures (Yeaman and Whitlock 2011). This disagreement may be attributable to the difference in timeframes between adaptation to a univariate environmental gradient and adaptation to a decoupling, multivariate gradient. Adaptation to a single, static gradient can proceed gradually, so may either favor large-effect alleles or allele-clusters in the long haul (Yeaman 2015; Yeaman 2022), once they have arisen by mutation, recombination, gene flow, or a combination thereof, or dispersed architectures (Bürger and Gimelfarb 2002; Kondrashov and Yampolsky 1996; Yeaman 2022; Yeaman and Whitlock 2011) in temporally fluctuating environments. However, the sudden onset of persistent environmental change in a population that is already locally adapted triggers a ‘race against time’, and it may be that the genomic architectures with optimal adaptive capacity are ‘middle ground’ architectures that comprise freely recombining loci with small-enough effect sizes to avoid large sudden declines in fitness from migration load but with large enough effect sizes to avoid the long wait times necessary for recombination to cluster many adaptive loci into larger-effect haplotypes. If this is true it would suggest an inherent tension between the architectures that might be expected to evolve in locally adapted populations prior to climate change and those most likely to facilitate adaptation to change – although the precise intragenomic positions of the loci involved, left simplistically symmetric in our simulations, may be influential in real-world systems.

A major challenge in simulation-based research is the complexity of the high-dimensional parameter space that could be explored. Informative studies can be constructed by focusing on a small set of key parameters while holding others at reasonable values, as we have done here. This nonetheless leaves unexplored a number of secondary parameters that can have non-negligible influence over the complex ecological phenomena of interest. In the case of evolutionary responses to climate change this provides areas for future research, including: population size (in Geonomics, an emergent property of a number of parameters, including carrying capacity and intrinsic population growth rate), which is a key determinant of the relative strengths of drift and natural selection (Murray *et al.* 2017) and of the wait time to emergence of recombinant haplotypes (Christiansen *et al.* 1998), among other important dynamics; movement behavior, a key factor embedded in the rudimentary migration-selection dynamics that lie at the heart of models like ours (Wright 1931; Haldane 1930; Barton 1999); allelic effect size distributions (Orr 1998), which are omitted here in favor of a single fixed effect size; and the spatial structure of the environment, including the geometries, slopes, and orientations of gradients (e.g., Benes and Bracken (2020)), and their rates of change. A number of other more complex evolutionary aspects could also be explored through a similar modeling framework, including pleiotropy (Thompson 2020) and epistasis, hybridization, and life history variation. Finally, important and conservation-relevant insight could emerge from the integration of other dimensions of climate change ecology, including range shifts (Weiss-Lehman and Shaw 2020), plasticity (Chevin *et al.* 2010) and range-wide variation in population density (Aitken and Whitlock 2013).

Adaptation is one of the main mechanisms by which species may persist under climate change. Given the multidimensional nature of ecological phenotypes and the decoupling of environmental gradients that climate change will frequently cause, adaptive responses will frequently be multivariate. Within this context, the potential influence of adaptive genomic

architecture on evolutionary outcomes has typically been neglected in conceptual and mechanistic models. Our results show that this neglect can conceal important variability in the nature and success of adaptive responses, including the strong influence of genomic architecture on the effectiveness of up-gradient gene flow and *in situ* adaptation, the magnitude and persistence of maladaptation, and the likelihood of concomitant demographic decline or evolutionary rescue. Further research in this arena, both simulated and empirical, can help refine our understanding of adaptation to climate change, both advancing evolutionary theory and facilitating the development of better-informed management strategies.

Methods

Simulation

We performed all of the simulations for this study using Geonomics (Terasaki Hart *et al.* 2021), a Python (Van Rossum and Drake, 1995) package for creating forward-time, agent-based, continuous-space landscape genomic simulations using arbitrarily complex life histories, environments, and environmental change scenarios. All of our simulated scenarios feature a species with two traits, each of which experiences selection on the basis of a different environmental variable. Both environmental variables are modeled as linear, horizontal gradients that initially span environmental values from 1 to 0, left to right. The fitness of individuals is a function of the difference between their local environmental values and their phenotypes, which are determined by the additive effects of multiple loci (i.e., without pleiotropy or epistasis) – a reasonable approximation of many traits of interest in real populations (Sella and Barton 2019).

Each simulation starts with a neutral burn-in period, ended by Geonomics' tests of temporal and spatial population stability, then runs for 2500 time steps non-neutral evolution, generating a pattern of local adaptation to the initial environment. After that, one of the environmental layers undergoes a change event in which the gradient's values change stepwise over a period of 250 time steps, resulting in a final gradient that spans values from 1 to 0.5, left to right. This creates a scenario in which the two environmental variables become decoupled, leading to the emergence of novel environments (i.e., sites occupying new vectors in two-dimensional environmental space). The purpose of this is to emulate a common phenomenon under climate change: the decoupling of multivariate environmental gradients, leading to the emergence of novel climates (Williams and Jackson 2007; Williams *et al.* 2007; Fitzpatrick *et al.* 2018). Importantly, this landscape arrangement generates heterogeneous rates of climate change, with the rate ranging from 0 at the leftmost edge to 0.5 per 250 time steps at the rightmost edge. This complicates interpretation of our results, but less so than in an alternative scenario with spatially homogeneous rates of change, which would generate an artefact of range expansion whose genomic signal would be superimposed on that of climate change adaptation. This latter approach is also of interest, and may be explored in future work, but the approach we chose here allows us to best isolate the evolutionary dynamics resulting from the components of genomic architecture that define our scenarios and hypotheses.

Next, we wrote a custom Python script that reads the template Geonomics parameters file, edits any parameters that vary across our scenarios, instantiates a model, runs a fixed number of iterations of that model, and outputs simulated genetic and other data. Our parameters of interest are the number of loci underlying each trait (parameter *n_loci*), the linkage between neighboring loci (i.e., the homogeneous recombination rate; parameter *recomb*), and the amount

of genotypic redundancy (parameter *redund*). The values we assigned to these parameters are provided in Table 1. Of note, in low-redundancy scenarios the n_loci values we specify produce many-to-one genotype-phenotype mappings at intermediate phenotypes that decline to one-to-one mappings at extreme phenotypes, whereas in high-redundancy scenarios the doubled n_loci values produce many-to-one mappings across all phenotypes between 0 and 1, inclusive – i.e., all phenotypes matching environmental values that occur somewhere on the landscape. The full factorial combinations of the chosen values of those parameters generate the set of simulation scenarios depicted in the tabular arrangements of Figs. 2, 3, 4, S2, S3, and S4. We used that script to run a set of batch jobs on the savio3 partition of UC Berkeley’s Savio system (each node has 96 GB RAM and 32, 2.1-GHz Skylake processors). For each scenario, we ran a total of 100 iterations of the scenario of interest, featuring a 250-time-step climate change period (henceforth, the ‘main’ scenarios), and 100 iterations of a paired null scenario without natural selection (henceforth, the ‘null’ scenarios). Given that Geonomics is a complex simulation framework, it features numerous other parameters, which we set at reasonable default values. Certain values of interest that might be parameter settings in some population-based, spatially-implicit simulation models are instead emergent properties within Geonomics’ individual-based, spatially-explicit framework; for example, the population size values we report herein emerge from the interaction of numerous parameters (e.g., the raster of local carrying capacities, the population intrinsic growth rate, the number of offspring per reproduction event, and the death rates resulting from the parameters controlling both density-dependent mortality and natural selection), rather than being set *a priori*. The complete set of Geonomics parameters and the values we assigned to them across all models (as well as green-highlighted notes denoting runtime and genomic-architecture parameter overridden by the main simulation scripts) are provided in Appendix 1.

Using a combination of internal Geonomics functions and custom Python code, we designed a set of data outputs from each model run, to visualize our results and test our series of hypotheses. We saved tables of individuals’ locations and phenotypes at both the beginning and the end of the climate change period. We also saved time series of population size, mean fitness, and mean phenotype of the trait adapted to the shifting gradient. We gathered this data at every time step, from 250 time steps before the onset of climate change, through the 250 time steps of the event, and continuing until 250 time steps after climate change completed. The final 250 time steps after climate change are unrealistic, but are useful for elucidating the nature of the changes that persisted to the end of the climate change period; we refer to this period as the ‘post-change period’.

We also saved subsampled data on the vector directions of gene flow occurring during climate change, by keeping all final individuals’ data pertaining to two randomly chosen loci that had positive effects on the trait adapted to the shifting environmental gradient. We restricted our sample in this way both to focus on loci expected to facilitate adaptation to increasing environmental values and thus to shift up-gradient, and to provide equal sample sizes across scenarios in downstream analysis (which was constrained to the number of positive-effect loci present in the low-polygenicity, low-redundancy scenarios). Data were collected using an internal Geonomics function that extracts data from the spatial pedigrees stored in the simulation’s *tskit* (Kelleher *et al.* 2018) data structures. We also calculated a single high-level summary metric per model that provided us a comparator for the relative amounts of up-gradient

gene flow that occurred across our scenarios, for use in hypothesis testing. We refer to the metric simply as ‘up-gradient gene flow’, and we calculate it as:

$$GF_{up} = \frac{\sum_i^n \cos \theta}{n}, \cos \theta \geq 0,$$

where angles are expressed counterclockwise from the right. The $\cos \theta \geq 0$ condition allows us to ignore leftward gene flow, which we expected to be low irrespective of scenario for the positive-effect loci for which we collected data, given that it would oppose the environmental shift so would be generally maladaptive.

Analysis

We analyzed all 18 scenarios' and all 100 iterations' data using visual summaries and companion statistical tests. All analysis and visualization was produced using custom scripts written in Python (Van Rossum and Drake 1995) and R (R Core Team 2021).

To test our first hypothesis about gene flow, we first produced a visualization of the directional distributions of gene flow in all 18 scenarios, comparing between main and null scenarios (Fig. 2). To create this visualization, we gathered directional data from each simulation for a random sample of the gene flow that occurred during the climate change event, using Geonomics' integration with *tskit* (Kelleher *et al.* 2018), which allows for temporal subsetting and output of the information contained the full spatial pedigree of a simulated population. We then fitted a mixture of 4 von Mises distributions to that data using the R package *movMF* (Hornik and Grün 2014), yielding 12 parameter estimates defining each simulation's fitted mixture distribution. For each of the 18 scenarios, we then plotted the probability density function described by the means of all length-12 vectors of fitted parameters. We did this separately for null scenarios and for main scenarios, then overlaid the main scenario (in red) on top of the nulls (in blue), providing a summary depiction of the nature of gene flow within each main scenario as compared to its null expectation.

Finally, we ran a simple linear regression of the main-null difference in up-gradient gene flow density as a function of the components of genomic architecture ($GF_{up,main} - GF_{up,null} \sim \alpha + \beta_p \text{polygenicity} + \beta_l \text{linkage} + \beta_r \text{redundancy} + \epsilon$), which we used to interpret our results and test the components of our first hypothesis. We model the genomic architecture components as numeric integer variables representing the factorial or exponential increases used to set their parameter values in simulations: polygenicity: low = 0, moderate = 1; high = 2 ($\text{polygenicity}_{variable} = \log_5((\text{polygenicity}_{simulated})/(4 \times \text{redundancy}_{variable}))$); linkage: low = 0, moderate = 1, high = 2 ($\text{linkage}_{variable} = -\log_{10}((\text{linkage}_{simulated})/0.5)$);

redundancy: low = 0, high = 1. We use this model, compared against the patterns revealed by Fig. 2, to test both components of our first hypothesis, expecting all scenarios' fitted values for the main-null difference in up-gradient gene flow to have 95% confidence intervals >0 for the main-null difference in up-gradient gene flow (component 1), and expecting the coefficients on the linkage and polygenicity terms of the model to be significantly positive and negative,

respectively (component 2).

To visually assess our second and third hypotheses, we created a series of plots comparing climate change-driven demographic shifts and maladaptation across all 18 scenarios and between null and main models. First, we plotted the null and main trajectories of two demographic metrics – mean fitness and population size – in each of our 18 scenarios. For each scenario and metric, and for both main and null models, we created ensemble datasets by combining all 100 iterations' output time series of the metric, then derived summary time series by calculating each time step's 100-iteration mean and 5th and 95th percentiles. We plotted the resulting summary time series in Fig. 3 (mean fitness) and Fig. S2 (population size), with null results again shown in blue and main results in red. We plotted the metrics starting from 250 time steps before the climate change event and running until 250 time steps after its completion, allowing us to evaluate the onset, course, and aftermath of the demographic responses. Additionally, we summarized all scenarios in a pair of box plots (one plot for low genotypic redundancy, one for high), with plots arranged by polygenicity along the x-axis, colored blue or red to denote null or main results, and shaded darker for increasing levels of linkage.

Next, to better understand changes in population size and distribution, we mapped before-after comparisons of population density across all 18 main scenarios (Fig. S4). Each population density map is calculated as the array of mean population densities at all cells on the landscape, averaged across all 100 main simulations of the map's corresponding scenario. Population densities are depicted as increasing from black to white, with values standardized across the entire plot grid.

Finally, to visualize maladaptation, we created an identically-structured grid of plots comparing each scenario's mean phenotypic distributions before and after climate change, including scatter plots of the density of individuals occurring across two-dimensional trait space and lines and wedges depicting the average maladaptation observed across each scenario's 100 iterations (Fig. 4). We refer to the wedge as 'persistent maladaptation', and we calculate it as the difference between: a.) the area within two-dimensional trait space that the population's phenotypic distribution would have needed to shift through during the climate change event, so as to remain optimally fit to its environment, and b.) the observed area of phenotypic shift within a scenario's 100 simulations. (We qualify this metric as 'persistent' to emphasize that it does not reflect transient maladaptation that arises but then resides during the period of climate change, but rather reflects only maladaptation that remains at the end of the climate change period.) To measure this area, we first determined the triangular area between the expected central tendency lines of the optimal bivariate phenotypic distributions before and after the climate change event; these are unambiguously determined by the model parameterization, because they are the lines connecting all of the discrete points in environmental space that occur on the pre- and post- change landscapes. Then, for each model run, we used ordinary least squares (OLS) to fit a central tendency line to the 100-iteration ensemble phenotypic distribution observed at the end of the climate change event (arithmetically fixing the y-intercept at the (1,1) point in phenotypic space, the unchanging phenotypic optimum at the leftmost extent of the landscape). The area of the wedge between the expected and observed post-change central tendency lines provides our measure of a scenario's persistent maladaptation. We plot before and after scatter plots of the ensemble datasets of individuals' two-dimensional phenotypes (binned to a grid of regular points for interpretability, with points shaded to depict relative densities of

individuals). We then overplot onto those scatter plots the central tendency lines of the expected (dotted lines) and observed (solid lines) phenotypic distributions, as well as the wedges of persistent maladaptation. In Fig. S3 we produce the same plot as Fig. 4, but using data from null simulations to demonstrate that all differences in maladaptation observed between scenarios were attributable to climate change, and with blue translucent wedges showing the expected area of phenotypic shift under each null scenario's corresponding main simulation, for reference.

To summarize our results, we ran identically structured simple linear regressions for each of our three response variables measuring population-level changes during the climate change event change in mean fitness, change in population size, and persistent maladaptation – with polygenicity, linkage, redundancy, and nullness serving as explanatory variables: $[\Delta N_t \mid \Delta fit \mid maladaptation] \sim \alpha + \beta_p polygenicity + \beta_l linkage + \beta_r redundancy + \beta_n nullness + \epsilon$. We model nullness as a binary categorical variable (null=0, main=1) and again model the genomic architecture components as numeric integer variables, as described above. We use these models, in combination with the cross-scenario trends visible in Fig. 3, Fig. S2, and Fig. 4, to test the components of our second hypothesis (polygenicity and linkage) and third (redundancy). Specifically, our second hypothesis predicts that the coefficients of the linkage and genicity terms are significantly non-zero and negative (for the fitness-change and population size-change models) or positive (for the maladaptation model) with increasing linkage and genicity, while our third hypothesis predicts significantly non-zero redundancy coefficients with the opposite signs.

Acknowledgments

We thank A. Bishop, T. Dawson, J. Frederick, N. Graham, M. Kelly, M. McElroy, E. Westeen, G. Wogan, and M. Yuan for feedback and guidance on various iterations of the simulations presented herein. We thank Berkeley Research Computing for providing access to the Savio computing cluster. We thank D. Ehrenfeld, N. Fefferman, M. Fitzpatrick, and L. Plough for cultivating an interest in conservation genetics. Lastly, we thank M. Terasaki Hart, C. Nemec-Hart, G. Hart, J. Hart, and M. Tylka for supporting and encouraging a lifetime of curiosity about nature, and C. S. aTunde Adjah, B. Evans, R. Pérez Joglar, Y. Y. Ma, and J. Redman for the pleasant solitude. This work was supported by a Berkeley Fellowship (to D.E.T.H.) and by National Science Foundation grant DEB1845682 (to I.J.W.).

Tables

component	level	parameter value
genotypic redundancy	low:	$redund = 1$
	high:	$redund = 2$
polygenicity	low:	$n_loci = 4 \times redund$
	mod:	$n_loci = 20 \times redund$
	high:	$n_loci = 100 \times redund$
linkage	low:	$recomb = 0.5$
	mod:	$recomb = 0.05$
	high:	$recomb = 0.005$

Table 1: Parameter values used for each of the three focal components of genomic architecture parameters. The full factorial combinations of these parameter values constitute the set of 18 simulation scenarios for which we present results. Of note, the interaction between the n_loci and $redund$ parameter values creates, in low-redundancy scenarios, genotype-phenotype mappings that are many-to-one for intermediate phenotypes but decline to one-to-one mappings at extreme phenotypes, but creates, in high-redundancy scenarios, many-to-one mappings across all phenotypes between 0 and 1, inclusive – i.e., across all phenotypes that are optimal within environmental values occurring somewhere on the landscape.

Figures




Figure 1: Conceptual model of adaptation to climate change. Above: Stacked, horizontal cross-sections of our square simulation landscape, shown for the shifting environmental gradient (e_1 , blue-red color ramp) and the stable gradient (e_2 , white-black color ramp), both before climate change (t_1) and after (t_2). Below: Bivariate fitness landscape of the traits adapted to the shifting and stable gradients, on axes e_1 and e_2 , respectively. Three example positions along the bivariate gradient (x_1 , x_2 , x_3) are delineated by thin vertical boxes on the physical landscape, both before climate change (gray) and after (yellow), and their corresponding fitness peaks are shown as color-matched kernels located along color-matched lines of the fitness optima that exist on the physical landscape before (t_1) and after (t_2) climate change. Shifts in local fitness peaks are shown as labeled arrows (x_2 , x_3); the environment at the far left of the physical landscape does not change, so x_1 's fitness peaks are overlapping and have no shift, whereas the environment at the far right of the physical landscape experiences the maximal rate of change, which is reflected in the shift in x_3 's fitness peaks. Note that fitness peaks are stylized and truncated for ease of depiction; in our simulations, fitness decreases as a linear function of an individual's distance from its phenotypic optimum, rather than the truncated Gaussian function depicted here.

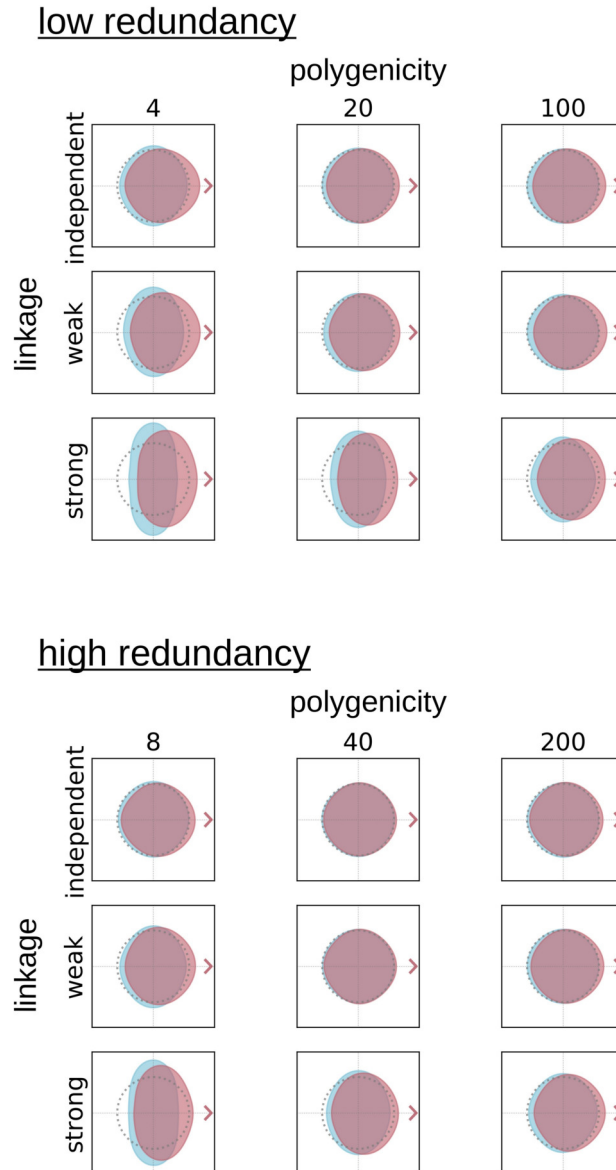


Figure 2: Comparison, across all 18 scenarios, of the distributions of gene-flow directions during the climate change period. Scenarios are organized into top and bottom sections for low and high redundancy, with rows in each section representing levels of linkage and columns representing polygenicity. Main scenarios (red) are compared against null scenarios (blue). Compass labels indicate directions of gene flow as it would be observed from a bird’s-eye view of the simulated landscape, with rightward (i.e., ‘up-gradient’) gene flow moving in the same direction as the shifting environmental gradient, and with upward and downward (i.e., ‘on contour’) gene flow being perpendicular to the environmental gradients. Down-gradient gene flow is expected to be maladaptive under all scenarios, explaining why it is universally suppressed relative to the null results. There is a general trend toward increasing on-contour gene flow and decreasing up-gradient gene flow with decreasing strength of linkage and increasing number of loci per trait.

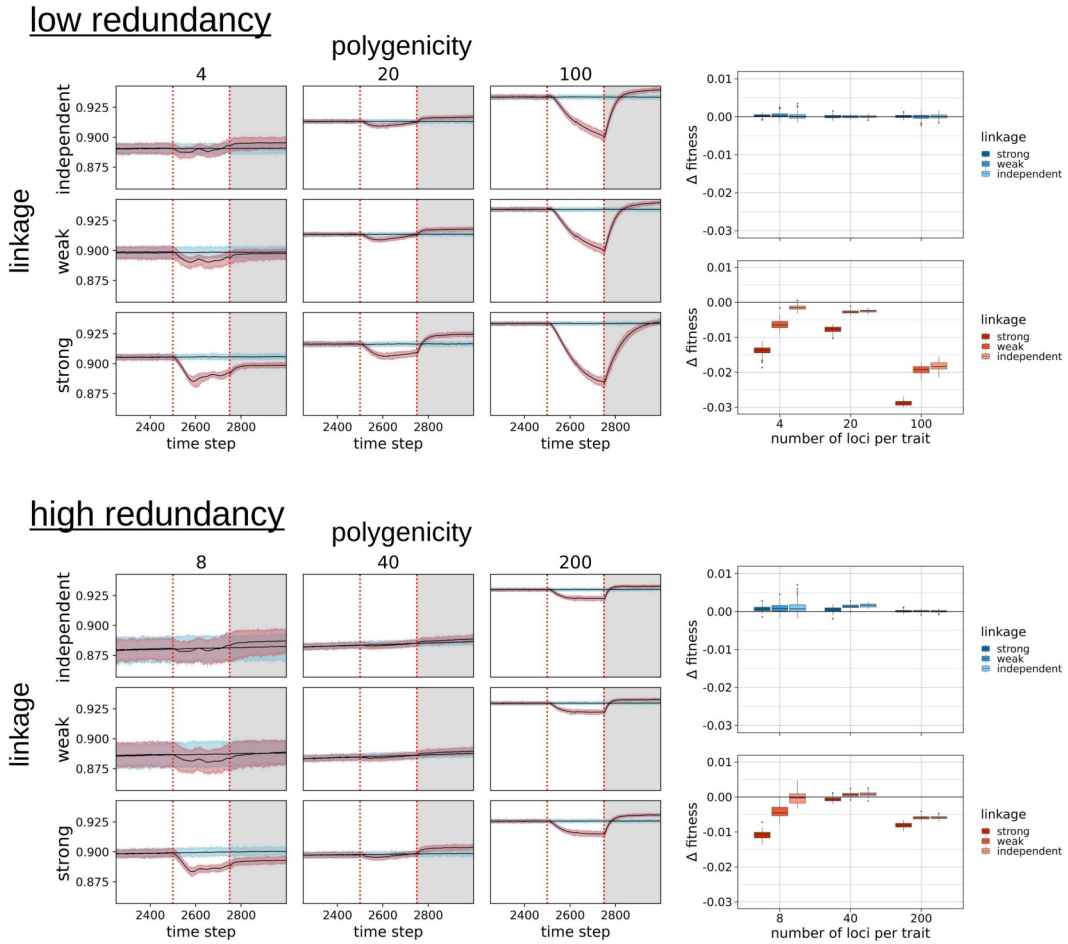
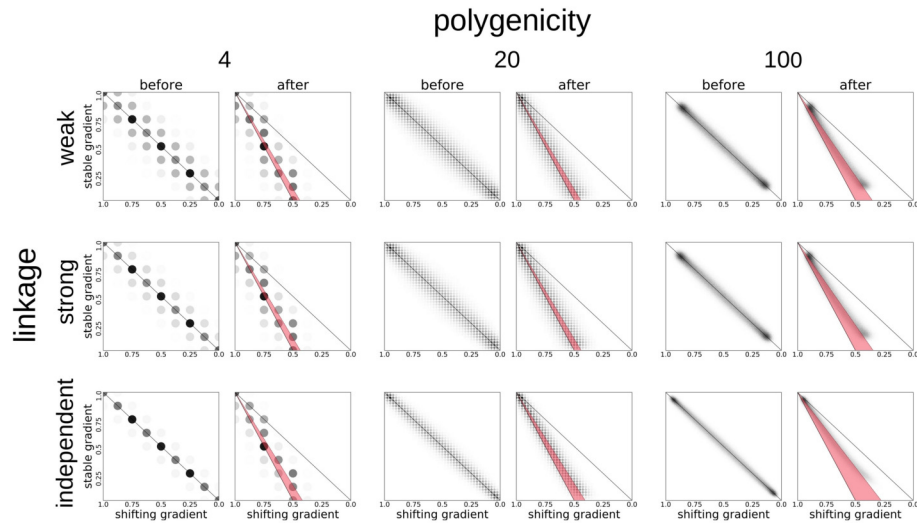


Figure 3: Left: Mean fitness dynamics for all scenarios during the 250 time steps before the climate change period and the 250 time steps during it (with the two periods divided by a red, dashed vertical line marking the onset of the climate change period). Scenarios are organized as in Fig. 2: top and bottom sections for low and high redundancy, with rows in each section representing levels of linkage and columns representing polygenicity. Right: Comparison of climate change-driven changes in mean fitness across scenarios. Null scenarios are plotted on the left in blue, and main scenarios are plotted on the right, in red. Within each plot, scenarios are divided by the number of loci per trait (x-axis) and by the strength of linkage (shade, with darker hues representing stronger linkage).

low redundancy



high redundancy

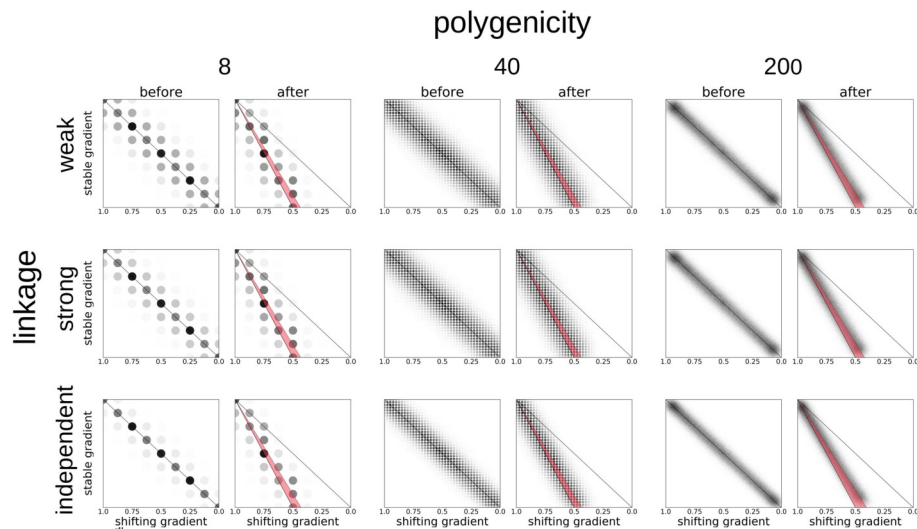


Figure 4: Comparison, across all 18 redundancy scenarios, of the observed versus expected phenotypic shift during the climate change period. Scenarios are organized as in Fig. 2: top and bottom sections representing low and high redundancy, with rows in each section representing levels of linkage and columns (in before-after pairs) representing polygenic. For each scenario, the left ('before') scatterplot shows the distribution of individuals' bivariate phenotypes before climate change begins, whereas the right ('after') scatterplot shows how the distribution has shifted by the end of the climate change period. The trait adapted to the shifting environmental gradient is distributed along the x axis, and the trait adapted to the stable gradient is distributed along the y axis. Scatterplots depict multi-model ensemble results for each scenario. The size and opacity of each point represents the number of individuals exhibiting that bivariate phenotype.

(Note that the gridded arrangement of the points in each scatterplot is a function of the number of loci per trait which, because locus effect sizes are fixed, directly determines the set of attainable, evenly-spaced, discrete phenotypic values. Because fewer loci per trait yields fewer possible phenotypes, individuals are grouped into fewer, larger phenotypic bins in the 4- and 20-locus scenarios.) Solid black lines delineate the shift in the phenotypic distributions' central tendencies that is expected to take place during the climate change period, dotted black lines depict the observed (OLS-fitted) phenotypic distributions' central tendencies at the 'during' and 'after' time steps, and translucent red wedges depict the differences between the expected and observed distributions (i.e., 'phenotypic shortfall', the response variable in our statistical tests).

References

- Ackerly, D. D., Loarie, S. R., Cornwell, W. K., Weiss, S. B., Hamilton, H., Branciforte, R., & Kraft, N. J. B. (2010). The geography of climate change: implications for conservation biogeography. *Diversity and Distributions*, *16*(3), 476–487. doi:10.1111/j.1472-4642.2010.00654.x
- Aitken, S. N., Yeaman, S., Holliday, J. A., Wang, T., & Curtis-McLane, S. (2008). Adaptation, migration or extirpation: climate change outcomes for tree populations. *Evolutionary Applications*, *1*(1), 95–111.
- Aitken, S. N., & Whitlock, M. C. (2013). Assisted Gene Flow to Facilitate Local Adaptation to Climate Change. *Annual Review of Ecology, Evolution, and Systematics*, *44*(1), 367–388. doi:10.1146/annurev-ecolsys-110512-135747
- Barghi, N., Tobler, R., Nolte, V., Jakšić, A. M., Mallard, F., Otte, K. A., ... Schlötterer, C. (2019-2). Genetic redundancy fuels polygenic adaptation in *Drosophila*. *PLoS Biology*, *17*(2), e3000128. doi:10.1371/journal.pbio.3000128
- Barghi, N., Hermisson, J., & Schlötterer, C. (2020). Polygenic adaptation: a unifying framework to understand positive selection. *Nature Reviews Genetics*, *21*(12), 769–781. doi:10.1038/s41576-020-0250-z
- Barton, N. H. (1999). Clines in polygenic traits. *Genetics Research*, *74*(3), 223–236. doi:10.1017/s001667239900422x
- Bell, G., & Gonzalez, A. (2011). Adaptation and evolutionary rescue in metapopulations experiencing environmental deterioration. *Science*, *332*(6035), 1327–1330.
- Bellis, J., Bourke, D., Maschinski, J., Heineman, K., & Dalrymple, S. (2020). Climate suitability as a predictor of conservation translocation failure. *Conservation Biology*, *34*(6), e30013518. doi:10.1111/cobi.13518
- Benes, K. M., & Bracken, M. E. S. (2020). Interactive effects of large- and local-scale environment gradients on phenotypic differentiation. *Ecology*, *101*(8), e03078. doi:10.1002/ecy.3078
- Bomblies, K., & Peichel, C. L. (2022). Genetics of adaptation. *Proceedings of the National Academy of Sciences*, *119*(30), e2122152119.
- Boyle, E. A., Li, Y. I., & Pritchard, J. K. (2017). An Expanded View of Complex Traits: From Polygenic to Omnigenic. *Cell*, *169*(7), 1177–1186. doi:10.1016/j.cell.2017.05.038
- Bürger, R., & Gimelfarb, A. (2002). Fluctuating environments and the role of mutation in maintaining quantitative genetic variation. *Genetics Research*, *80*(1), 31–46.
- Capblancq, T., Fitzpatrick, M. C., Bay, R. A., Exposito-Alonso, M., & Keller, S. R. (2020). Genomic Prediction of (Mal)Adaptation Across Current and Future Climatic Landscapes. *Annual Review of Ecology, Evolution, and Systematics*, *51*(1), 1–25. doi:10.1146/annurev-ecolsys-020720-042553
- Chevin, L.-M., Lande, R., & Mace, G. M. (2010). Adaptation, Plasticity, and Extinction in a Changing Environment: Towards a Predictive Theory. *PLoS Biology*, *8*(4), e1000357. doi:10.1371/journal.pbio.1000357
- Christiansen, F. B., Otto, S. P., Bergman, A., & Feldman, M. W. (1998). Waiting with and without recombination: the time to production of a double mutant. *Theoretical Population Biology*, *53*(3), 199–215.

- Crimmins, S. M., Dobrowski, S. Z., Greenberg, J. A., Abatzoglou, J. T., & Mynsberge, A. R. (2011). >Changes in Climatic Water Balance Drive Downhill Shifts in Plant Species' Optimum Elevations. *Science*, 331(6015), 324–327. doi:10.1126/science.1199040
- Daly, C., Conklin, D. R., & Unsworth, M. H. (2010). Local atmospheric decoupling in complex topography alters climate change impacts. *International Journal of Climatology*, 30(12), 1857–1864. doi:10.1002/joc.2007
- Feder, J. L., Gejji, R., Yeaman, S., & Nosil, P. (2012). Establishment of new mutations under divergence and genome hitchhiking. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1587), 461–474.
- Felsenstein, J. (1976). The theoretical population genetics of variable selection and migration. *Annual Review of Genetics*, 10(1), 253–280.
- Fitzpatrick, M. C., Blois, J. L., Williams, J. W., Nieto-Lugilde, D., Maguire, K. C., & Lorenz, D. J. (2018). How will climate novelty influence ecological forecasts? Using the Quaternary to assess future reliability. *Global Change Biology*, 24(8), 3575–3586. doi:10.1111/gcb.14138
- Guillaume, F. (2011). Migration-induced phenotypic divergence: the migration--selection balance of correlated traits. *Evolution: International Journal of Organic Evolution*, 65(6), 1723–1738.
- Haldane, J. B. S. (1930). A mathematical theory of natural and artificial selection.(Part VI, Isolation.). *Mathematical Proceedings of the Cambridge Philosophical Society*, 26:2, 220–230. Cambridge University Press.
- Harrisson, K. A., Pavlova, A., Telonis-Scott, M., & Sunnucks, P. (2014). Using genomics to characterize evolutionary potential for conservation of wild populations. *Evolutionary Applications*, 7(9), 1008–1025. doi:10.1111/eva.12149
- Hornik, K., & Grün, B. (2014). movMF: An R Package for Fitting Mixtures of von Mises-Fisher Distributions. *Journal of Statistical Software*, 58(10), 1–31. doi:10.18637/jss.v058.i10
- Kelleher, J., Thornton, K. R., Ashander, J., & Ralph, P. L. (2018). Efficient pedigree recording for fast population genetics simulation. *PLOS Computational Biology*, 14(11), e1006581. doi:10.1371/journal.pcbi.1006581
- Kondrashov, A. S., & Yampolsky, L. Y. (1996). High genetic variability under the balance between symmetric mutation and fluctuating stabilizing selection. *Genetics Research*, 68(2), 157–164.
- Láruson, Á. J., Yeaman, S., & Lotterhos, K. E. (2020). The importance of genetic redundancy in evolution. *Trends in Ecology & Evolution*, 35(9), 809–822.
- Le Corre, V., & Kremer, A. (2012). The genetic differentiation at quantitative trait loci under local adaptation. *Molecular Ecology*, 21(7), 1548–1566.
- Lenormand, T. (2002). Gene flow and the limits to natural selection. *Trends in Ecology & Evolution*, 17(4), 183–189. doi:10.1016/s0169-5347(02)02497-7
- Loarie, S. R., Duffy, P. B., Hamilton, H., Asner, G. P., Field, C. B., & Ackerly, D. D. (2009). The velocity of climate change. *Nature*, 462(7276), 1052–1055. doi:10.1038/nature08649
- Manceau, M., Domingues, V. S., Linnen, C. R., Rosenblum, E. B., & Hoekstra, H. E. (2010). Convergence in pigmentation at multiple levels: mutations, genes and function.

- Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1552), 2439–2450.
- Martin, A., & Orgogozo, V. (2013). The loci of repeated evolution: a catalog of genetic hotspots of phenotypic variation. *Evolution*, 67(5), 1235–1250.
- Murray, G. G. R., Soares, A. E. R., Novak, B. J., Schaefer, N. K., Cahill, J. A., Baker, A. J., Demboski, J.R., Doll, A., Da Fonseca, R. R., Fulton, T. L., Gilbert, M. T. (2017). Natural selection shaped the rise and fall of passenger pigeon genomic diversity. *Science*, 358(6365), 951–954.
- Nicotra, A. B., Beever, E. A., Robertson, A. L., Hofmann, G. E., & O’Leary, J. (2015). Assessing the components of adaptive capacity to improve conservation and management efforts under global change. *Conservation Biology*, 29(5), 1268–1278. doi:10.1111/cobi.12522
- Orr, H. A. (1998). The population genetics of adaptation: the distribution of factors fixed during adaptive evolution. *Evolution*, 52(4), 935–949.
- Pritchard, J. K., & Di Rienzo, A. (2010a). Adaptation--not by sweeps alone. *Nature Reviews Genetics*, 11(10), 665–667.
- Pritchard, J. K., Pickrell, J. K., & Coop, G. (2010b). The genetics of human adaptation: hard sweeps, soft sweeps, and polygenic adaptation. *Current Biology*, 20(4), R208–R215.
- R Core Team. (2021). *R: A Language and Environment for Statistical Computing*. Retrieved from R Foundation for Statistical Computing website: <https://www.R-project.org/>
- Rees, J. S., Castellano, S., & Andrés, A. M. (2020). The genomics of human local adaptation. *Trends in Genetics*, 36(6), 415–428.
- Rockman, M. V. (2012). The QTN program and the alleles that matter for evolution: all that’s gold does not glitter. *Evolution: International Journal of Organic Evolution*, 66(1), 1–17.
- Savolainen, O., Lascoux, M., & Merilä, J. (2013). Ecological genomics of local adaptation. *Nature Reviews Genetics*, 14(11), 807–820.
- Schiffers, K., Bourne, E. C., Lavergne, S., Thuiller, W., & Travis, J. M. J. (2013). Limited evolutionary rescue of locally adapted populations facing climate change. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1610), 20120083. doi:10.1098/rstb.2012.0083
- Sella, G., & Barton, N. H. (2019). Thinking about the evolution of complex traits in the era of genome-wide association studies. *Annual Review of Genomics and Human Genetics*, 20, 461–493.
- Slatkin, M. (1987). Gene flow and the geographic structure of natural populations. *Science*, 236(4803), 787–792.
- Terasaki Hart, D. E., Bishop, A. P., & Wang, I. J. (2021). Geonomics: forward-time, spatially explicit, and arbitrarily complex landscape genomic simulations. *Molecular Biology and Evolution*, msab175-. doi:10.1093/molbev/msab175
- Thomas, C. D. (2010). Climate, climate change and range boundaries. *Diversity and Distributions*, 16(3), 488–495.
- Thompson, K. A. (2020). Experimental Hybridization Studies Suggest That Pleiotropic Alleles Commonly Underlie Adaptive Divergence between Natural Populations. <https://doi.org/10.1086/708722>, 196(1), 000–000. doi:10.1086/708722
- Tigano, A., & Friesen, V. L. (2016). Genomics of local adaptation with gene flow. *Molecular Ecology*, 25(10), 2144–2164.

- Van Rossum, G., & Drake, F. L., Jr. (1995). *Python reference manual*. Centrum voor Wiskunde en Informatica Amsterdam.
- Vilas, A., Pérez-Figueroa, A., Quesada, H., & Caballero, A. (2015). Allelic diversity for neutral markers retains a higher adaptive potential for quantitative traits than expected heterozygosity. *Molecular Ecology*, *24*(17), 4419–4432. doi:10.1111/mec.13334
- Wade, A. A., Hand, B. K., Kovach, R. P., Luikart, G., Whited, D. C., & Muhlfeld, C. C. (2017). Accounting for adaptive capacity and uncertainty in assessments of species' climate-change vulnerability. *Conservation Biology*, *31*(1), 136–149. doi:10.1111/cobi.12764
- Wang, I. J., & Bradburd, G. S. (2014). Isolation by environment. *Molecular Ecology*, *23*(23), 5649–5662. doi:10.1111/mec.12938
- Weiss-Lehman, C., & Shaw, A. K. (2020-1). Spatial Population Structure Determines Extinction Risk in Climate-Induced Range Shifts. *The American Naturalist*, *195*(1), 31–42. doi:10.1086/706259
- Williams, J. W., Jackson, S. T., & Kutzbach, J. E. (2007). Projected distributions of novel and disappearing climates by 2100 AD. *Proceedings of the National Academy of Sciences*, *104*(14), 5738–5742. doi:10.1073/pnas.0606292104
- Williams, J. W., & Jackson, S. T. (2007). Novel climates, no-analog communities, and ecological surprises. *Frontiers in Ecology and the Environment*, *5*(9), 475–482. doi:10.1890/070037
- Wright, S. (1931). Evolution in Mendelian populations. *Genetics*, *16*, 97–159.
- Yeaman, S., & Whitlock, M. C. (2011). THE GENETIC ARCHITECTURE OF ADAPTATION UNDER MIGRATION–SELECTION BALANCE. *Evolution*, *65*(7), 1897–1911. doi:10.1111/j.1558-5646.2011.01269.x
- Yeaman, S. (2015). Local adaptation by alleles of small effect. *The American Naturalist*, *186*(S1), S74–S89.
- Yeaman, S. (2022). Evolution of polygenic traits under global vs local adaptation. *Genetics*, *220*(1). doi:10.1093/genetics/iyab134

Supplemental Figures

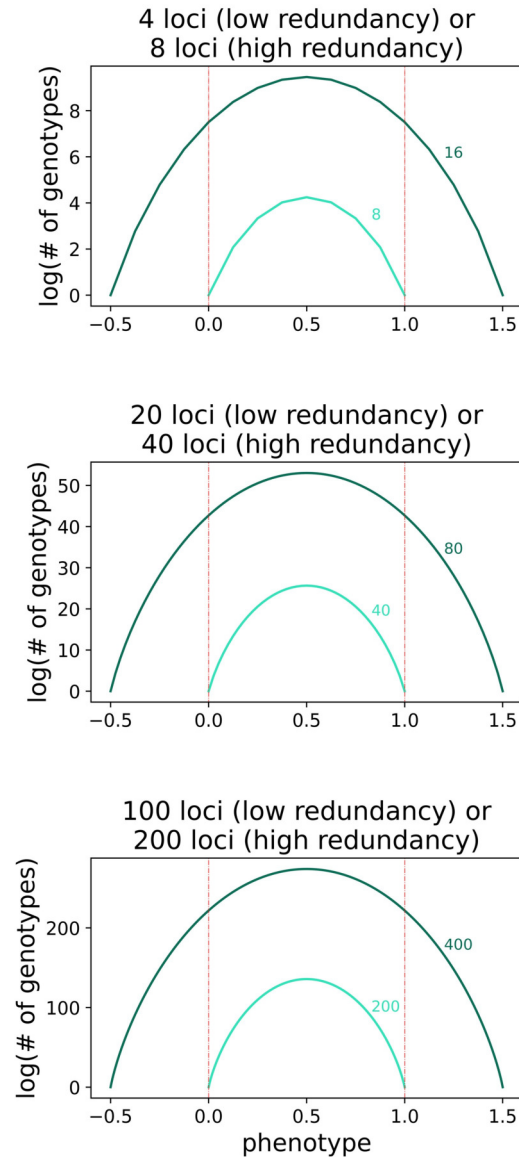


Figure S1: Depiction of genotypic redundancy for all simulated polygenicity values. Phenotypic values are plotted along the x-axis, and the natural log of the number of genotypes that yield each phenotypic value is plotted along the y-axis. Polygenicities corresponding to low-redundancy scenarios are plotted and labeled in light teal, and those corresponding to high-redundancy scenarios in dark teal. The minimum and maximum environmental values on the landscape are represented by dotted vertical lines. The number of genotypes corresponding to each phenotype is calculated using a custom adaptation of Eqxn. ii, Box 1 in (Láruson *et al.* 2020) implemented for a diploid species and fitted to the numerical conventions used by Geomics.

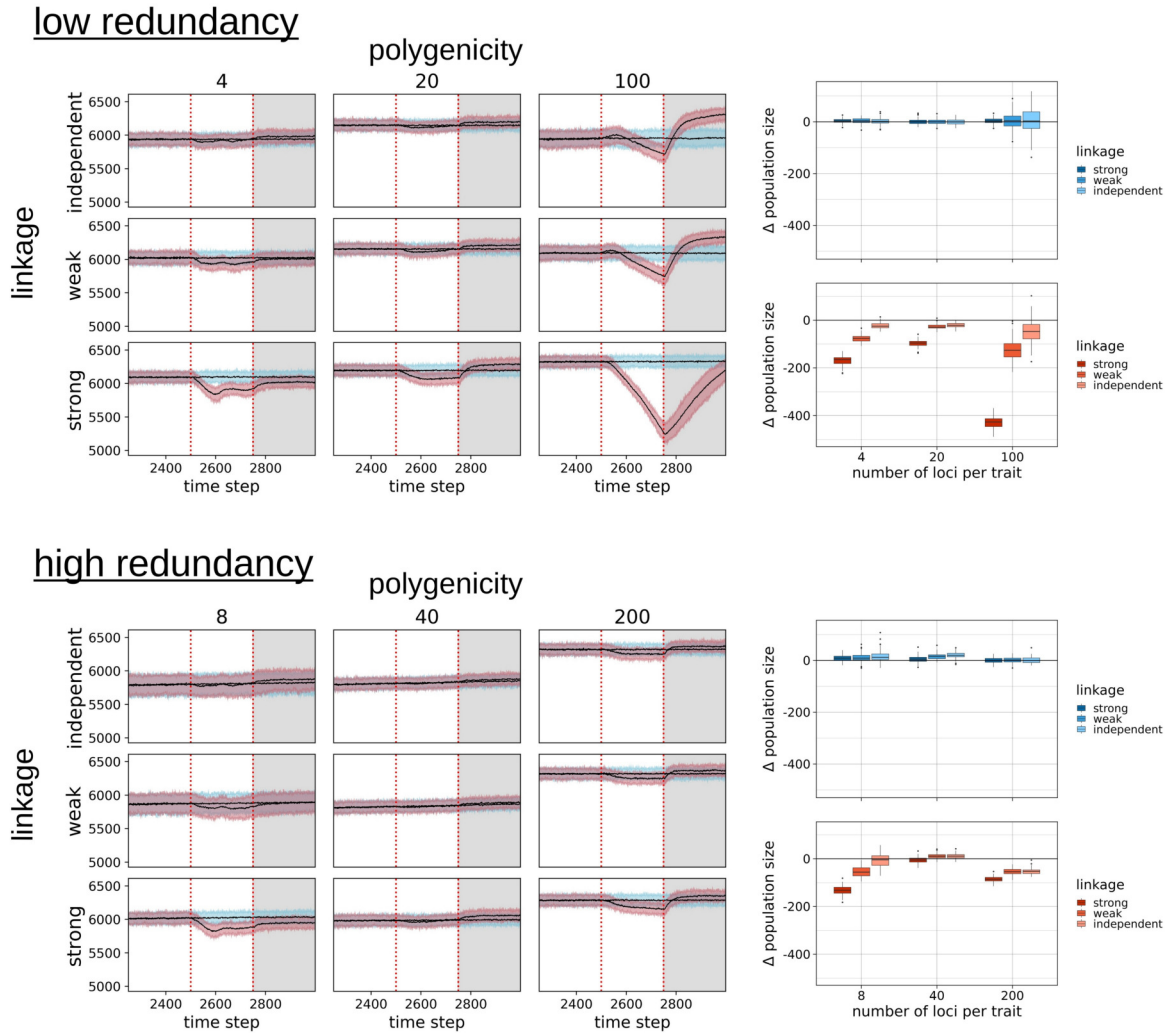


Figure S2: Left: Population size dynamics for all scenarios during the 250 time steps before the climate change period and the 250 time steps during it (with the two periods divided by a red, dashed vertical line marking the onset of the climate change period). Scenarios are organized as in Fig. 2: top and bottom sections representing low and high redundancy, with rows in each section representing levels of linkage and columns representing polygenicity. Both means (black lines) and variability envelopes (5th percentile to 95th percentile) are shown, with scenario type depicted by color (main scenarios in red, null scenarios in blue). Right: Comparison of climate change-driven changes in mean population size across scenarios. Null scenarios are plotted on the left in blue, and main scenarios are plotted on the right, in red. Within each plot, scenarios are divided by the number of loci per trait (x-axis) and by the strength of linkage (shade, with darker hues representing stronger linkage).

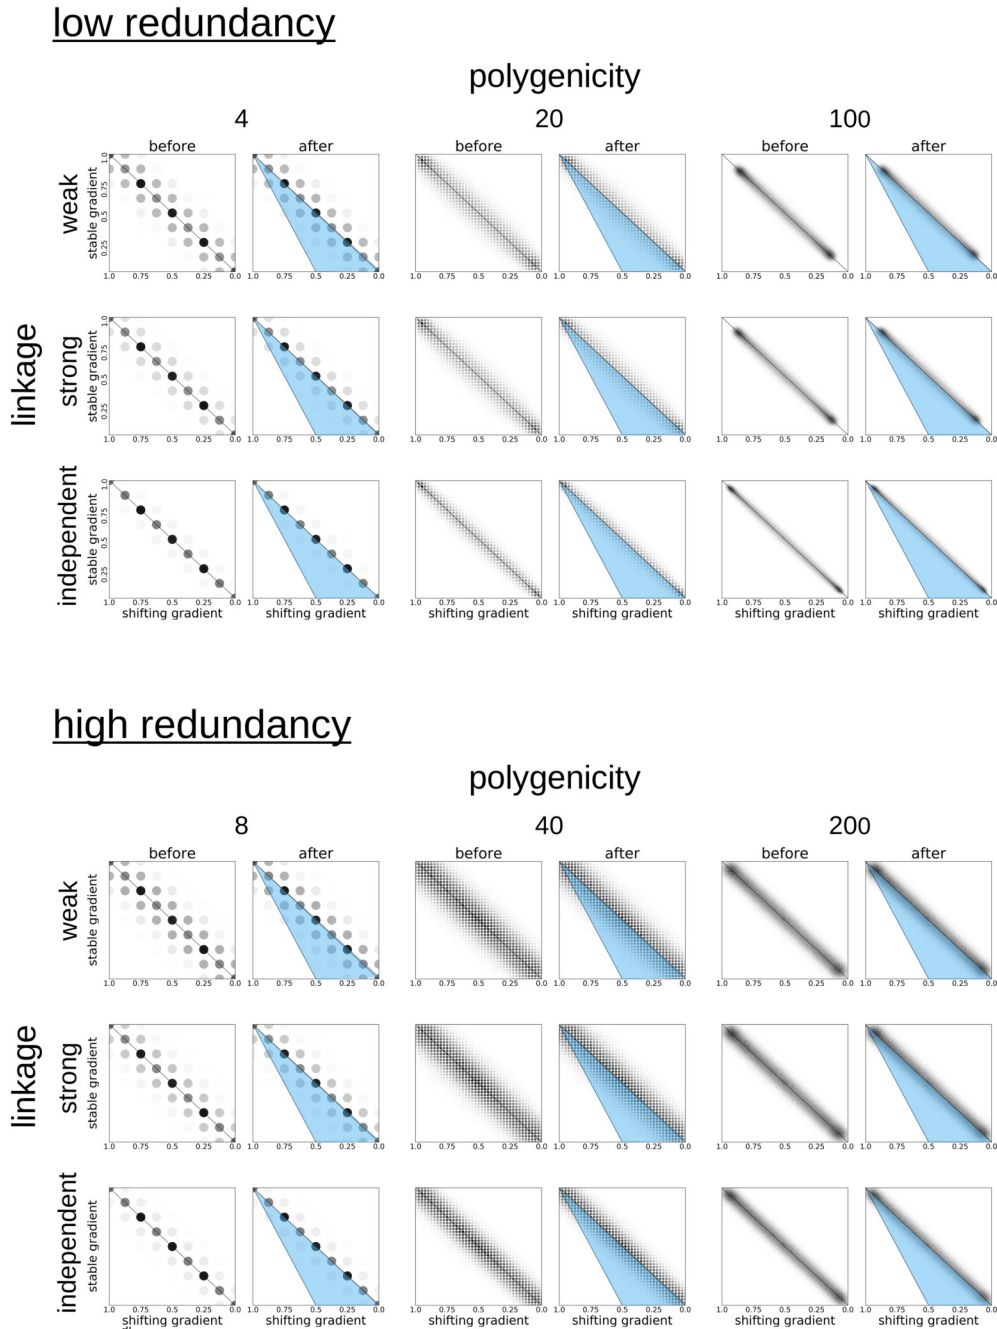


Figure S3: Maps depicting shifts in population density during climate change for all null simulations. Scenarios are organized as in Fig. 2: top and bottom sections representing low and high redundancy, with rows in each section representing levels of linkage and columns (in before-after pairs) representing polygenicity. As well as showing local extinction in the low-redundancy, high-polygenicity, strong-linkage scenario (bottom right of top section), these maps also show moderate simulation edge effects and density banding in the low-polygenicity scenarios because of the mismatch between environmental and phenotypic resolutions.

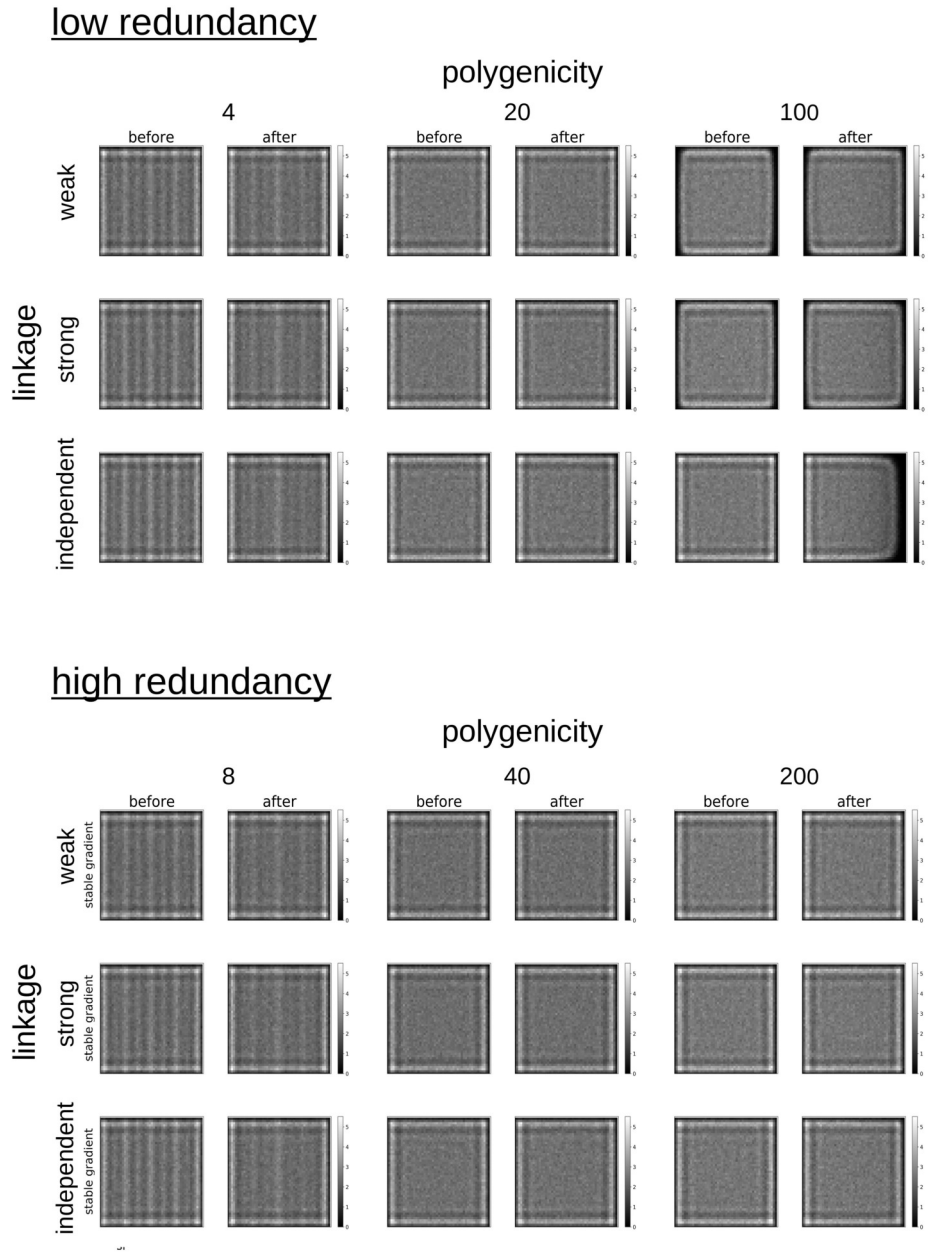


Figure S4: Maps depicting shifts in population density during climate change for all scenarios.

Appendix 1: Template parameters file used for all simulations (with key parameters annotated in green)

```
# template_params.py

import numpy as np
import pandas as pd
import os

# set up the landscape
b4 = np.vstack([np.linspace(1, 0, 50) for _ in range(50)])
af = np.vstack([np.linspace(1, 0.5, 50) for _ in range(50)])
stable = np.vstack([np.linspace(1, 0, 50) for _ in range(50)])
K = np.ones((50,50))

# NOTE: time_steps.CSV USED TO SET CLIMATE CHANGE
#       TO START AT TIME STEP 2500
#       AND FINISH AT TIME STEP 2750
if os.getcwd().split('/')[1] == 'home':
    steps = pd.read_csv('/home/deth/Desktop/CAL/research/projects/sim/'
                       'ch2/climate_change_adaptation_and_genomic_arch/sim/'
                       'time_steps.CSV')
else:
    steps = pd.read_csv('/global/scratch/users/drewhart/'
                       'ch2/climate_change_adaptation_and_genomic_arch/sim/'
                       'time_steps.CSV')
# set time when environmental change begins
change_T = int(steps[steps['name']=='start']['num'].values[0])
# set time when environmental change ends
T = int(steps[steps['name']=='end']['num'].values[0])

# show the landscape, for debugging, if requested
debug_landscape = False
if debug_landscape:
    import matplotlib.pyplot as plt
    fig = plt.figure()
    ax1 = fig.add_subplot(221)
    im1 = ax1.imshow(b4, cmap='spring', vmin=0, vmax=1)
    plt.colorbar(im1)
    ax3 = fig.add_subplot(223)
    im3 = ax3.imshow(af, cmap='spring', vmin=0, vmax=1)
    plt.colorbar(im3)
    ax2 = fig.add_subplot(222)
    im2 = ax2.imshow(stable, cmap='winter', vmin=0, vmax=1)
    plt.colorbar(im2)
    ax4 = fig.add_subplot(224)
    im4 = ax4.imshow(K, cmap='autumn', vmin=0, vmax=1)
    plt.colorbar(im4)
    plt.show()

# This is a parameters file generated by Geonomics
# (by the gnx.make_parameters_file() function).

#       #       : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : #
#       #: : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : #
#       #: : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : #
#       #: : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : #
#       # GGGGG :EEEE: 00000  NN  NN  00000  MM  MM  IIIIII  CCCCC  SSSSS #
# GG  EE  00  00  NNN  NN  00  00  MM  MM  II  CC  SS  #
# GG  EE  00  00  NN  N  NN  00  00  MMM  MMM  II  CC  SSSSS #
# GG  GGG  EEEE  00  00  NN  NNN  00  00  MM  M  MM  II  CC  SS  #
# GG  G  EE  00  00  NN  NN  00  00  MM  MM  II  CC  SSS  #
# GGGGG :EEEE: 00000  NN  NN  00000  MM  MM  IIIIII  CCCCC  SSSSS #
#       #       : : : : : : : : : : : : : : : : : : : : : : : : : : : : : #
#       #: : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : : #
```

```

#   :::           ::::  ::           ::::  #
#   ::           ::::  #           #
#   #           #           #           #
#   #   ::   ::   ::           #           #

params = {
#####

#####
#### LANDSCAPE ####
#####

  'landscape': {

#####
#### main ####
#####

    'main': {
      #x,y (a.k.a. j,i) dimensions of the Landscape
      'dim':          (50,50),
      #x,y resolution of the Landscape
      'res':          (1,1),
      #x,y coords of upper-left corner of the Landscape
      'ulc':          (0,0),
      #projection of the Landscape
      'prj':          None,
    }, # <END> 'main'

#####
#### layers ####
#####

    'layers': {

      #layer name (LAYER NAMES MUST BE UNIQUE!)
      'shift': {

#-----#
#--- layer num. 0: init parameters ---#
#-----#

        #initiating parameters for this layer
        'init': {

          #parameters for a 'defined'-type Layer
          'defined': {
            #raster to use for the Layer
            'rast':          b4,
            #point coordinates
            'pts':           None,
            #point values
            'vals':          None,
            #interpolation method {None, 'linear', 'cubic',
            #'nearest'}
            'interp_method': None,

            }, # <END> 'defined'

          }, # <END> 'init'

#-----#
#--- layer num. 0: change parameters ---#
#-----#

        #landscape-change events for this Layer
        'change': {

          0: {
            #array or file for final raster of event, or directory

```

```

        #of files for each stepwise change in event
        'change_rast':      af,
        #starting timestep of event
        'start_t':         change_T,
        #ending timestep of event
        'end_t':           T,
        #number of stepwise changes in event
        'n_steps':         T-change_T,
    }, # <END> event 0

    }, # <END> 'change'

}, # <END> layer num. 0

#layer name (LAYER NAMES MUST BE UNIQUE!)
'stable': {

#-----#
#--- layer num. 1: init parameters ---#
#-----#

#initiating parameters for this layer
'init': {

    #parameters for a 'defined'-type Layer
    'defined': {
        #raster to use for the Layer
        'rast':      stable,
        #point coordinates
        'pts':       None,
        #point values
        'vals':      None,
        #interpolation method {None, 'linear', 'cubic',
        #'nearest'}
        'interp_method':      None,

    }, # <END> 'defined'

    }, # <END> 'init'

}, # <END> layer num. 1

#layer name (LAYER NAMES MUST BE UNIQUE!)
'K': {

#-----#
#--- layer num. 2: init parameters ---#
#-----#

#initiating parameters for this layer
'init': {

    #parameters for a 'defined'-type Layer
    'defined': {
        #raster to use for the Layer
        'rast':      K,
        #point coordinates
        'pts':       None,
        #point values
        'vals':      None,
        #interpolation method {None, 'linear', 'cubic',
        #'nearest'}
        'interp_method':      None,

    }, # <END> 'defined'

```

```

    }, # <END> 'init'
  }, # <END> layer num. 2

#layer name (LAYER NAMES MUST BE UNIQUE!)
'move': {

#-----#
#--- layer num. 2: init parameters ---#
#-----#

#initiating parameters for this layer
'init': {

#parameters for a 'defined'-type Layer
'defined': {
  #raster to use for the Layer
  'rast': np.ones((50,50)),
  #point coordinates
  'pts': None,
  #point values
  'vals': None,
  #interpolation method {None, 'linear', 'cubic',
  #'nearest'}
  'interp_method': None,

  }, # <END> 'defined'

}, # <END> 'init'

}

#### NOTE: Individual Layers' sections can be copy-and-pasted (and
#### assigned distinct keys and names), to create additional Layers.

} # <END> 'layers'
}, # <END> 'landscape'

#####

#####
#### COMMUNITY ####
#####
'comm': {
  'species': {

#species name (SPECIES NAMES MUST BE UNIQUE!)
'spp_0': {

#-----#
#--- spp num. 0: init parameters ---#
#-----#

'init': {
  #starting number of individs
  'N': 1000,
  #carrying-capacity Layer name
  'K_layer': 'K',
  #multiplicative factor for carrying-capacity layer
  'K_factor': 2.5,
  }, # <END> 'init'

#-----#

```



```

#--- spp num. 0: mating parameters ---#
#-----#

'mating'      : {
  #age(s) at sexual maturity (if tuple, female first)
  'repro_age':      0,
  #whether to assign sexes
  'sex':           False,
  #ratio of males to females
  'sex_ratio':     1/1,
  #whether P(birth) should be weighted by parental dist
  'dist_weighted_birth': False,
  #intrinsic growth rate
  'R':            0.5,
  #intrinsic birth rate (MUST BE 0<=b<=1)
  'b':            0.5,
  #expectation of distr of n offspring per mating pair
  'n_births_distr_lambda': 1,
  #whether n births should be fixed at n_births_dist_lambda
  'n_births_fixed': True,
  #radius of mate-search area
  'mating_radius': 5,
  #whether individs should choose nearest neighs as mates
  'choose_nearest_mate': False,
  #whether mate-choice should be inverse distance-weighted
  'inverse_dist_mating': False,
}, # <END> 'mating'

#-----#
#--- spp num. 0: mortality parameters ---#
#-----#

'mortality'   : {
  #maximum age
  'max_age':      None,
  #min P(death) (MUST BE 0<=d_min<=1)
  'd_min':        0,
  #max P(death) (MUST BE 0<=d_max<=1)
  'd_max':        1,
  #width of window used to estimate local pop density
  'density_grid_window_width': None,
}, # <END> 'mortality'

#-----#
#--- spp num. 0: movement parameters ---#
#-----#

'movement': {
  #whether or not the species is mobile
  'move':         True,
  #mode of distr of movement direction
  'direction_distr_mu': 0,
  #concentration of distr of movement direction
  'direction_distr_kappa': 0,
  #mean of distr of movement distance
  'movement_distance_distr_param1': 0.25,
  #variance of distr of movement distance
  'movement_distance_distr_param2': 0.5,
  #movement distance distr to use ('lognormal','levy','wald')
  'movement_distance_distr': 'wald',
  #mean of distr of dispersal distance
  'dispersal_distance_distr_param1': 0.5,
  #variance of distr of dispersal distance
  'dispersal_distance_distr_param2': 0.5,
  #dispersal distance distr to use ('lognormal','levy','wald')
  'dispersal_distance_distr': 'wald',
}, # <END> 'movement'

```

```

#-----#
#--- spp num. 0: genomic architecture parameters ---#
#-----#

'gen_arch': {
  #file defining custom genomic arch
  'gen_arch_file':      None,
  #num of loci
  'L':                  1000,
  #value to use for fixed starting allele freqs (None to draw)
  'start_p_fixed':     0.5,
  #whether to start neutral locus freqs at 0
  'start_neut_zero':   True,
  #genome-wide per-base neutral mut rate (0 to disable)
  'mu_neut':           0,
  #genome-wide per-base deleterious mut rate (0 to disable)
  'mu_delet':          0,
  #shape of distr of deleterious effect sizes
  'delet_alpha_distr_shape': 0.2,
  #scale of distr of deleterious effect sizes
  'delet_alpha_distr_scale': 0.2,
  #NOTE: MAIN SCRIPT OVERRIDES THE FOLLOWING TWO PARAMS
  #TO SET RECOMBINATION RATES TO A FIXED VALUE OF
  #0.5, 0.05, OR 0.005 FOR INDEPENDENT,
  #WEAK, OR STRONG LINKAGE VALUES
  #alpha of distr of recomb rates
  'r_distr_alpha':     1000,
  #beta of distr of recomb rates
  'r_distr_beta':      1e3,
  #whether loci should be dominant (for allele '1')
  'dom':               False,
  #whether to allow pleiotropy
  'pleiotropy':        False,
  #custom fn for drawing recomb rates
  'recomb_rate_custom_fn': None,
  #number of recomb paths to hold in memory
  'n_recomb_paths_mem': int(1e4),
  #total number of recomb paths to simulate
  'n_recomb_paths_tot': int(1e5),
  #num of crossing-over events (i.e. recombs) to simulate
  'n_recomb_sims':     100_000,
  #whether to generate recombination paths at each timestep
  'allow_ad_hoc_recomb': False,
  #whether to save mutation logs
  'mut_log':           False,
  #whether to jitter recomb bps, to correctly track num_trees
  'jitter_breakpoints': False,
  #whether to use tskit (to record full spatial pedigree)
  'use_tskit': True,
  #time step interval for simplication of tskit tables
  'tskit_simp_interval': 100,

  'traits': {

    #-----#
    #---trait 0 parameters ---#
    #-----#
    #trait name (TRAIT NAMES MUST BE UNIQUE!)
    'trait_0': {
      #trait-selection Layer name
      'layer':      'shift',
      #polygenic selection coefficient
      'phi':        1,
      #NOTE: MAIN SCRIPT CHANGES NEXT PARAM TO 4, 20, OR 100
      #FOR LOW-REDUNDANCY SCENARIOS OF DIFF. POLYGENICITY,
      #OR 8, 40, OR 200 FOR HIGH-REDUNDANCY SCENARIOS
      #number of loci underlying trait

```

```

'n_loci':          50,
#mutation rate at loci underlying trait
'mu':              0,
#mean of distr of effect sizes
'alpha_distr_mu' : 0,
#variance of distr of effect size
'alpha_distr_sigma': 0,
#max allowed magnitude for an alpha value
'max_alpha_mag':  None,
#curvature of fitness function
'gamma':           1,
#whether the trait is universally advantageous
'univ_adv':        False
}, # <END> trait 0

#-----#
#---trait 1 parameters ---#
#-----#
#trait name (TRAIT NAMES MUST BE UNIQUE!)
'trait_1': {
  #trait-selection Layer name
  'layer':          'stable',
  #polygenic selection coefficient
  'phi':            1,
  #NOTE: MAIN SCRIPT CHANGES NEXT PARAM TO 4, 20, OR 100
  #FOR LOW-REDUNDANCY SCENARIOS OF DIFF. POLYGENICITY,
  #OR 8, 40, OR 200 FOR HIGH-REDUNDANCY SCENARIOS
  #number of loci underlying trait
  'n_loci':         50,
  #mutation rate at loci underlying trait
  'mu':             0,
  #mean of distr of effect sizes
  'alpha_distr_mu' : 0,
  #variance of distr of effect size
  'alpha_distr_sigma': 0,
  #max allowed magnitude for an alpha value
  'max_alpha_mag':  None,
  #curvature of fitness function
  'gamma':          1,
  #whether the trait is universally advantageous
  'univ_adv':       False
}, # <END> trait 1

#### NOTE: Individual Traits' sections can be copy-and-pasted (and
#### assigned distinct keys and names), to create additional Traits.

    }, # <END> 'traits'

  }, # <END> 'gen_arch'

}, # <END> spp num. 0

#### NOTE: individual Species' sections can be copy-and-pasted (and
#### assigned distinct keys and names), to create additional Species.

    }, # <END> 'species'

  }, # <END> 'comm'

#####
#####
#### MODEL ####
#####

```

```

'model': {
  # NOTE: NEXT PARAM OVERRIDDEN BY MAIN SCRIPT
  #total Model runtime (in timesteps)
  'T': 100000,
  #min burn-in runtime (in timesteps)
  'burn_T': 30,
  #seed number
  'num': None,

  #####
  #### iterations parameters ####
  #####
  'its': {
    #num iterations
    'n_its': 1,
    #whether to randomize Landscape each iteration
    'rand_landscape': False,
    #whether to randomize Community each iteration
    'rand_comm': False,
    #whether to randomize GenomicArchitectures each iteration
    'rand_genarch': True,
    #whether to burn in each iteration
    'repeat_burn': False,
  }, # <END> 'iterations'

  #####
  #### data-collection parameters ####
  #####
  'data': {
    'sampling': {
      #sampling scheme {'all', 'random', 'point', 'transect'}
      'scheme': 'all',
      #sample size at each point, for point & transect sampling
      'n': 1000,
      #coords of collection points, for point sampling
      'points': None,
      #coords of transect endpoints, for transect sampling
      'transect_endpoints': None,
      #num points along transect, for transect sampling
      'n_transect_points': None,
      #collection radius around points, for point & transect sampling
      'radius': None,
      #when to collect data
      'when': [change_T-1,
              int((change_T-1+T-1)/2),
              T-1],
      #whether to save current Layers when data is collected
      'include_landscape': False,
      #whether to include fixed loci in VCF files
      'include_fixed_sites': True,
    },
    'format': {
      #format for genetic data {'vcf', 'fasta'}
      'gen_format': ['vcf'],
      #format for vector geodata {'csv', 'shapefile', 'geojson'}
      'geo_vect_format': 'csv',
      #format for raster geodata {'geotiff', 'txt'}
      'geo_rast_format': 'geotiff',
      #format for files containing non-neutral loci
      'nonneut_loc_format': 'csv',
    },
  }, #<END> 'data'

} # <END> 'model'

}# <END> params

```

Chapter 3

Global phenology mapping reveals regional heterogeneity, intercontinental convergence, and asynchrony in Mediterranean and tropical climates

by

Drew E. Terasaki Hart¹

Lauren Di Maggio²

Thao-Nguyen Bui¹

Ian J. Wang¹

¹ Department of Environmental Science, Policy, and Management, University of California, Berkeley CA, USA

² Department of Statistics, University of California, Berkeley CA, USA

Abstract

As I have demonstrated in chapter 2, the spatial nature of populations can have an important influence on their temporal evolutionary dynamics. Yet, the reverse is also true: The complex temporal dynamics of populations can play an important, albeit underappreciated role in determining the spatial dynamics and patterns of evolution. As one example, geographic patterns in the timing of plant productivity – i.e., in land surface phenology (LSP) – not only convey critical information about environmental controls on plant function and carbon cycling but also indicate the potential for spatial asynchrony in reproductive phenology, and thus for genetic isolation and divergence between conspecific populations. Despite the major implications for ecophysiology, ecosystems ecology, and evolutionary biology, most LSP mapping methodologies struggle to describe the full diversity of terrestrial biome phenologies, and LSP asynchrony mapping methodologies are scarcely developed. Here, we first apply a multivariate, generalized, and robustly-validated LSP mapping methodology, based on simple harmonic regression, to a global time series archive of MODIS near-infrared reflectance of vegetation (NIRV, a proxy of plant productivity). The result reveals surprising LSP diversity, including both regional patterns of heterogeneity that are corroborated by prior findings and intercontinental patterns of convergence that recapitulate major bioclimatic and biogeographic gradients. Next, we present a global map of LSP asynchrony, then use machine learning to explore regional variability in potential climatic and physiographic drivers. We find that high LSP asynchrony occurs in the world's five Mediterranean climate regions, where it appears driven by precipitation asynchrony and spatial variability in vegetation structure, and in tropical montane regions, where minimum temperature asynchrony and precipitation asynchrony appear to be interacting drivers. Finally, we use an ensemble of regressions within high-asynchrony regions to demonstrate that phenological asynchrony between climatically similar sites is most frequent at lower latitudes, suggesting that phenological asynchrony is most likely to cause allochronic divergence in the tropics.

Introduction and Background

Terrestrial plant communities exhibit wide variation in annual rhythms of productivity, a collective result of species' adaptive responses to the broad diversity of abiotic environments (Tang *et al.* 2016). The spatiotemporal patterns this creates, known as land surface phenology (LSP), convey rich ecophysiological information on the bioclimatic controls of canopy function (Adole *et al.* 2019) and thus the global carbon cycle. Characterizing these patterns in order to study the nature and drivers of LSP continues to grow in importance as global change advances and phenologies shift in response (Piao *et al.* 2019; Richardson *et al.* 2013; Tang *et al.* 2016), portending wide-ranging consequences for conservation and land management (Ettinger *et al.* 2022; Morellato *et al.* 2016). A deeper understanding of LSP is crucial for improving the ability of earth system models to simulate current and future phenologies (Fu *et al.* 2020; Peñuelas *et al.* 2009; Richardson *et al.* 2012; Tang *et al.* 2016) and for improving understanding of the feedbacks of phenological shifts on components of the global climate system (Bonan 2008; Fu *et al.* 2020; Lovejoy and Nobre 2018; Notaro *et al.* 2011; Salati *et al.* 1979) and carbon cycle (Richardson *et al.* 2013), including key uncertainties (S. Wang *et al.* 2022) like interannual variation in xeric biome productivity (Broich *et al.* 2014; Buitenwerf *et al.* 2015; Poulter *et al.* 2014) and spatial variation in tropical forest source-sink dynamics (Gatti *et al.* 2021; Harris *et al.* 2014). Yet, most LSP research has focused on the estimation of univariate statistics designed for unimodal time series, limiting our ability to describe the full complexity of Earth's LSP patterns, especially in many tropical and arid regions characterized by subtle or multimodal seasonal patterns that remain poorly understood (Fu *et al.* 2020; Garonna *et al.* 2016; Lloyd 1990; Richardson *et al.* 2013; Tang *et al.* 2016; L. Zeng *et al.* 2020; Sakai and Kitajima 2019). The historical lack of strong remote sensing proxies of photosynthesis has compounded this limitation, forcing prior analyses to use standard vegetation indices (e.g., the enhanced or normalized difference vegetation indices) with limited sensitivity to seasonal variation in evergreen ecosystems (A. Chen *et al.* 2022; Walther *et al.* 2016; Y. Zeng *et al.* 2022). Multivariate analysis of new, stronger proxies of photosynthesis, including near-infrared reflectance of vegetation (Badgley *et al.* 2017, 2019; Huang *et al.* 2019; Y. Zeng *et al.* 2022) and sun-induced chlorophyll fluorescence (X. Li and Xiao 2019; Piao *et al.* 2019; Sun *et al.* 2017; Tang *et al.* 2016; Y. Zeng *et al.* 2022), offers the potential for globally-consistent, data-rich, comparative insight into the seasonality of terrestrial ecosystems.

By providing a biological signal of the predominant seasonal environmental fluctuations that control the annual phenologies of a wide range of taxa, spatial patterns in LSP also convey valuable information for evolutionary biogeography. According to the Asynchrony of Seasons Hypothesis (ASH; (Martin *et al.* 2009), spatial variation in seasonal timing can decouple reproductive phenologies between populations, creating allochronic isolation (Hendry and Day 2005) that accelerates genetic divergence and perhaps even speciation (Coyne and Orr 2004; Taylor and Friesen 2017). The ASH posits this to be most common in the tropics and thus proposes it as a potential mechanism generating the latitudinal diversity gradient. This is because tropical species' phenologies are more often determined by seasonal timing of precipitation or solar radiation – which can decouple across short geographic distances because of complex interactions between topography and synoptic-to-mesoscale airflow, and can do so even between sites with similar climatologies or habitat conditions (Cavelier *et al.* 1996; Moore *et al.* 2005; Scholl *et al.* 2007; Thomé *et al.* 2021; Zhisheng *et al.* 2015) – rather than the common high-

latitude cues of temperature and photoperiod seasonality, which are generally synchronized across broad geographic areas. Allochronic isolation resulting from this phenomenon would likely be amplified by the tight altitudinal ranges and bioclimatic niches that characterize many tropical montane species (Freeman *et al.* 2022; Ghalambor *et al.* 2006; Janzen 1967). Thus, whereas chapter 2 provides an example of the spatial nature of an evolving system influencing its temporal dynamics, this chapter provides an example of the less-appreciated inverse situation: the potential for a system's temporal complexity to influence its evolutionary dynamics. Observational and genetic evidence for the ASH is mixed (Gamba and Muchhala 2020; Guarnizo *et al.* 2022; Moore *et al.* 2005; Quintero *et al.* 2014; Thomé *et al.* 2021), but global geographic patterns of climatic asynchrony are still poorly resolved (Guarnizo *et al.* 2022; Martin *et al.* 2009), global patterns and drivers of phenological asynchrony remain unknown, and a core implication of the ASH argument remains unassessed: that phenological asynchrony is less dependent on climatological differences in high-asynchrony tropical regions. Here, we leverage recent remote sensing advancements and high-performance computing to provide an unprecedented global analysis of the spatial variability and asynchrony of LSP. First, we use harmonic regression in Google Earth Engine (Gorelick *et al.* 2017) to fit pixelwise characteristic annual phenology curves to a rigorously quality-filtered, 10-year time series (2010-2019) of MODIS-derived (Vermote and Wolfe, 2015a, 2015b) NIR_v data. We use multivariate analysis of the results to compute and visualize a global map of characteristic LSP (henceforth, 'LSP map') at 0.05° (~5.5 km) resolution, then examine emergent spatial patterns in the context of previously described global and regional land cover patterns and climate-phenology relationships. Next, we produce and validate a global map of spatial phenological asynchrony (henceforth, 'LSP asynchrony map') to examine patterns of asynchrony at different spatial scales, explore variation in regional drivers of LSP asynchrony, and test a key implication of the ASH.

Results and Discussion

Phenology mapping

Our global LSP map reveals clear, coherent patterns from regional to intercontinental scales, demonstrating the broad ecological value of a globally-consistent, multivariate approach to LSP analysis. Empirical orthogonal function (EOF) analysis of the global set of characteristic annual LSP curves fitted by harmonic regression of NIR_v data suggests that the majority of the variation (95.6%) in earth's diverse LSP regimes is explained by three, orthogonal modes of spatiotemporal variation. The predominant mode (70.0% of total variation) largely reflects the north-south hemispheric seasonal dipole, but embedded within it is a clear signal of intercontinental convergence across major marine climate regions, including the five Mediterranean climate zones and their neighboring semi-arid shrublands, as well as the fringing wet-forest regions in the extreme east of Brazil and along the coast of Somalia, Kenya, and Tanzania (Fig. S1). Embedded within modes two (17.7%) and three (7.9%) appear to be the remaining seasonally-wet regions of the global tropical and subtropical monsoon systems (Zhisheng *et al.* 2015) and finer patterns related to agricultural production.

Combination of these three major modes into a color-composite visualization (Fig. 1) depicts the global diversity and patterns of LSP in unprecedented clarity. (To aid interpretation across latitudes we transformed all EOF modes that exhibit north-south dipoles using a two-part

weighted sum across the inter-tropical convergence zone (ITCZ), earth's functional climatic division between north and south hemispheres, but we provide untransformed color-composites of both components of the weighted sum in Fig. S2, for comparison.) At the broadest scales, intercontinental convergence stands out as a pattern of similar LSP gradients and mosaics occurring within similar geographic and climatic contexts. One striking example is the convergence between earth's more temperate Mediterranean-climate regions (California, coastal Chile, and the Mediterranean proper), where non-forest vegetation takes on predominantly green-blue hues representing phenological maxima in late winter and spring whereas forests, predominantly montane, display later-peaking orange-red hues (e.g., Fig. 1b) – corroborating and extending across continents the 'double peak' seasonality described by Turner *et al.* (2020). Within the remaining, warmer Mediterranean regions of the South African cape (Fig. 1e) and southern and southwestern Australia (Fig. 1g), non-forest in the milder coastal climate zones is similarly green and late-winter/spring-peaking, but isolated wet subtropical forests stand out in peach-orange and peak in summer, and vegetation in the hot subtropical inland climates quickly grades to winter-peaking dark blues, purples, and pinks roughly indicative of dry woodlands, shrublands and deserts, and grasslands, respectively. Another example of convergence occurs at the southeastern edges of tropical land masses, including northern Madagascar (Fig. 1f) and the Cape York Peninsula of Australia (Fig. 1h), where coastal wet forests (orange) with photosynthetic peaks not far from the summer solstice quickly grade westward into seasonally dry forests, shrublands, and savannas (pink and purple) with phenologies peaking late in the rainy season. Convergence also emerges between similar agricultural regions – for example, between the United States' corn belt (the rose-colored region that encompasses much of the upper midwest and extends along the Mississippi River in Fig. 1; (Zhong *et al.* 2016)) and the maize-rich regions to the north of the Po River in northern Italy (Fig. 1i), both of which contrast with peach-orange regional forests.

Complex patterns of regional LSP heterogeneity also emerge from Fig. 1, which we interpret as being attributable to a mixture of climatic, community compositional, and physiographic and ecophysiological mechanisms. In many regions, phenological heterogeneity appears to be driven by geographic variation in climate – for example, in Southern California, Northern Baja California, and western Sonora and Sinaloa (Fig. 1b), a sharp transition zone distinguishes the Mediterranean winter-monsoon climate zone to the north (lime green to blue) from the summer-monsoon climate that prevails in the south (pink to purple), and this stark division shows striking, incidental agreement with the monsoon climate zone whose orographic forcing was recently described (Boos and Pascale, 021). This climate-phenology relationship often appears partially convolved with a signal of plant community structure – for example, in North America's Great Basin (Fig. 1a), we observe a marked and significant advance of spring onset in areas more heavily invaded by cheatgrass and other annuals (lime-green pixels) compared to clusters of cells that peak one (olive-green) to two months (orange) later. The early-peaking clusters have significantly higher average estimated annual herbaceous cover (19.74% according to data from Maestas *et al.* (2020) than both the mid- (9.59%; Tukey's Honest Significant Difference (HSD) P-value: 0.00) and the late-peaking cluster (6.00%; HSD P-value: 0.00), providing a striking match to the cheatgrass, sagebrush, and montane phenologies fitted by Bradley *et al.* (Bradley *et al.* 2007) using a more sophisticated algorithm. Similarly, heterogeneity occurs in Northern Italy (Fig. 1i) between montane forest (peach-orange) and non-

forest vegetation (light green) as well as the Po River Valley's patchwork of agricultural activity (d'Andrimont *et al.* 2021), which is composed of mixed crops and pasture to the south of the River (also light green) and a string of maize-dominated regions to the north (rose). In other regions, however, phenological heterogeneity appears to be driven by a suite of local and largely non-climatic influences on plant ecophysiology. For example, stark patterns in South Florida (Fig. 1c) suggest a combined influence of microtopography and hydrology that makes sense in light of previous efforts to map vegetation (Homer *et al.* 2015) and CO₂ fluxes (C. Zhang *et al.* 2021): much of the Everglades sawgrass marsh (dark green-blue) exhibits peak photosynthesis during the winter dry season, when water levels are lowest, while the wooded wetland region to the north and west (light green) shows a quick spring peak that may partly reflect cypress leaf-out, and remaining areas of drained, urban, and upland vegetation, along with southwest mangrove forests (all orange-red), show a broader peak in the summer wet season. Similarly, microtopography or ecohydrology could underpin the distinction between the dual-peak phenologies observed in forests surrounding the mouth of the Amazon River (indigo and maroon; Fig. 1d) and the nearly unimodal phenologies (light blue) that occur in non-forest vegetation that is both naturally occurring (shrubland in the northwest) and anthropogenic (deforested areas at the mouth of the river and south of it), especially if different plant architectures and, thus, different root foraging dynamics cause different seasonal patterns of water availability for different plant communities subjected to an identical climate.

We validated the LSP map using both a shorter (2014-2017), independent SIF dataset from Orbiting Carbon Observatory 2 (OCO-2) (Yu *et al.* 2019) and a global set of gross primary productivity (GPP) time series from publicly-available flux tower datasets (FLUXNET; (Knoben 2019; Pastorello *et al.* 2020)). The NIR_v-SIF comparison shows that the NIR_v-derived LSP map closely matches its SIF-derived counterpart (R^2 between NIR_v and SIF fitted seasonal time series has a pixel-wise median of 0.872), but reveals low between-dataset correlation in some drier grassland, shrubland, and woodland biomes (Fig. 2a). We feel that any disagreement between these two datasets is unlikely to be an artefact of the ANN-derived SIF map we chose to use because the ANN-interpolated seasonality within OCO-2 orbital gaps matches that of an independent TROPOMI SIF dataset (Fig. S3, S4). Subsequent validation of both the NIR_v and SIF LSP datasets against flux-tower gross primary productivity (GPP) data is strong in most biomes but is again mixed in roughly the same subset of drier biome types (Fig. 2b, 2d). We believe a large fraction of this biome-sensitivity results from mismatched temporal scales, given that seasonally dry biomes are generally characterized by large interannual variability in productivity (Broich *et al.* 2014; Buitenwerf *et al.* 2015; Poulter *et al.* 2014), leading to frequent phenological disagreement between the time periods covered by the GPP time series available at flux towers, which vary both in dates and length, and the fixed period covered by our remote sensing datasets. The fact that the bulk of the poorly validating datasets have shorter time series than our 10-year MODIS NIR_v dataset and even our 4½-year interpolated OCO-2 SIF dataset (Fig. 2c, 2e) supports this interpretation. However, we also suspect that some of the observed biome-sensitivity derives from mismatched spatial scales, given that many of the biomes in question are characterized by high spatial heterogeneity in vegetation structure and thus in phenology (Richardson *et al.* 2013), often at spatial scales finer than our 0.05° pixels. It is thus possible that local asynchrony in LSP (Turner *et al.* 2020) causes a mismatch between the local GPP dynamics recorded at flux towers and the NIR_v dynamics of the majority land cover that is

captured by our coarser pixels (L. Zeng *et al.* 2020). This interpretation is supported by the sensitivity of tower sites' validation performance to the distance between them and their nearest available NIR_v pixels. Both temporal and spatial mismatches could also help explain the heightened divergence between the NIR_v and SIF datasets in these regions, to the extent that the distinct ranges of time and the non-coregistered footprints of the raw sensor data would cause differential mixes of vegetation types to propagate through to each of our datasets. Nonetheless, it is also plausible that a portion of the biome-sensitivity we observe could result from a weaker ability of our remote-sensing metrics to reflect productivity (as measured in GPP) in some regions; both NIR_v and SIF show strong correlation with canopy productivity at seasonal-to-annual scales (Pickering *et al.* 2022; J. Zhang *et al.* 2022), but a large portion of the signal in these metrics is structural (e.g., leaf clumping, leaf-angle distribution; Dechant *et al.* 2020), and these metrics' shortcomings are being actively researched, particularly with respect to the estimation of GPP (Kim *et al.* 2021; Yang *et al.* 2022; M. Zhang *et al.* 2022). To gain both ecophysiological and remote-sensing insight, future work should aim to discriminate between cases where the divergence of our NIR_v and SIF phenologies from GPP reflects actually divergent photosynthetic dynamics – for example, differences in the temporal dynamics of water-availability experienced under different plant architectures or soil conditions (Phillips *et al.* 2016) or in the environmental stress induced by different near-surface climates (Bonan 2008; Cohn *et al.* 2019; Sampaio *et al.* 2007) – and cases where it instead reflects structural decoupling of our remote sensing metrics from the true seasonality of canopy productivity. Candidate regions for such work could include places where our map registers starkly different phenologies between closely juxtaposed vegetation types whose occurrence is anthropogenic rather than being driven by topoclimate, such that *in situ* work could determine if the differential remote sensing signal across different canopies exposed to the same climate is being driven by actual differences in seasonal productivity or just by differences in seasonal canopy structure.

Phenological asynchrony

We estimate spatial asynchrony as the slope of the relationship between a.) the Euclidean distances between a pixel's fitted characteristic annual time series and the fitted time series of all of its neighbors, and b.) the geographic distances between a pixel and all its neighbors (explained in detail in Fig. 3a). (Note that the resulting asynchrony metric is technically expressed in change in units of the target variable per change in geographic distance, but the units are not especially informative and the magnitude of the mapped values is a function of the neighborhood radius used for the calculation, such that values be considered arbitrary and only informative when considered as relative values within a map.) Global mapping of LSP asynchrony using this metric (Fig. 3b) reveals high-asynchrony hotspots centered in temperate Mediterranean climates and neighboring xeric regions, in tropical montane regions, and in the 'arc of deforestation' and other parts of the Amazon basin. All NIR_v-derived maps, irrespective of the neighborhood radii used to derive them, validate well against independent, SIF-derived maps ($R^2=0.43$ for asynchrony calculated at a 50 km radius, 0.49 at 100 km, and 0.52 at 150 km; see Fig. S5). We display all 6 phenological asynchrony maps in Figs. S6 and S7, but in our subsequent discussion focus on the 100 km-radius NIR_v LSP map and all 100 km-radius climatic asynchrony maps, both because the daily-composite, uninterpolated, 10-year NIR_v dataset provides more robust information than the SIF dataset, and because we find that global patterns of asynchrony are

largely insensitive to the choice of neighborhood radius size (R^2 between NIR_v maps at different radii are 0.815 for the 50-150 km comparison, 0.899 for 50-100 km, and 0.953 for 100-150 km; Table S1 provides these R^2 s for all radius comparisons and all phenological and climatic asynchrony maps produced).

To explore potential drivers of global LSP asynchrony patterns, we constructed a random forest model predicting LSP asynchrony as a function of seven covariates (detailed in Table S2): the spatial asynchrony of five major climatological variables (mean minimum and maximum temperature, mean precipitation, mean climatic water deficit, and mean cloud cover; in Figs. S8-12, we show these maps for all three neighborhood radii), neighborhood entropy of vegetation structure, and topographic complexity (as measured by the vector ruggedness metric). Results show a strong ability to predict the broad patterns of phenological asynchrony, irrespective of model structure ($R^2=0.81$ and $RMSE=0.44$, for our ‘main model’, predicting the 100 km-neighborhood NIR_v LSP asynchrony map and including geographic coordinates as covariates; we ran models for all three neighborhood radii, for both phenology datasets, and both including and excluding geographic coordinates as covariates, and chose this as our main model after finding that our results are largely insensitive to phenology dataset, asynchrony neighborhood radius, and geographic coordinate inclusion; see Table S3). Our results attribute global patterns of LSP asynchrony to three predominant covariates: precipitation asynchrony; minimum temperature asynchrony; and neighborhood entropy in vegetation structure (Fig. 4a-b). To regionally interpret our global model’s results, we map the local influence of these three predominant covariates using SHAP values (Lundberg and Lee 2017) (Fig. 4c). We then combine these three SHAP maps into the summary map in Fig. 4d, in which the hue of a pixel (yellow, cyan, magenta) indicates its predominant covariate, the value (i.e., brightness) indicates the scale of predominance, and the saturation (i.e., color richness) is scaled to LSP asynchrony, such that low-asynchrony locations fade to white.

The summary map in Fig. 4d, especially when considered in the light of our phenology visualization (Fig. 1), suggests a few generalized explanations for the observed global pattern of LSP asynchrony. First, asynchrony is low in: moist temperate regions (e.g., eastern North America; eastern Asia; Fig. 3b), where consistent water availability but long, harsh winters synchronize seasonal LSP across space and across vegetation types; and in most tropical regions without substantial topographic variation, where consistent warmth and solar energy allow broad regional patterns of precipitation seasonality to similarly synchronize seasonal LSP across space and across vegetation types. Second, asynchrony is generally high in global Mediterranean climate zones, where distinctly structured vegetation communities show decoupled phenologies (Turner *et al.* 2020), perhaps because of extrinsic topoclimatic controls (e.g., altitudinal variation in growing season) and intrinsic physiological controls (e.g., differential plant structure and root architecture could lead to different seasonal patterns of water availability and thus to different resource-use strategies). A similar mechanism could be operating in the western Amazon arc of deforestation, but further research is needed to understand the nature of the dramatically different LSP between forest and juxtaposed non-forest in this anomalous region. Third, precipitation asynchrony is a major cause of LSP asynchrony across a number of regions globally, including not only the semi-arid mid-latitude regions that neighbor Mediterranean climates (e.g., the Mediterranean-monsoon transition zone of North America’s desert southwest and the mountains of northwest Africa, the Middle East, Central Asia, and eastern Australia) but also some montane

tropical regions characterized by strong precipitation gradients (e.g., eastern Brazil, the Afromontane, eastern Madagascar, northeast Queensland). Finally, asynchrony is generally high across most other tropical montane regions, where minimum temperature asynchrony is sometimes the top covariate, but where covariate predominance is low in many locations (black regions of the tropics in Fig. 4d). The apparent importance of minimum temperature is a surprising result, given that tropical seasonality is most commonly considered in the context of precipitation, but we also find that the high temperature asynchrony we map in the tropics recapitulates the findings of Martin *et al.* (2009), albeit at higher resolution, and it is entirely plausible that minimum temperature is a previously unrecognized control on phenology in these regions, at least at some elevations or in some biomes. However, we also note that low covariate predominance could indicate that complex driver interactions underpin LSP in these regions and that more nuanced climatic LSP drivers could be acting in these regions but could be correlated with minimum temperature asynchrony and omitted from our model. Finally, it is possible that many species phenologies are unrelated to climate in these regions, such that the species dominance of a pixel's signal determines that pixel's fitted phenology and, when aggregated across space, generates a signal of high asynchrony; however, we do not believe this to be the case, not only because we would expect it to create 'static' in these regions in the Fig. 1 RGB map, which we do not see.

The ASH suggests that tropical montane asynchrony could be driven partially by differences in phenology between different climatic zones, as would be expected in temperate zones, but could also be decoupled from climate to the extent that climatically similar sites can have divergent seasonalities. To probe this possibility, we test the hypothesis that the strength of the relationship between the pairwise climatic and phenological distances between sites is positively correlated with absolute latitude. We first used a clustering algorithm and a hull-delineation algorithm (together referred to as our 'regionalization algorithm') to empirically describe a global set of high-asynchrony regions (Fig. 5a, S8a). We drew sets of random points within those regions and ran a matrix regression of the form $dist_{phen} \sim dist_{clim} + dist_{geog}$ on each point set, then tested the hypothesis that $dist_{clim}$ coefficients from regional models are positively correlated with regions' mean absolute latitudes. Because of the potential sensitivity of our results to the parameterization of our regionalization algorithm, we ran this analysis for all 27 combinations of reasonable ranges of values for each of the three regionalization hyperparameters parameters (the maximum geographic distance between two points that the DBSCAN clustering algorithm could consider to be in the same neighborhood, *epsilon*: 2, 3.5, 5; the number of samples required within a neighborhood for a point to be considered as a cluster's core, *min_samples*: 0.3, 0.45, 0.6; and the value controlling the maximum complexity of a region's alpha-complex polygon, *alpha*: 0.25, 0.75, 1.25). Summarizing all parameterizations, we found that the magnitude of the relationship between climatic and seasonal distances varies regionally (Fig. 5a) but that the overall pattern provides a clear result: phenological distance between sites is less related to climatic distance at lower latitudes (Fig. 5b-c; Monte Carlo P-value = 0.00). This supports a core implication of the ASH (Martin *et al.* 2009), suggesting that tropical species – already often more range-restricted by topoclimate (Ghalambor *et al.* 2006; Janzen 1967; Sheldon *et al.* 2018) and occurring at lower population densities more susceptible to genetic drift (terSteege *et al.* 2013) – are likely to experience additive, allochronic isolation

because of phenological asynchrony between sites that are climatologically equivalent but phenologically out of sync.

Indeed, our global asynchrony map identifies high-asynchrony regions in many of earth's continental biodiversity hotspots (Myers *et al.* 2000), including not only tropical montane regions (e.g., tropical Andes, Brazil's Atlantic Forest, and the Afromontane region) but also Mediterranean and semiarid regions (e.g., California Floristic Zone, Cape Floristic Province, Chile, Western Australian, and the Mediterranean and Caucasus). This striking pattern is consistent with the possibility that seasonal asynchrony has played a role in the generation of broad biogeographic patterns of species diversity, as predicted by the ASH. Our maps can help to identify sites where future phenological research can quantify the predictive power of LSP for species-level reproductive cycles and where population genetic data can measure the extent of allochronic isolation between phenologically asynchronous populations. Compelling support for this use comes from the fact that the local seasonal phenologies reported in our dataset show strong agreement with the local climatic seasonalities and seasonal breeding phenologies reported in the only two field-based population genetic ASH studies we are aware of (Moore *et al.* 2005; Thomé *et al.* 2021) and that precipitation is both the predominant driver of climatic asynchrony identified in those studies and an important covariate in our analysis within these regions (Fig. 4d).

Conclusion

Our LSP analysis provides not only unprecedented global understanding of variation in canopy phenology at an ecosystem scale but also a globally-consistent demonstration of the utility of long-term remote sensing archives for description of characteristic LSP, even in semi-arid and arid biomes typified by high interannual variability in carbon cycling and low signal:noise ratio of remotely-sensed photosynthesis. We identify coherent spatial patterns of intercontinental LSP convergence and regional divergence, and contextualize those patterns within a broader understanding of spatial bioclimatic variation. Finally, our work provides not only the first moderate-resolution global maps of climatic and LSP asynchrony but also new insight into the drivers of LSP asynchrony and global patterns in its independence from climatic variability, lending strong support to a core prediction implication of the ASH. We believe our findings provide a strong foundation for future work in many promising and largely uncharted areas of research, including development of higher-resolution, regional LSP and LSP asynchrony maps, higher-resolution and regionally-focused analyses of the drivers and climatological dependence of LSP asynchrony, and evolutionary biogeography research that can more explicitly consider and account for the potential effects of phenological asynchrony on allochronic genetic isolation between conspecific populations.

Methods

Land surface phenology (LSP) datasets

Both the near-infrared reflectance of vegetation (NIR_v) and sun-induced chlorophyll fluorescence (SIF) datasets were preprocessed on Google Earth Engine (GEE; Gorelick *et al.* 2017). We use a 10-year archive of MODIS-derived NIR_v data as our main dataset. Following best-practices for estimation of patterns at the annual timescale, we chose the MCD43A4 v006 dataset (Schaaf and Wang 2015), a nadir bidirectional reflectance distribution function (BRDF) adjusted reflectance (Y. Zeng *et al.* 2022) 16-day temporal composite (L. Zeng *et al.* 2020). We used the version of this data that is publicly available in the GEE data catalog. We did not carry out topographic correction because the scale of our analysis (0.05°; ~5.5 km) is sufficiently coarse that spatial averaging is expected to remove topographic bias (R. Chen *et al.* 2022; Y. Zeng *et al.* 2022).

To help validate our NIR_v maps, we ran some of our main analyses identically but using the global, gridded SIF dataset produced by Yu *et al.* (2019). This is a spatially contiguous time series dataset, interpolated by artificial neural net (ANN) from the spatially discontinuous SIF data measured along Orbiting Carbon Observatory 2 (OCO-2) orbital swaths. This dataset was rigorously validated, internally and externally, by its authors, who found that it accurately captured the global patterns present in the original OCO-2 retrievals and that it explained 81% of the variation in contemporaneous Chlorophyll Fluorescence Imaging Spectrometer (CFIS) aerial measurements taken beneath OCO-2 orbits and 72% of the variation in measurements not beneath orbits. We downloaded this data from the Distributed Active Archive Center for Biogeochemical Dynamics (https://daac.ornl.gov/VEGETATION/guides/Global_High_Res_SIF_OCO2.html), then ingested it into GEE.

Given that the SIF dataset interpolates across orbital gaps but that the paper describing the dataset did not explicitly validate the seasonal phenological patterns of the interpolated data, we compared seasonality in the interpolated, orbital-gap data to seasonality in another, coarser-resolution SIF dataset collected by the Tropospheric Monitoring Instrument (TROPOMI). To do this, we extracted SIF time series from the ANN-interpolated dataset at a sample of random points drawn within OCO-2 orbital gaps in three tropical realms (the Neotropics, tropical Africa, and Indo-Pacific and Australia; Fig. S1), then compared those values to contemporaneous time-series extracted from another high-resolution, satellite SIF dataset derived from data from the TROPOMI sensor (Köhler *et al.* 2018; Köhler and Frankenberg, 2020). We used tropical regions for this validation because their lack of a pronounced thermal winter creates the possibility that seasonality there exhibits spatially varying patterns that are not accurately recovered by spatial interpolation from orbital swaths. If the interpolated dataset adequately captures the true seasonal patterns of SIF within OCO-2 orbital gaps then its time series should explain the bulk of the variation in the TROPOMI time series. Indeed, we found very high correlation between these orbital-gap datasets ($R^2 = 0.87$; Fig S2).

Data filtering

In an attempt to exclude locations where our LSP-mapping methodology would return inaccurate or artificial results, we produced both LSP maps (NIR_v and SIF) using an extensive filtering pipeline that removed invalid land cover, data-deficient pixels, and pixels with

statistically insignificant fitted LSP. For land-cover filtering, we used the GEE data catalog asset for the 500-meter resolution, yearly, global MODIS land cover dataset (MCD12Q1.006; (Friedl and Sulla-Menashe 2019), classified according to the Annual International Geosphere-Biosphere Programme's (IGBP) classification scheme ('Land Cover Type 1'). To avoid fitting LSP noise resulting from land cover/land use change, we removed any pixels at our analysis resolution of 0.05° , or ~ 5.5 km (henceforth, 'target resolution') within which the most common (i.e., mode) land cover type of the 500-meter MODIS pixels was not the same in at least 8 of the 10 years of the NIR_v time series. Next, to avoid fitting LSP curves to data with no expected productivity, we iterated over each year in an LSP dataset and dropped data for any target-resolution pixels where $> 20\%$ of the 500-meter pixels for that year were classed as permanent snow and ice, barren, or water bodies.

Next, because we observed a tendency for the harmonic regression procedure (described below) to interpolate spurious second LSP peaks during extended periods of missing data (e.g., during high-latitude winters), we removed any target-resolution pixels whose data did not satisfy a set of strict non-missingness criteria. First, we dropped any pixels whose LSP time series had $> 50\%$ missing data, a simple step to help remove sites with data dropout because of substantial cloud contamination or regular MODIS QC problems. Next, we removed any pixels with a monthly mean proportion of non-missing daily NIR_v data < 0.1 . Finally, we removed any pixels whose Pielou's evenness (Pielou 1966) was < 0.8 ; we calculated Pielou's evenness, $J' = H'/H'max$, by calculating H' (i.e., Shannon's diversity index; (Shannon 1948) using 12 values, each value being a monthly average proportion of non-missing daily data over the 10-year NIR_v archive. Manual inspection of fitted LSP patterns after applying this series of filtering steps confirmed successful removal of locations previously producing spurious results.

Finally, after running the global harmonic regression described in the subsequent section, we used permutation tests to filter out pixels whose fitted seasonality was insignificant. To do this, we ran a number of additional harmonic regressions, each time permuting the LSP time-series stack so as to scramble any true temporal (i.e. seasonal) pattern. The image of R^2 values was extracted from each permuted regression and compared to the R^2 values from the true harmonic regression, and the proportion of permutations for which the permuted time series' R^2 exceeded the unpermuted R^2 was stored, providing a map of empirical P-values for the unpermuted fitted LSP patterns. For the phenology portion of the study, we ran this test using 25 permutations at every pixel globally (because of computational limitations), then filtered out any pixels with a P-value \geq a one-sided significance threshold of $\alpha=0.01$.

We used the series of filtering steps described above to produce both the NIR_v and SIF datasets for which we then carried out the LSP mapping and validation analyses described below. This allowed us to retain in our LSP mapping results both agricultural land cover and more contiguous coverage of semi-arid regions, improving the utility and visibility of the resulting LSP summaries. However, to maximally avoid sensitivity of our LSP asynchrony analyses to anthropogenic and spurious results we made the conservative choice to implement two additional data-filtering steps prior to producing LSP asynchrony maps. First, because we found that the permutation-testing procedure yielded more conservative results when run on the shorter and less temporally resolved SIF dataset – i.e., dropped many more pixels in semi-arid regions, cloudy tropical mountains, and other regions of marginal data quality – we used 100 permutations of the SIF data and the same one-sided significance threshold ($\alpha=0.01$) to produce a second

significance mask, then removed from both the NIR_v and SIF datasets any pixels flagged as insignificant in that mask. Second, because our asynchrony analysis explores the potential climatic drivers of LSP we chose to remove any target-resolution pixels whose land cover suggested that their phenologies could be influenced by irrigation or other management practices, and thus could be at least partly decoupled from climatic control. To do this, we removed from both the NIR_v and SIF LSP datasets used to calculate asynchrony maps any years' data at target-resolution pixels where > 20% of the 500-meter pixels in that year's MODIS land cover dataset were classed as croplands, urban and built-up lands, or cropland/natural vegetation mosaics.

Estimation and validation of seasonal phenology

To estimate the seasonality of stand-level photosynthesis we ran a harmonic regression model, vectorized across all 0.05° pixels in the global NIR_v and SIF maps, in which each pixel's time series is predicted as a function of time. For each variable we used the following model:

$$val = \beta_0 + \beta_1 \sin(t_{ann}) + \beta_2 \cos(t_{ann}) + \beta_3 \sin(t_{sem}) + \beta_4 \cos(t_{sem}) + \varepsilon$$

where t is the linear time component (days from the start of the time series), and t_{ann} and t_{sem} are circular time expressed on annual and semiannual frequencies (i.e., day of the year expressed in radians, where 2π radians corresponds to the last day of the year for t_{ann} , or to middle and last days of the year for t_{sem}). We then retained all resulting coefficient maps except β_0 , yielding a stack of five coefficient maps that describes the detrended, characteristic annual LSP fitted at each pixel globally. (Note that this operation is algebraically equivalent to detrending the time series and then running a Fourier transform that includes both the annual and semiannual frequency components.) We chose to include both the annual and semiannual frequencies in the harmonic regression in order to strike a balance between complexity and overfitting. We expected that complex annual LSP patterns were likely to occur in locations around the globe that have bimodal annual seasonal patterns of precipitation (Knoben 2019) and deep winter freeze; indeed, preliminary analysis revealed numerous regions with stronger bimodal than unimodal annual LSP patterns (i.e. regions containing many pixels whose R² values were higher in semiannual-only harmonic regression models than in annual-only models) in precisely those regions. The chosen linear combination of annual and semiannual components is complex enough to allow us to represent annual LSP curves that are unimodal, evenly bimodal (two equal peaks and troughs), or unevenly bimodal (featuring major and minor peaks and troughs), but not more complex, and thus avoids overfitting by excluding unfounded higher frequencies.

We first validated the annual NIR_v LSP map by calculating a map of R² values between each pixel's fitted annual NIR_v and SIF LSP curves. To do this for a given pixel, we first predict the pixel's characteristic annual NIR_v LSP as a daily time series by matrix-multiplying the pixel's vector of five retained NIR_v coefficients with a 365x5 matrix in which rows represent days of the year and columns contain those days expressed in the units of the regression model, explained above. We then do the same operation for the corresponding pixel in the SIF map. Finally, we calculate the R² of the two resulting 365-day, detrended, fitted time series. (We

henceforth refer to the characteristic LSP time series thus obtained as “fitted LSP curves”.) Doing this for all non-missing pixels globally produces the validation map displayed in Fig. 2a. We next separately validated the fitted LSP curves for both LSP datasets (NIR_v, SIF) against gross primary productivity (GPP) curves that we fitted to FLUXNET2015 (Knoben 2019; Pastorello *et al.* 2020) daily GPP data following an identical harmonic regression methodology as above. We used data from any tower: 1.) located within a valid, unmasked pixel in our LSP datasets or ≤ 2 pixels (queen’s neighborhood) distant from a valid, non-masked pixel; and 2.) providing at least one full year of publicly-available GPP data at the time of download (October 11, 2021), resulting in 237 GPP datasets for validation of the NIR_v map and 234 datasets for validation of the SIF map. We first manually downloaded all available such datasets from the FLUXNET server. We then used Python to iteratively load each tower’s dataset and implement the harmonic regression model as long as at least one year of contiguous GPP data was available. We used the fitted coefficients from each such regression and the matrix multiplication previously described to generate a 365-day fitted characteristic GPP curve for each flux tower. Finally, for each tower, we calculated the R^2 s between the tower’s fitted GPP curve and each of the fitted LSP curves corresponding to the tower (i.e., either pertaining to the pixel where the tower is located or, if that pixel is masked in our LSP dataset, the nearest valid, non-masked pixel within a 2 pixel-wide box that surrounds it). We summarize this validation procedure across all tower datasets by producing, for each LSP dataset: 1.) a scatter plot of the LSP-GPP R^2 values overplotted on the Whittaker biomes (Whittaker 1970) (Fig. 2B,D), to depict bioclimatic patterns in validation performance; and 2.) a scatterplot of LSP-GPP R^2 versus GPP time series length (Fig. 2C,E), to depict the relationship between GPP data availability and validation performance.

Visualization of seasonal phenology

To visualize the global variability of seasonal phenology present in the results of our harmonic regression, we used three-channel color visualization of the results of a dimensionality-reduction analysis to produce a single global map. First, we used Python to run empirical orthogonal functions (EOF) analysis (i.e., principal components analysis (PCA) of spatiotemporal data) to reduce the dimensionality of our global set of fitted LSP curves (a data cube of dimensionality $i \times j \times 365$, where i and j are the latitudinal and longitudinal dimensions of our map and 365 is the length of the fitted characteristic daily time series). We ran the EOF analysis using the square root of the cosine of the latitude as area weights for pixel data, per standard practice.

Finding that the first three EOFs cumulatively explain 95.6% of the total variation in our global, we used the red, green, and blue (RGB) color channels to display them, thus visualizing the bulk of global canopy phenological variability within a single map. Given that EOFs 1 and 2 have embedded within them both the unremarkable north-south hemispheric seasonality dipole and the hemisphere-independent patterns of interest (e.g., monsoon-driven dynamics), we devised a method to transform the raw EOF maps prior to RGB visualization, and hence to represent global phenological variability in a consistent color scheme irrespective of the north-south hemispheric divide. To do this, using WebPlotDigitizer (Rohatgi 2022) we first manually digitized a geospatial vector file of the mean intertropical convergence zone (ITCZ) in boreal summer (June, July, August) and winter (December, January, February) as presented by Zhisheng

et al. (2015). Then we used that file to calculate a single, annual-mean ITCZ geospatial vector by averaging the boreal summer and winter ITCZ latitudes at evenly spaced longitudes across the globe. Finally, for each of EOFs 1 and 2, we constructed a synthetic, transformed map by calculating $w \times EOF + (1-w) \times (1-EOF)$, where w varies from 1 in the northern hemisphere to 0 in the southern hemisphere and transitions linearly from 1 to 0 in a 10 degree band latitudinally bracketing the annual mean ITCZ. We chose to use the ITCZ as the latitudinal boundary across which to transform the EOF maps for visualization because it serves as a much more natural climatic dividing line between the northern and southern hemispheres than does the equator. We heuristically tuned this transformation method so as to minimize the inevitable artefactual color-warping it creates; nonetheless, to help interpret the main result (displayed in Fig. 1) across the region surrounding the ITCZ, where color-warping occurs, we provide in Fig. S2 a pair of equivalent RGB maps created using untransformed EOF maps and using EOF maps transformed as $(1-EOF)$. Despite the unavoidable color-warping this visualization method creates, we still feel that this single-map multivariate summary provides a valuable global comparison of LSP patterns, and we emphasize that because this method is used only for visualization it thus has no influence on any analytical results of this study.

To highlight focal regions with patterns of interest, we manually defined the bounding boxes of each of the focal-region maps within Fig. 1 (a-i), then plotted each focal region's map using the same RGB visualization displayed in the global map. To depict the characteristic annual phenological time series corresponding to each of the major colors in a region, we first used mini-batch K-means clustering (a version of the standard K-means clustering algorithm that reduces computational burden by using only a fixed-size random subsample of the full dataset at each iteration) to cluster the focal region's RGB map into K colors, for $K=1:10$, then visually inspected the focal region's scree plot (% variance explained as a function of K) to visually determine the K value nearest to the elbow, and hence the most parsimonious value of K for clustering phenologies within the region. We then used K-means clustering for the chosen K value to assign each pixel to one of K clusters. Finally, for each cluster, we calculated a summary time series as the composite of the daily medians of all of the cluster's pixels' original time series. We used these results to provide a summary comparison of the range of fitted LSP curves represented by the RGB map within a focal region: to do this, we plotted each cluster's summary time series as a line plot, colored by the RGB value corresponding to the cluster's center. Noting regional phenological variability in the Great Basin of the United States (Fig. 2B) that appeared to match the cheatgrass-invaded, sagebrush, and montane phenologies presented by Bradley *et al.* (2007), we used ancillary data from Maestas *et al.* (2020), aggregated to the 0.05° resolution of our map, to calculate the average estimated percent annual herbaceous cover in each of the three predominant clusters fitted by our analysis. To support our interpretation of the three clusters as annual-invaded regions, sagebrush, and montane vegetation, which we based on the differences in their average annual herbaceous cover and visual comparison to Bradley *et al.* (2007), we used ANOVA to test the significance of the differences across all three clusters, followed by a Tukey Honest Significant Difference (HSD) test to test for significant pairwise differences between the clusters.

Calculation of phenological asynchrony

Although GEE is designed for carrying out a set of common spatial analysis operations at scale, it is not well suited to the calculation of custom neighborhood metrics such as our spatial asynchrony metric. Thus we exported the results of our filtered harmonic regression as a series of *TFRecord* (TensorFlow Record) files, each containing a stack of images of the regression coefficients for a number of rectangular geographic regions. Earth Engine accepts a ‘kernel width’ argument for the export of these files, which determines the extent of overlap between neighboring regions. We set the kernel width to be double our largest target neighborhood size (300 km = 2×150 km), thus outputting a series of files whose new asynchrony images could be independently calculated. These new asynchrony files were processed using parallel Julia (Bezanson *et al.* 2017) code on a supercomputer (UC Berkeley’s Savio cluster, using the savio3 partition, each node of which contains a 2.1 GHz Skylake processor with 32 cores and 96 GB of RAM).

We calculated a global image of our asynchrony metric as a series of regional tiles, and each tile was calculated pixel-wise, using an algorithm founded upon that described in Martin *et al.* (2009):

1. Use the regression coefficients to calculate the 365-day annual time series of fitted SIF or NIR_v values for the focal pixel (following the matrix-multiplication operation explained above);
2. Identify all pixels whose centerpoints are within the chosen neighborhood radius of the focal pixel (the ‘neighbor pixels’);
3. For each neighbor pixel:
 1. Calculate the fitted SIF or NIR_v time series (again using the same matrix multiplication);
 2. Calculate and save the 365-dimensional Euclidean phenological distance between that time series and the focal pixel’s time series (after standardizing both);
 3. Calculate and save the geographic (geodesic) distance to the focal pixel;
4. Calculate asynchrony as the slope of the overall regression of neighbor-wise Euclidean phenological distances on neighbor-wise geographic distances;
5. Store the asynchrony value (as well as the R² and sample size of the regression used to devise it).

(This algorithm is visualized in Fig. 3a, and the methods used to visualize it are described in the subsequent section.) Finally, we used Python to mosaic the resulting collection of asynchrony image TFRecord tiles into a global map (displayed in Fig. 3b).

Asynchrony conceptual diagram

To depict the calculation of asynchrony, we first simulated harmonic-regression output for a low-asynchrony region (as a 5-layer stack of mean coefficient values with rasters of low-relative-magnitude Gaussian noise added to them) and for a high-asynchrony region (as a 5-layer stack of mean coefficient values with large-relative-magnitude spatially autocorrelated noise added to them, using neutral landscape models (NLMs) generated by the *nlmpy* package (Etherington *et al.* 2015)). We reduced each 5-layer simulated map to a single-layer map by calculating each pixel’s 365-day fitted phenological time series from the pixel’s simulated vector of 5 harmonic regression coefficients, then calculating the day of the year when the simulated variable reaches its peak. We display the resulting day-of-peak maps, colored using a cyclical

colormap, in the left column of Fig. 3a, showing a low-asynchrony region in the upper row and a high-asynchrony region in the lower row, and we superimpose on each a translucent circular neighborhood within which the center pixel's asynchrony metric will be calculated. We display line plots of all of an asynchrony neighborhood's pixels' fitted time series in the center column. We then display scatter plots of an asynchrony neighborhood's pixels' geographic distance-phenological distance relationships in the right column, along with the zero-intercept simple linear regressions fitted to those scatter plots, the slopes of which serve as the center pixel's asynchrony metric (as explained in the previous section). We use orange and red stars (left and right columns) and orange and red curves (center column) to track the behavior of a pair of focal neighbor pixels at nearer and further geographic distance from the focal pixel (i.e., the black star at the centers of the maps in the left column).

Calculation of climatic and physiographic covariates

For our random forest model exploring the potential drivers of phenological asynchrony (explained in the next section) we used workflows combining GEE, Julia, GDAL, and Python to produce a number of physiographic and environmental covariates. We saved all covariates as GeoTIFF rasters, aggregated to the resolution of our response variable (NIR_v , 0.05°), for downstream statistical analysis in R (R Core Team 2021). (For a detailed synopsis of the covariates, see Table S2.)

First, we used the 64-year archive of TerraClimate data (Abatzoglou *et al.* 2018) to develop a set of maps describing the asynchrony of climatic factors potentially driving phenological asynchrony: minimum and maximum monthly temperature, monthly precipitation, climatic water deficit. We supplemented this with a map of asynchrony in cloud cover by extracting the internal cloud algorithm flag bit (bit 10 of the 1 km reflectance data QA band) from the MODIS Aqua and Terra daily 1 km global surface reflectance archive (MYD09GA.006 (Vermote and Wolfe 2015b) and MOD09GA.006 (Vermote and Wolfe 2015a)). We accessed both the TerraClimate and MODIS datasets using the GEE data catalog, reprojected them to match our target resolution (0.05°), processed them using the same harmonic regression workflow described above for the LSP datasets (minus the LSP-specific data filtering pipeline), then downloaded the set of overlapping TFRecord image tiles and processed them on UC Berkeley's Savio cluster using the same parallelized Julia code as used for the LSP asynchrony maps. This resulted in a series of three asynchrony maps for each climatic variable, one for each of our three neighborhood radii (50 km, 100 km, 150 km).

To allow the random forest to capture phenological asynchrony between juxtaposed and structurally distinct vegetation communities that is not attributable to climatic asynchrony, we used GEE to create a global map of entropy in vegetation structure within 100 km neighborhoods (henceforth, the 'vegetation entropy map'). To do this, we used a 10-year series (matching the 10 years of the NIR_v dataset) of reclassified annual MODIS IGBP 500 m land cover (Friedl and Sulla-Menashe 2019; the same dataset used in the LSP data-filtering workflow described above). We reclassified land cover into categories of forest (IGBP classes 1-5: evergreen or deciduous broadleaf or needleleaf forests and mixed forest), shrubland (IGBP classes 6 and 7: closed and open shrublands), savanna (IGBP classes 8 and 9: woody savannas and savannas), grassland (IGBP class 10), or permanent wetland (IGBP class 11); we masked out IGBP classes 12 and above (croplands and cropland mosaics, urban and built-up lands, permanent snow and ice,

barren, and water bodies) so that the information captured by this covariate would match the information reflected in the LSP asynchrony response variable from which the same categories were removed. Next, we reduced the stack of 10 years' reclassified vegetation structure maps into a single map representing the most common class (i.e., mode), then for purposes of computational tractability we aggregated that map to 0.05° using the mode. Finally, we calculated the entropy of the vegetation structure classes within each pixel's circular (100 km radius) neighborhood, producing the vegetation entropy covariate map.

To capture the potential importance of topographic complexity for driving phenological asynchrony, we downloaded a global map of the vector ruggedness metric (Sappington *et al.* 2007) as a final covariate. We chose this over other measures of topographic complexity because of its reduced correlation with slope. We downloaded the layer produced by Amatulli *et al.* (2018) from <http://www.earthenv.org/topography>, choosing the SRTM-derived, median-aggregated layer at 50 km so that each 50 km pixel could serve as a proxy for a true neighborhood metric, providing information about the area surrounding our ~5.5 km (0.05°) target-resolution LSP response variables.

Drivers of phenological asynchrony

To explore the potential drivers of LSP asynchrony, we constructed a random forest model (using R's *ranger* package; Wright and Ziegler 2017) predicting LSP asynchrony using the set of neighborhood-metric covariates described above (and detailed in Table S2): $LSP.asy_{neigh} \sim ppt.asy_{neigh} + tmp.min.asy_{neigh} + tmp.max.asy_{neigh} + def.asy_{neigh} + cld.asy_{neigh} + veg.ent + vrm.med [+ x + y]$, where *LSP.asy* is asynchrony of LSP datasets (using either the NIR_v or SIF dataset), *neigh* indicates the asynchrony neighborhood radius, *ppt.asy* is precipitation asynchrony, *tmp.min.asy* and *tmp.max.asy* are minimum and maximum temperature asynchrony, *def.asy* is climatic water deficit asynchrony, *cld.asy* is cloud cover asynchrony, *veg.ent* is vegetation structural entropy, *vrm.med* is the median vector ruggedness metric, and *x* and *y* are longitude and latitude (brackets indicating optional inclusion). We chose the random forest algorithm for this analysis because of its ability to robustly model non-linear relationships, suited to our expectation that phenological asynchrony would be driven by different and potentially interacting factors in different regions of the globe. We developed a comprehensive and conservative modeling workflow, described in detail below, then ran this workflow once for each combination of LSP dataset (NIR_v, SIF), neighborhood radius (50 km, 100 km, 150 km), and coordinate inclusion (geographic coordinates either included or excluded as covariates; we check sensitivity to this because of lack of consensus about how to handle spatial data within random forests; J. Li *et al.* 2011; Sekulić *et al.* 2020). This produced a final set of 12 models, the results of which we summarize in Table S3. Because we found that salient results are largely insensitive to choice of LSP dataset, neighborhood radius, and coordinate inclusion, we chose the 100 km, NIR_v-based, coordinates-included model as the main model to summarize in the text.

Prior to producing final results, we prepared the modeling data, tuned hyperparameters, and carried out feature selection. First, we coprojected the response and covariate rasters in a metric projection (EPSG: 3857) to ensure coordinates are expressed in meters, then stacked them and extracted their values at all non-masked pixels. Next, we carried out comprehensive hyperparameter tuning (Boehmke and Greenwell 2019), assessing model performance as a

function of five random forest tuning parameters (number of trees per forest: *n_{tree}* = 150, 200, 250, 300; fraction of observations to use in each tree, for tree decorrelation: *sample.fraction* = 0.3, 0.55, 0.8; minimum number of observations that can be captured by a node: *min.node.size* = 1, 3, 5, 10; size of random subset of variables from which to choose each node's split variable: *mtry*=2, 4; and whether to sample with replacement: *replace* = true, false) and a function of the fraction of the full global dataset used for modeling (*subset.frac* = 0.05, 0.005; drawn as a random subsample, quartile-stratified by the LSP response variable, to reduce the computational strain imposed by the size of the modeling dataset but not cause excessive information loss). We included geographic coordinates in all models used for hyperparameter tuning, as we intended to retain them in the main model assessed in the text (unless we found that predominant results were highly sensitive to their inclusion). We used as a performance metric the root mean squared error (RMSE) of the model fitted to a 60% training split of the subsampled global dataset, but found that the RMSE of the predictions made on the 40% test split yielded the same set of optimum-performance hyperparameter choices. Lastly, before running the final set of models we confirmed that none of our subsetted datasets contained highly collinear variables (i.e., $R^2 \geq 0.75$), and we used the Boruta feature-selection algorithm (Kursa and Rudnicki 2010) to select our final feature set (but found no features to be dropped).

We produced the final 12 models of results (Table S3) by constructing each model using the optimum hyperparameters indicated by our tuning results (*n_{tree}* = 300, *sample.fraction* = 0.8, *min.node.size* = 1, *mtry* = 3, *replace* = false, and *subset.frac* = 0.05). To assess each model, we calculated two variable importance metrics (permutation-based: *ranger*'s default importance metric, using the comparison between the accuracy of out-of-bag sample predictions and the same accuracy calculated after permuting covariate values, averaged over all trees; and 'SHAP': absolute Shapley Additive Explanations (SHAP) values (Lundberg and Lee 2017) summed across all predictions in a model's training dataset), as well as two metrics of overall model performance (R^2 and RMSE). Finally, we use each model to make predictions at each pixel in our global dataset and to calculate prediction errors (mapped in Fig. S13), as well as SHAP values (to help with spatial model interpretability). Noting minimal variability across models in the top-importance covariates (Table S3), we summarize the main model (100-km NIR_v, coordinates included) in the text and in Fig. 4, including all covariates' importance metrics (Fig. 4a), maps of the covariate maps (Fig. 4b) and SHAP values for the three top-importance (by SHAP value) covariates (Fig. 4c), and a SHAP-value summary map (Fig. 4d). The summary map combines the three SHAP maps shown in Fig. 4c into a single map in which each pixel's hue (yellow, cyan, or magenta) indicates the top covariate at that pixel (by absolute SHAP value), saturation indicates the scale of predominance (with pixels grading toward white as the standard deviation of their SHAP values approaches 0), and value (i.e., brightness) indicates magnitude of LSP asynchrony (such that the lowest-asynchrony locations fade to black), such that color indicates the likely predominant driver within a high-asynchrony area, white whiter colors indicating greater uncertainty about driver predominance or stronger driver codominance.

Regional relationships between climatic distance and phenological distance

To test the hypothesis that phenological asynchrony is less dependent on climatic difference at low latitudes than at higher latitudes, we ran an ensemble analysis. Each sub-analysis in the ensemble first uses clustering to delineate a global set of high-asynchrony regions,

then uses matrix regression to estimate the slope of the relationship between climatic distance and phenological distance (henceforth, the climate-phenology correlational strength) within each region, then finally assesses the relationship between regions' mean latitudes and their climate-phenology correlational strengths. To avoid sensitivity to the specific values chosen for the three parameters used in the region-delineation process we repeated sub-analyses across combinations of low, middle, and high values for those parameters, then used Monte Carlo analysis to yield final summary results across the ensemble.

First, to delineate high-asynchrony regions, we first converted our NIR_v LSP asynchrony map into a map of maximum-asynchrony pixels by setting all pixels \geq the 95th percentile asynchrony value to 1 and masking all other pixels. We then used the DBSCAN clustering algorithm (Ester *et al.* 1996) to cluster those high-asynchrony pixels into high-asynchrony clusters. We chose the DBSCAN algorithm because of its ability to robustly identify arbitrary-shape clusters around the high-density centers of a point set, without forcing all points to have cluster assignments, which was a good match for the noisiness of our asynchrony metric and thus of the maximum-asynchrony map. Finally, we used the alpha-complex algorithm (a straight-line edge variant of the alpha-hull algorithm), implemented in Python by the Alpha Shape Toolbox (Bellock 2021), to delineate those clusters as high-asynchrony regions. This algorithm allowed us to relax the convexity and contiguity assumptions of other hull-determination algorithms, and thus to flexibly delineate complexly-shaped regions (e.g., mountain arcs) without inevitably including all intervening geographic area as would occur with a convex hull.

To assess the relationship between a region's mean latitude and the strength of correlation between climatic and phenological distances within that region, we first standardized and stacked each of the 19 WorldClim bioclimatic variables (Fick and Hijmans 2017) and standardized our global map of 365-day fitted LSP curves. Then, for each delineated region, we executed the following algorithm:

1. Draw a set of 1000 random points within the region, then subset them to all N points that fall within non-masked NIR_v pixels;
2. Calculate the matrix of pairwise phenological distances ($dist_{phen}$) between all N points (calculated as 365-dimensional pairwise Euclidean distances between points' paired time series);
3. Calculate the matrix of pairwise climatic distances ($dist_{clim}$) between all N points (calculated as 19-dimensional pairwise Euclidean distances between points' vectors of bioclimatic values);
4. Calculate the matrix of pairwise geographic distances ($dist_{geog}$) between all N points (calculated as Great Circle distances);
5. Standardize all three variables (so that coefficients of all regressions are beta coefficients, and hence comparable), then run a multiple-matrix regression with randomization (MMRR; Wang 2013) on the resulting matrices, using the formula $dist_{phen} \sim dist_{clim} + dist_{geog}$;
6. Save the $dist_{clim}$ coefficient from the MMRR, as well as the region's mean pairwise standardized climatic distance and its mean absolute latitude

Given the possible sensitivity of this analysis to the specific values chosen for the two main DBSCAN parameters (epsilon (eps): the maximum geographic distance between two points

that can be considered to be in the same neighborhood, expressed in decimal degrees in our analysis; minimum number of samples (*min_samples*): the number of samples required within a neighborhood for a point to be considered as a core point) and the main alpha-complex parameter (*alpha*: a value controlling how edge members are chosen, and thus determining the maximum complexity of the hull's edge), we reran the above analysis for all combinations of reasonable low, middle, and high values for each of the three parameters: *eps* = 2, 3.5, 5; *min_samples* = 0.3, 0.45, 0.6; *alpha* = 0.25, 0.75, 1.25.

As a final step, we then summarized the ensemble results of all 27 ($=3^3$) region-delineation parameterizations by running the overarching, ensemble simple linear regression model $coeff_{dist,clim} \sim \beta_{lat}|lat_{med}|$, using β_{lat} to test the null hypothesis that the regions' climate-phenology correlational strength is equal across mean absolute latitude. Because this regression violates the assumption that independent variable samples are IID – each point represents a clustered and delineated high-asynchrony region, and those regions can overlap across distinct parameterizations of the sub-analyses – we use Monte Carlo analysis to generate an empirical P-value for β_{lat} in the ensemble linear regression model: We run 1000 iterations of the same regression, each time permuting the vector of $|lat_{med}|$ values, then calculate an empirical P-value as the fraction of the 1000 simulated β_{lat} that are at least as extreme as the true β_{lat} . We display the results of the ensemble linear regression as scatter plot and trend line (Fig. 5b), and display the results of the Monte Carlo analysis as a histogram depicting the empirical distribution of simulated β_{lat} and a vertical line depicting the true β_{lat} value relative to that distribution (Fig. 5c). To provide a spatially explicit geographic interpretation of the results of this analysis, we also map a summary of the ensemble results as a hexbin map (Fig. 5a), with the color of each hexbin indicating the mean β_{lat} of all high-asynchrony regions whose delineated alpha-complex polygon spatially intersects the hexbin.

Figures

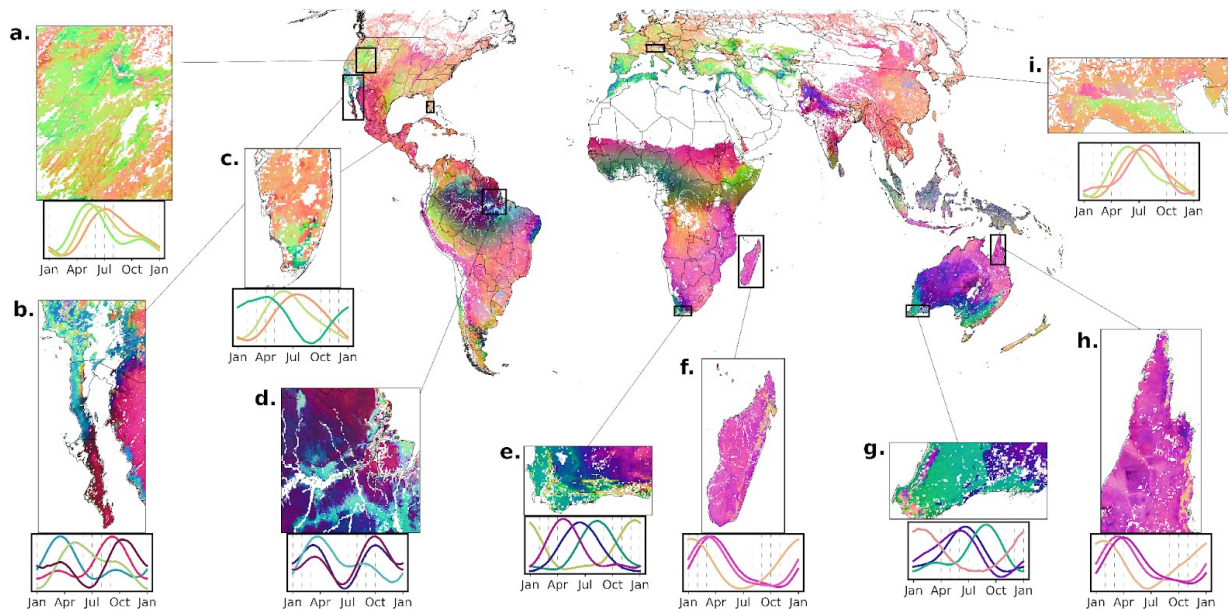


Figure 1: Phenological diversity: Global LSP visualization. Colors result from depicting in red, green, and blue the top three modes of our global empirical orthogonal functions (EOF) analysis (after folding across the intertropical convergence zone (ITCZ) any EOFs that display a north-south hemispheric dipole, for inter-hemispheric comparability; see methods). These three EOFs together explain >90% of the variation in the annual phenologies fitted to our filtered, global time series of MODIS near-infrared reflectance of vegetation (NIR_v). a-i: Zoomed maps of the focal regions discussed in the text, paired with line plots depicting the median fitted phenologies of the major color clusters occurring in the regions. Clusters were derived using K-means clustering, with K chosen by visual inspection of scree plots and regional maps; they facilitate visual interpretation but have no influence over any quantitative or statistical results. Regions (and the color-interpretation mentioned in the text) include: a. Great Basin, USA (green: cheatgrass; olive: sagebrush; orange: montane); b. Southern California, USA and Baja California, Mexico (lime green to blue: Mediterranean winter-monsoon climate region; pink to purple: summer-monsoon climate region) c. Southern Florida, USA (dark green-blue: Everglades sawgrass marsh; light green: wooded wetland; orange-red: drained, urban, upland, and mangrove); d. mouth of the Amazon River, Brazil (indigo and maroon: forest; light blue: non-forest); e. Cape Region, South Africa (green: non-forest Mediterranean vegetation; peach: isolated wet forests; dark blue, purple, and pink: dry woodland, shrubland and desert, and grassland); f. Madagascar (orange: coastal wet forests; pink and purple: seasonally dry forests, shrublands, and savannas); g. Southwest Australia (color-interpretation same as for e.); h. Cape York Peninsula, Australia (color-interpretation same as for f.); i. northern Italy and surroundings (rose: maize-rich agricultural regions; green: other agriculture, and non-forest vegetation; peach-orange: montane forest).

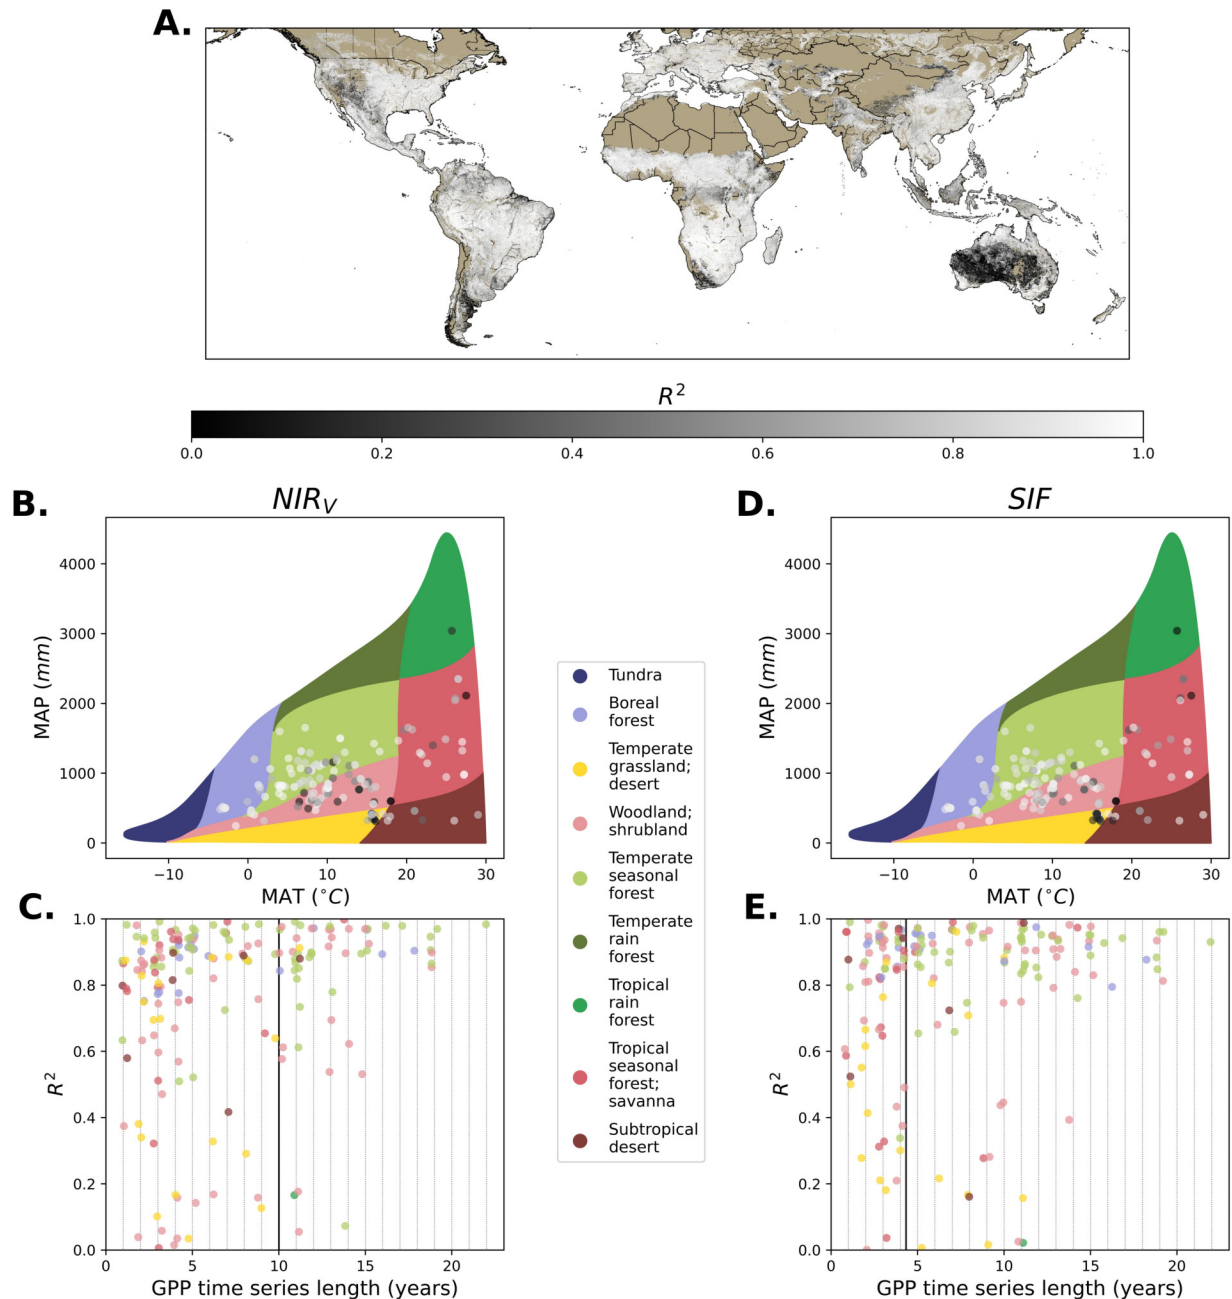


Figure 2: FLUXNET validation of NIR_v phenology: A. Agreement between NIR_v and SIF fitted phenologies. Each pixel shows the R^2 value (scaled from 0=black to 1=white) between the characteristic annual phenology curves fitted to that cell's NIR_v and SIF datasets. Areas of poor validation mainly cluster in temperate-to-hot semi-arid ecosystems, especially in the southern hemisphere, where interannual carbon-cycle variability depends largely on stochastic precipitation (Broich *et al.* 2014; Buitenwerf *et al.* 2015; Poulter *et al.* 2014) and may drive decorrelation between the characteristic LSP curves fitted to our 10-year NIR_v and 4.25-year SIF

samples. B. and D. Flux-tower validation performance across global biomes, for the NIR_v (B.) and SIF (D.) datasets. Each flux tower's site is plotted in the environmental space defined by the site's mean annual temperature (MAT) and mean annual precipitation (MAP), with the Whittaker biomes (Whittaker 1970) plotted beneath for context. Sites are colored by the R² (again scaled from 0=black to 1=white) between the characteristic annual phenology fitted to the remote sensing data at that site (or up to two map pixels away, ~11 km, distant) and the characteristic annual GPP phenology fitted, using identical methods, to the site's daily flux tower data. As in part A, the worst-performing flux tower sites also fall mostly within semi-arid and arid biomes typified by large interannual C-cycle variability. C. and E. Flux-tower validation performance of the NIR_v (C.) and SIF (E.) datasets, as a function of the length of available GPP time series. The worst-performing sites in our validation analysis frequently feature GPP time series shorter than the remote sensing time series used to estimate seasonal phenology (10 years for NIR_v; 4.25 years for SIF), and do not necessarily overlap or nest within the time period covered by the remote sensing time series, again likely contributing to poor validation because of differential temporal sampling of these regions' characteristically stochastic interannual photosynthetic variability.

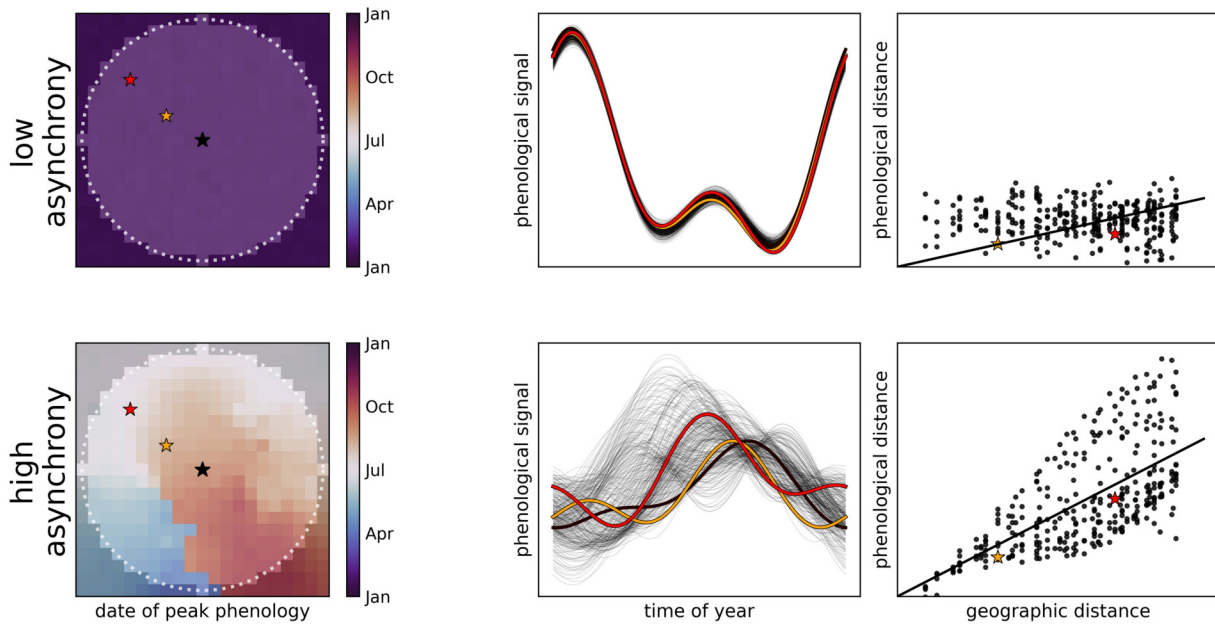
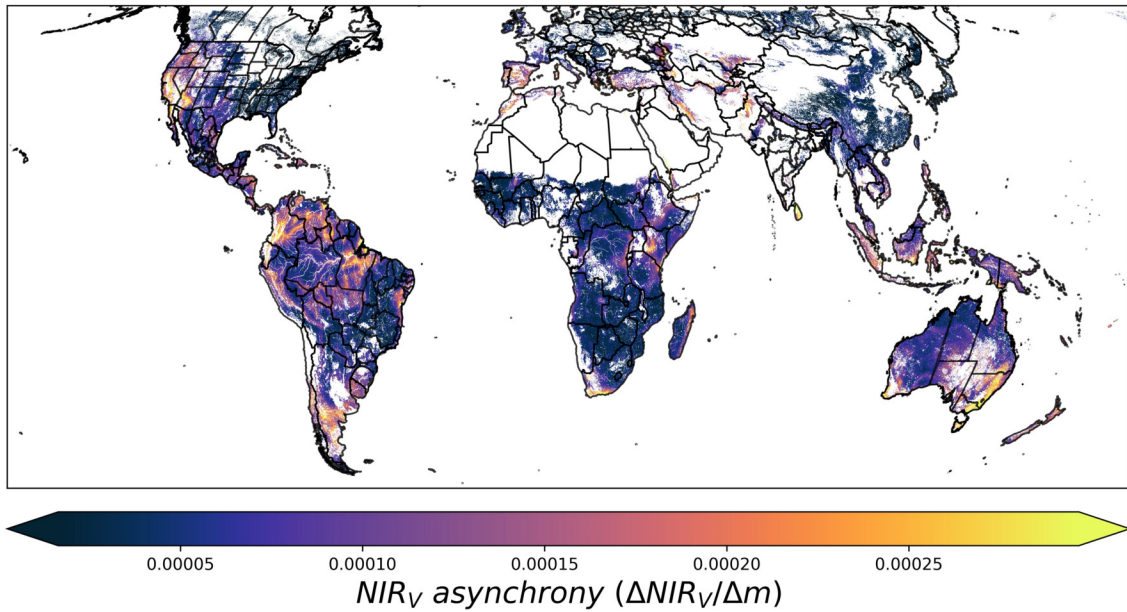
A.**B.**

Figure 3: Phenological asynchrony: A: Conceptual diagram depicting the stepwise calculation (from left to right) of spatial asynchrony for a focal pixel located in either of two simulated regions, one of low asynchrony (upper row) and one of high asynchrony (lower row). Maps in the left column depict spatial heterogeneity in phenology (depicted here as the day of the year of peak seasonality, as shown on the colorbar). The central focal pixel (black star) is the pixel for

which an asynchrony value is being calculated, using an analysis that considers the pairwise comparisons between the focal pixel and each other pixel inside the neighborhood (white dashed line). The line plots in the middle column depict the characteristic annual phenologies pertaining to each of the neighbor pixels inside the circular neighborhood (set of dim gray curves) and the focal pixel (bold black curve). Scatter plots in the right column depict the relationship between pairwise geographic distances and pairwise phenological distances, where each distance is calculated between the focal pixel and another pixel somewhere in the neighborhood. The slope of the trend line fitted to one of the scatter plots in the right column (by simple linear regression) is taken as the focal pixel's asynchrony metric. To demonstrate why this is so, two neighbor pixels – at shorter (orange) and further (red) geographic distances from the focal pixel – are tracked across the series of plots. In the low-asynchrony region, distance from the focal pixel does little to change the date of a neighbor pixel's peak phenology (top row, left) and little to change the Euclidean distance between the phenology curves at the focal and neighbor pixels (top row, center), such that the scatter plot for all neighbor pixels remains relatively flat (top row, right). However, because of the large phenological heterogeneity in a high-asynchrony region, neighbor pixels further from the focal pixel often fall in locations with considerably different phenologies from the focal pixel (bottom row, left); as neighbor pixels get further from the focal pixel they tend to have phenologies whose curves have greater Euclidean distances from the focal pixel's curve (bottom row, center; note that the orange and then red curves appear increasingly warped with respect to the bold black curve), and thus the full set of pairwise pixel comparisons builds up a geographic-phenological distance scatter plot with a strong slope that is indicative of the region's high asynchrony (bottom row, right).

B. Global LSP asynchrony map. This map shows the result of applying the asynchrony metric explained in A. to the annual LSP curve fitted at each pixel in our filtered, global dataset of MODIS near-infrared reflectance of vegetation (NIR_v). Brighter colors indicate increasing spatial asynchrony of land-surface phenology. (By dimensional analysis, asynchrony is expressed in change in units of the variable, which here (NIR_v) is unitless, per change in geographic distance (meters), but the magnitude of the mapped values is a function of the neighborhood radius used for the calculation and thus should be considered arbitrary, with only the relative values within a map being informative.)

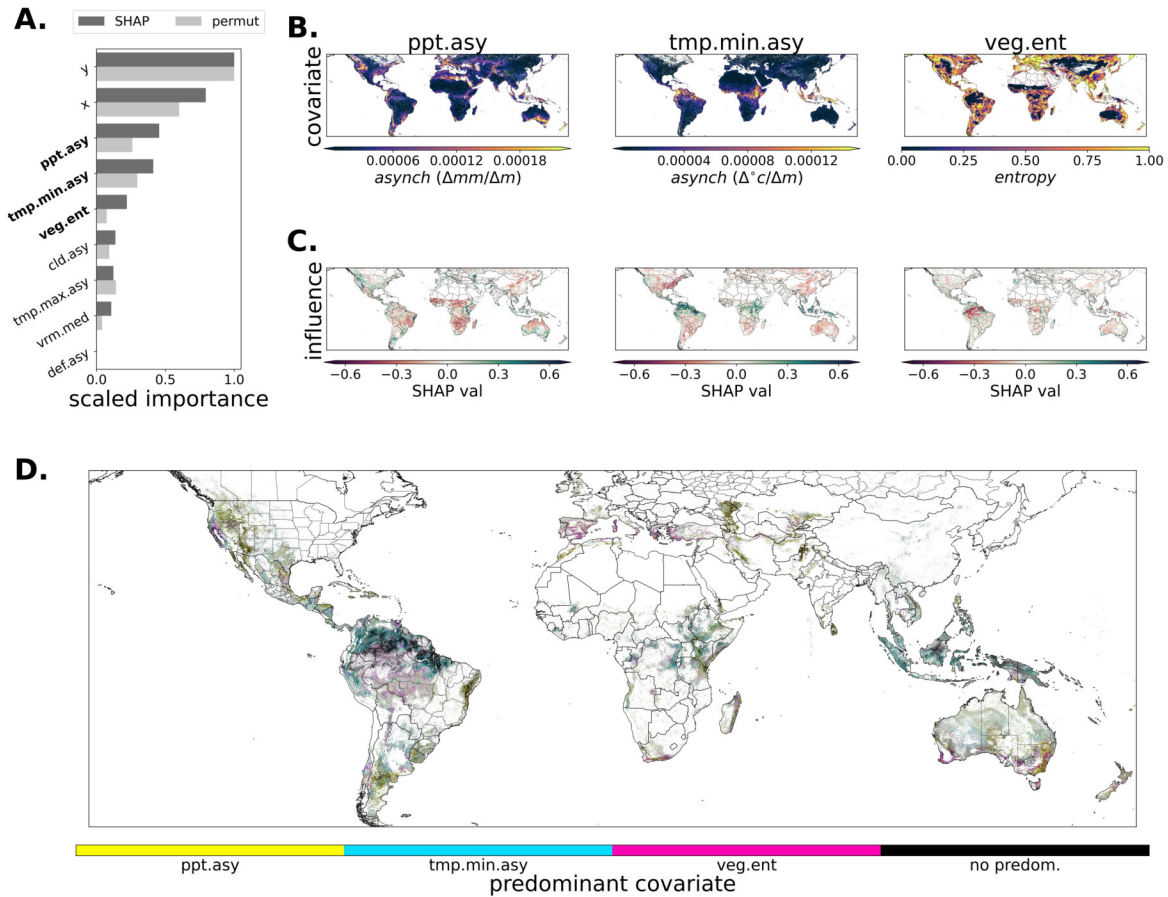


Figure 4: Correlates of spatial phenological asynchrony: A. Variable importance results for the main model (predicting 100-km neighborhood NIR_V LSP asynchrony, including geographic coordinates as covariates). Variable importance is shown using two metrics (mean SHAP values, in dark gray, and permutation-based importance, in light gray). Variables are ordered top-down by decreasing SHAP importance, and are: longitude (‘y’), latitude (‘x’), precipitation asynchrony (‘ppt.asy’), minimum temperature asynchrony (‘tmp.min.asy’), vegetation structural entropy (‘veg.ent’), cloud cover asynchrony (‘cld.asy’), maximum temperature asynchrony (‘tmp.max.asy’), median vector ruggedness metric (‘vrn.med’), and climatic water deficit asynchrony (‘def.asy’). Covariates are explained in detail in Table S2. B. Maps of the top three covariates in the main model: ppt.asy, tmp.min.asy, and veg.ent. These three top variables were largely consistent across all models run, with only vegetation structural entropy being outranked by cld.asy and tmp.max.asy depending on the importance metric and model (see Table S3 for all variable importance results). C. Maps of the top three covariates’ local Shapley (SHAP) values. For a given covariate’s map, each pixel value indicates the magnitude and direction of the covariate’s influence on the random forest’s predicted LSP asynchrony at that pixel; thus, each map provides an interpretation of the overall importance of a covariate and its regionally variable influence. D. Summary interpretation map. We combine the three SHAP maps shown in part C. into a single map in which each pixel’s hue (yellow, cyan, or magenta) indicates the top covariate

at that pixel (by absolute SHAP value), value indicates the scale of predominance (with pixels grading toward black as the standard deviation of their SHAP values approaches 0), and saturation (i.e., brightness) indicates magnitude of LSP asynchrony (such that the lowest-asynchrony locations fade toward white, down to the 50th percentile of asynchrony, beyond which all pixels are masked out for ease of interpretability). Thus, color indicates the likely predominant driver of LSP asynchrony within a high-asynchrony area, and blacker areas indicate greater uncertainty about driver predominance or stronger driver codominance.

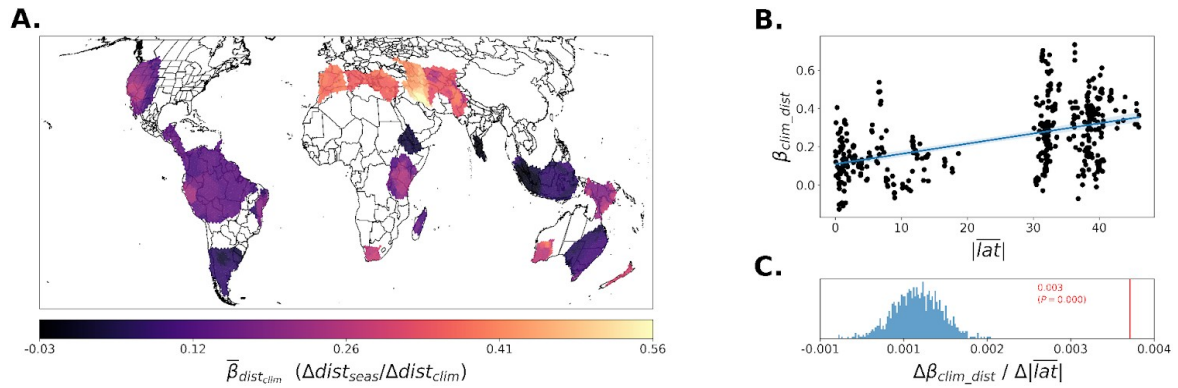


Figure 5: Latitudinal variation in the climate-dependence of phenological asynchrony

Summary results of an ensemble of MMRR models, run for high-asynchrony regions identified and delineated worldwide, estimating the strength of the relationship between seasonal and climatic differences (i.e., the magnitude of the fitted parameter β_{dist_clim}). The weaker this relationship within a given high-asynchrony region, the more climatically similar are sites with divergent phenologies, increasing the likelihood of allochronic isolation between separate populations of species with even narrow climatic niches. A. Geographic summary of results. Each hexbin's color represents the mean β_{dist_clim} value of all high-asynchrony regions whose polygons intersect the hexbin. B. Ensemble results (one point per delineated high-asynchrony region) reveal a pronounced positive relationship between β_{dist_clim} mean absolute latitude. C. Monte Carlo simulation across the ensemble results finds that the positive slope in part B. (0.003) is significantly greater than the slopes derived from permuted datasets (empirical P-value = 0.000).

Acknowledgments

We thank D. Ackerly, L. Anderegg, A. Bishop, T. Dawson, J. Frederick, N. Graham, M. Kelly, M. Kling, N. Muchhala, P. Papper, A. Turner, E. Westeen, G. Wogan, and M. Yuan for feedback and guidance on various iterations of the simulations presented herein. We thank Berkeley Research Computing for providing access to the Savio computing cluster. Lastly, we thank M. Terasaki Hart, C. Nemec-Hart, G. Hart, J. Hart, and M. Tylka for supporting and encouraging a lifetime of curiosity, and D. Byrne, A. Jamal, C. Mingus, and Y. Y. Ma for the good company. D.E.T.H. was supported by an Emerging Challenges in Tropical Science Graduate Student Fellowship from the Organization for Tropical Studies, by a Tinker Field Research Grant from the UC Berkeley Center for Latin American Studies, by a research equipment grant from IdeaWild, and by a Berkeley Fellowship. I.J.W. was supported by a National Science Foundation grant DEB1845682.

References

- Abatzoglou, J. T., Dobrowski, S. Z., Parks, S. A. & Hegewisch, K. C. (2018). TerraClimate, a high-resolution global dataset of monthly climate and climatic water balance from 1958–2015. *Scientific Data*, 5(1), 170191. <https://doi.org/10.1038/sdata.2017.191>
- Adole, T., Dash, J., Rodriguez-Galiano, V. & Atkinson, P. M. (2019). Photoperiod controls vegetation phenology across Africa. *Communications Biology*, 2(1), 391. <https://doi.org/10.1038/s42003-019-0636-7>
- Amatulli, G., Domisch, S., Tuanmu, M. N., Parmentier, B., Ranipeta, A., Malczyk, J., & Jetz, W. (2018). A suite of global, cross-scale topographic variables for environmental and biodiversity modeling. *Scientific data*, 5(1), 1-15.
- Badgley, G., Anderegg, L. D. L., Berry, J. A. & Field, C. B. (2019). Terrestrial gross primary production: Using NIR V to scale from site to globe. *Global Change Biology*, 25(11), 3731–3740. <https://doi.org/10.1111/gcb.14729>
- Badgley, G., Field, C. B. & Berry, J. A. (2017). Canopy near-infrared reflectance and terrestrial photosynthesis. *Science Advances*, 3(3), e1602244. <https://doi.org/10.1126/sciadv.1602244>
- Bellock, K. E. (2021). *Alpha shape toolbox*. <https://pypi.org/project/alphashape/>
- Bezanson, J., Edelman, A., Karpinski, S. & Shah, V. B. (2017). Julia: A Fresh Approach to Numerical Computing. *SIAM Review*, 59(1), 65–98. <https://doi.org/10.1137/141000671>
- Boehmke, B. & Greenwell, B. (2019). *Hands-On Machine Learning with R*. <https://doi.org/10.1201/9780367816377>
- Bonan, G. B. (2008). Forests and Climate Change: Forcings, Feedbacks, and the Climate Benefits of Forests. *Science*, 320(5882), 1444–1449. <https://doi.org/10.1126/science.1155121>
- Boos, W. R. & Pascale, S. (2021). Mechanical forcing of the North American monsoon by orography. *Nature*, 599(7886), 611–615. <https://doi.org/10.1038/s41586-021-03978-2>
- Bradley, B. A., Jacob, R. W., Hermance, J. F. & Mustard, J. F. (2007). A curve fitting procedure to derive inter-annual phenologies from time series of noisy satellite NDVI data. *Remote Sensing of Environment*, 106(2), 137–145. <https://doi.org/10.1016/j.rse.2006.08.002>
- Broich, M., Huete, A., Tulbure, M. G., Ma, X., Xin, Q., Paget, M., Restrepo-Coupe, N., Davies, K., Devadas, R. & Held, A. (2014). Land surface phenological response to decadal climate variability across Australia using satellite remote sensing. *Biogeosciences*, 11(18), 5181–5198. <https://doi.org/10.5194/bg-11-5181-2014>
- Buitenwerf, R., Rose, L. & Higgins, S. I. (2015). Three decades of multi-dimensional change in global leaf phenology. *Nature Climate Change*, 5(4), 364–368. <https://doi.org/10.1038/nclimate2533>
- Cavelier, J., Solis, D. & Jaramillo, M. A. (1996). Fog interception in montane forests across the Central Cordillera of Panamá. *Journal of Tropical Ecology*, 12(3), 357–369. <https://doi.org/10.1017/s026646740000955x>
- Chen, A., Meng, F., Mao, J., Ricciuto, D. & Knapp, A. K. (2022). Photosynthesis phenology, as defined by solar-induced chlorophyll fluorescence, is overestimated by vegetation indices in the extratropical Northern Hemisphere. *Agricultural and Forest Meteorology*, 323, 109027. <https://doi.org/10.1016/j.agrformet.2022.109027>

- Chen, R., Yin, G., Zhao, W., Xu, B., Zeng, Y., Liu, G. & Verger, A. (2022). TCNIRv: Topographically-corrected nearinfrared reflectance of vegetation for tracking gross primary production over mountainous areas. *IEEE Transactions on Geoscience and Remote Sensing*, *PP(99)*, 1–1. <https://doi.org/10.1109/tgrs.2022.3149655>
- Cohn, A. S., Bhattarai, N., Campolo, J., Crompton, O., Dralle, D., Duncan, J. & Thompson, S. (2019). Forest loss in Brazil increases maximum temperatures within 50 km. *Environmental Research Letters*, *14(8)*, 084047. <https://doi.org/10.1088/1748-9326/ab31fb>
- Coyne, J. A. & Orr, H. A. (2004). *Speciation* (Vol. 37). Sinauer Associates.
- d’Andrimont, R., Verhegghen, A., Lemoine, G., Kempeneers, P., Meroni, M. & Velde, M. van der. (2021). From parcel to continental scale – A first European crop type map based on Sentinel-1 and LUCAS Copernicus in-situ observations. *Remote Sensing of Environment*, *266*, 112708. <https://doi.org/10.1016/j.rse.2021.112708>
- Dechant, B., Ryu, Y., Badgley, G., Zeng, Y., Berry, J. A., Zhang, Y., Goulas, Y., Li, Z., Zhang, Q., Kang, M., Li, J. & Moya, I. (2020). Canopy structure explains the relationship between photosynthesis and sun-induced chlorophyll fluorescence in crops. *Remote Sensing of Environment*, *241*, 111733. <https://doi.org/10.1016/j.rse.2020.111733>
- Ester, M., Kriegel, H.-P., Sander, J. & Xu, X. (1996). A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. *KDD-96 Proceedings*.
- Etherington, T. R., Holland, E. P. & O’Sullivan, D. (2015). NLMpy: a python software package for the creation of neutral landscape models within a general numerical framework. *Methods in Ecology and Evolution*, *6(2)*, 164–168. <https://doi.org/10.1111/2041-210x.12308>
- Ettinger, A. K., Chamberlain, C. J. & Wolkovich, E. M. (2022). The increasing relevance of phenology to conservation. *Nature Climate Change*, *12(4)*, 305–307. <https://doi.org/10.1038/s41558-022-01330-8>
- Fick, S. E. & Hijmans, R. J. (2017). WorldClim 2: new 1-km spatial resolution climate surfaces for global land areas. *International Journal of Climatology*, *37(12)*, 4302–4315. <https://doi.org/10.1002/joc.5086>
- Freeman, B. G., Strimas-Mackey, M. & Miller, E. T. (2022). Interspecific competition limits bird species’ ranges in tropical mountains. *Science*, *377(6604)*, 416–420. <https://doi.org/10.1126/science.abl7242>
- Friedl, M. & Sulla-Menashe, D. (2019). *MCD12Q1 MODIS/Terra+Aqua Land Cover Type Yearly L3 Global 500m SIN Grid V006 [Data set]*. NASA EOSDIS Land Processes DAAC. <https://doi.org/10.5067/MODIS/MCD12Q1.006>
- Fu, Y., Li, X., Zhou, X., Geng, X., Guo, Y. & Zhang, Y. (2020). Progress in plant phenology modeling under global climate change. *Science China Earth Sciences*, *63(9)*, 1237–1247. <https://doi.org/10.1007/s11430-019-9622-2>
- Gamba, D. & Muchhala, N. (2020). Global patterns of population genetic differentiation in seed plants. *Molecular Ecology*, *29(18)*, 3413–3428. <https://doi.org/10.1111/mec.15575>
- Garonna, I., Jong, R. & Schaepman, M. E. (2016). Variability and evolution of global land surface phenology over the past three decades (1982–2012). *Global Change Biology*, *22(4)*, 1456–1468. <https://doi.org/10.1111/gcb.13168>

- Gatti, L. V., Basso, L. S., Miller, J. B., Gloor, M., Domingues, L. G., Cassol, H. L. G., Tejada, G., Aragão, L. E. O. C., Nobre, C., Peters, W., Marani, L., Arai, E., Sanches, A. H., Corrêa, S. M., Anderson, L., Randow, C. V., Correia, C. S. C., Crispim, S. P. & Neves, R. A. L. (2021). Amazonia as a carbon source linked to deforestation and climate change. *Nature*, 595(7867), 388–393. <https://doi.org/10.1038/s41586-021-03629-6>
- Ghalambor, C. K., Huey, R. B., Martin, P. R., Tewksbury, J. J. & Wang, G. (2006). Are mountain passes higher in the tropics? Janzen’s hypothesis revisited. *Integrative and Comparative Biology*, 46(1), 5–17. <https://doi.org/10.1093/icb/icj003>
- Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D. & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202, 18–27. <https://doi.org/10.1016/j.rse.2017.06.031>
- Guarnizo, C. E., Montoya, P., Quintero, I. & Cadena, C. D. (2022). Population divergence associated with spatial asynchrony in precipitation in Neotropical frogs. *Journal of Biogeography*. <https://doi.org/10.1111/jbi.14484>
- Harris, I., Jones, P. D., Osborn, T. J. & Lister, D. H. (2014). Updated high-resolution grids of monthly climatic observations – the CRU TS3.10 Dataset. *International Journal of Climatology*, 34(3), 623–642. <https://doi.org/10.1002/joc.3711>
- Hendry, A. P. & Day, T. (2005). Population structure attributable to reproductive time: isolation by time and adaptation by time. *Molecular Ecology*, 14(4), 901–916. <https://doi.org/10.1111/j.1365-294x.2005.02480.x>
- Homer, C. G., Dewitz, J. A., Yang, L., Jin, S., Danielson, P., Xian, G., Coulston, J., Herold, N. D., Wickham, J. D. & Megown, K. (2015). Completion of the 2011 National Land Cover Database for the conterminous United States: Representing a decade of land cover change information. *Photogrammetric Engineering and Remote Sensing*, 81(5), 345–354. https://www.asprs.org/a/publications/pers/2015journals/PERS_May_2015/HTML/index.html
- Huang, M., Piao, S., Ciais, P., Peñuelas, J., Wang, X., Keenan, T. F., Peng, S., Berry, J. A., Wang, K., Mao, J., Alkama, R., Cescatti, A., Cuntz, M., Deurwaerder, H. D., Gao, M., He, Y., Liu, Y., Luo, Y., Myneni, R. B., ... Janssens, I. A. (2019). Air temperature optima of vegetation productivity across global biomes. *Nature Ecology & Evolution*, 3(5), 772–779. <https://doi.org/10.1038/s41559-019-0838-x>
- Janzen, D. H. (1967). Why Mountain Passes are Higher in the Tropics. *The American Naturalist*, 101(919), 233–249. <https://doi.org/10.1086/282487>
- Kim, J., Ryu, Y., Dechant, B., Lee, H., Kim, H. S., Kornfeld, A. & Berry, J. A. (2021). Solar-induced chlorophyll fluorescence is non-linearly related to canopy photosynthesis in a temperate evergreen needleleaf forest during the fall transition. *Remote Sensing of Environment*, 258, 112362. <https://doi.org/10.1016/j.rse.2021.112362>
- Knoben, W. J. M. (2019). Global bimodal precipitation seasonality : A systematic overview. *International Journal of Climatology*, 39(September 2017), 558–567. <https://doi.org/10.1002/joc.5786>
- Köhler, P. & Frankenberg, C. (2020). *Ungridded TROPOMI SIF (at 740nm) (1.0) [Data set]*. CaltechDATA. <https://doi.org/10.22002/D1.1347>
- Köhler, P., Frankenberg, C., Magney, T. S., Guanter, L., Joiner, J. & Landgraf, J. (2018). Global Retrievals of Solar-Induced Chlorophyll Fluorescence With TROPOMI: First Results and

- Intersensor Comparison to OCO-2. *Geophysical Research Letters*, 45(19), 10,456-10,463. <https://doi.org/10.1029/2018gl079031>
- Kursa, M. B. & Rudnicki, W. R. (2010). Feature Selection with the Boruta Package. *Journal of Statistical Software*, 36(11).
- Li, J., Heap, A. D., Potter, A. & Daniell, J. J. (2011). Application of machine learning methods to spatial interpolation of environmental variables. *Environmental Modelling & Software*, 26(12), 1647–1659. <https://doi.org/10.1016/j.envsoft.2011.07.004>
- Li, X. & Xiao, J. (2019). A Global, 0.05-Degree Product of Solar-Induced Chlorophyll Fluorescence Derived from OCO-2, MODIS, and Reanalysis Data. *Remote Sensing*, 11(5), 517. <https://doi.org/10.3390/rs11050517>
- Lloyd, D. (1990). A phenological classification of terrestrial vegetation cover using shortwave vegetation index imagery. *International Journal of Remote Sensing*, 11(12), 2269–2279. <https://doi.org/10.1080/01431169008955174>
- Lovejoy, T. E. & Nobre, C. (2018). Amazon Tipping Point. *Science Advances*, 4(2), eaat2340. <https://doi.org/10.1126/sciadv.aat2340>
- Lundberg, S. & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. *ArXiv*.
- Maestas, J., Jones, M., Pastick, N. J., Rigge, M. B., Wylie, B. K., Garner, L., Crist, M., Homer, C., Boyte, S. & Witacre, B. (2020). *Annual Herbaceous Cover across Rangelands of the Sagebrush Biome: U.S. Geological Survey data release*. <https://doi.org/10.5066/P9VL3LD5>
- Martin, P., Bonier, F., Moore, I. & Tewksbury, J. (2009). Latitudinal variation in the asynchrony of seasons: implications for higher rates of population differentiation and speciation in the tropics. *Ideas in Ecology and Evolution*, 2. <https://doi.org/10.4033/iee.2009.2.3.n>
- Moore, I. T., Bonier, F. & Wingfield, J. C. (2005). Reproductive asynchrony and population divergence between two tropical bird populations. *Behavioral Ecology*, 16(4), 755–762. <https://doi.org/10.1093/beheco/ari049>
- Morellato, L. P. C., Alberton, B., Alvarado, S. T., Borges, B., Buisson, E., Camargo, M. G. G., Cancian, L. F., Carstensen, D. W., Escobar, D. F. E., Leite, P. T. P., Mendoza, I., Rocha, N. M. W. B., Soares, N. C., Silva, T. S. F., Staggemeier, V. G., Streher, A. S., Vargas, B. C. & Peres, C. A. (2016). Linking plant phenology to conservation biology. *Biological Conservation*, 195, 60–72. <https://doi.org/10.1016/j.biocon.2015.12.033>
- Myers, N., Mittermeier, R. A., Mittermeier, C. G., Fonseca, G. A. B. da & Kent, J. (2000). Biodiversity hotspots for conservation priorities. *Nature*, 403.
- Notaro, M., Chen, G. & Liu, Z. (2011). Vegetation Feedbacks to Climate in the Global Monsoon Regions*. *Journal of Climate*, 24(22), 5740–5756. <https://doi.org/10.1175/2011jcli4237.1>
- Pastorello, G., Trotta, C., Canfora, E., Chu, H., Christianson, D., Cheah, Y.-W., Poindexter, C., Chen, J., Elbashandy, A., Humphrey, M., Isaac, P., Polidori, D., Reichstein, M., Ribeca, A., Ingen, C. van, Vuichard, N., Zhang, L., Amiro, B., Ammann, C., ... Papale, D. (2020). The FLUXNET2015 dataset and the ONEFlux processing pipeline for eddy covariance data. *Scientific Data*, 7(1), 225. <https://doi.org/10.1038/s41597-020-0534-3>
- Peñuelas, J., Rutishauser, T. & Filella, I. (2009). Phenology Feedbacks on Climate Change. *Science*, 324(5929), 887–888. <https://doi.org/10.1126/science.1173004>
- Phillips, R. P., Ibáñez, I., D'Orangeville, L., Hanson, P. J., Ryan, M. G. & McDowell, N. G. (2016). A belowground perspective on the drought sensitivity of forests: Towards

- improved understanding and simulation. *Forest Ecology and Management*, 380, 309–320. <https://doi.org/10.1016/j.foreco.2016.08.043>
- Piao, S., Liu, Q., Chen, A., Janssens, I. A., Fu, Y., Dai, J., Liu, L., Lian, X., Shen, M. & Zhu, X. (2019). Plant phenology and global climate change: Current progresses and challenges. *Global Change Biology*, 25(6), 1922–1940. <https://doi.org/10.1111/gcb.14619>
- Pickering, M., Cescatti, A. & Duveiller, G. (2022). Sun-Induced Fluorescence as a Proxy of Primary Productivity across Vegetation Types and Climates. *Biogeosciences Discussions*, 2022, 1–33. <https://doi.org/10.5194/bg-2021-354>
- Pielou, E. C. (1966). The measurement of diversity in different types of biological collections. *Journal of Theoretical Biology*, 13, 131–144. [https://doi.org/10.1016/0022-5193\(66\)90013-0](https://doi.org/10.1016/0022-5193(66)90013-0)
- Poulter, B., Frank, D., Ciais, P., Myneni, R. B., Andela, N., Bi, J., Broquet, G., Canadell, J. G., Chevallier, F., Liu, Y. Y., Running, S. W., Sitch, S. & Werf, G. R. van der. (2014). Contribution of semi-arid ecosystems to interannual variability of the global carbon cycle. *Nature*, 509(7502), 600–603. <https://doi.org/10.1038/nature13376>
- Quintero, I., González-Caro, S., Zalamea, P.-C. & Cadena, C. D. (2014). Asynchrony of Seasons: Genetic Differentiation Associated with Geographic Variation in Climatic Seasonality and Reproductive Phenology. *The American Naturalist*, 184(3), 352–363. <https://doi.org/10.1086/677261>
- R Core Team. (2021). *R: A Language and Environment for Statistical Computing*. <https://www.R-project.org/>
- Richardson, A. D., Anderson, R. S., Arain, M. A., Barr, A. G., Bohrer, G., Chen, G., Chen, J. M., Ciais, P., Davis, K. J., Desai, A. R., Dietze, M. C., Dragoni, D., Garrity, S. R., Gough, C. M., Grant, R., Hollinger, D. Y., Margolis, H. A., McCaughey, H., Migliavacca, M., ... Xue, Y. (2012). Terrestrial biosphere models need better representation of vegetation phenology: results from the North American Carbon Program Site Synthesis. *Global Change Biology*, 18(2), 566–584. <https://doi.org/10.1111/j.1365-2486.2011.02562.x>
- Richardson, A. D., Keenan, T. F., Migliavacca, M., Ryu, Y., Sonnentag, O. & Toomey, M. (2013). Climate change, phenology, and phenological control of vegetation feedbacks to the climate system. *Agricultural and Forest Meteorology*, 169, 156–173. <https://doi.org/10.1016/j.agrformet.2012.09.012>
- Rohatgi, A. (2022, September 15). *WebPlotDigitizer, Version 4.6*. <https://automeris.io/WebPlotDigitizer>
- Sakai, S., & Kitajima, K. (2019). Tropical phenology: Recent advances and perspectives. *Ecological Research*, 34(1), 50–54.
- Salati, E., Dall'Olio, A., Matsui, E. & Gat, J. R. (1979). Recycling of water in the Amazon Basin: An isotopic study. *Water Resources Research*, 15(5), 1250–1258. <https://doi.org/10.1029/wr015i005p01250>
- Sampaio, G., Nobre, C., Costa, M. H., Satyamurty, P., Soares-Filho, B. S. & Cardoso, M. (2007). Regional climate change over eastern Amazonia caused by pasture and soybean cropland expansion. *Geophysical Research Letters*, 34(17). <https://doi.org/10.1029/2007gl030612>
- Sappington, J. M., Longshore, K. M. & Thompson, D. B. (2007). Quantifying Landscape Ruggedness for Animal Habitat Analysis: A Case Study Using Bighorn Sheep in the

- Mojave Desert. *Journal of Wildlife Management*, 71(5), 1419–1426.
<https://doi.org/10.2193/2005-723>
- Schaaf, C. & Wang, Z. (2015). *MCD43A4 MODIS/Terra+Aqua BRDF/Albedo Nadir BRDF Adjusted Ref Daily L3 Global - 500m V006 [Data set]*. NASA EOSDIS Land Processes DAAC. <https://doi.org/10.5067/MODIS/MCD43A4.006>
- Scholl, M. A., Giambelluca, T. W., Gingerich, S. B., Nullet, M. A. & Loope, L. L. (2007). Cloud water in windward and leeward mountain forests: The stable isotope signature of orographic cloud water. *Water Resources Research*, 43(12).
<https://doi.org/10.1029/2007wr006011>
- Sekulić, A., Kilibarda, M., Heuvelink, G. B. M., Nikolić, M. & Bajat, B. (2020). Random Forest Spatial Interpolation. *Remote Sensing*, 12(10), 1687. <https://doi.org/10.3390/rs12101687>
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
- Sheldon, K. S., Huey, R. B., Kaspari, M. & Sanders, N. J. (2018). Fifty Years of Mountain Passes: A Perspective on Dan Janzen’s Classic Article. *The American Naturalist*, 191(5), 553–565. <https://doi.org/10.1086/697046>
- Sun, Y., Frankenberg, C., Wood, J. D., Schimel, D. S., Jung, M., Guanter, L., Drewry, D. T., Verma, M., Porcar-Castell, A., Griffis, T. J., Gu, L., Magney, T. S., Köhler, P., Evans, B. & Yuen, K. (2017). OCO-2 advances photosynthesis observation from space via solar-induced chlorophyll fluorescence. *Science*, 358(6360).
<https://doi.org/10.1126/science.aam5747>
- Tang, J., Körner, C., Muraoka, H., Piao, S., Shen, M., Thackeray, S. J. & Yang, X. (2016). Emerging opportunities and challenges in phenology: a review. *Ecosphere*, 7(8).
<https://doi.org/10.1002/ecs2.1436>
- Taylor, R. S. & Friesen, V. L. (2017). The role of allochryny in speciation. *Molecular Ecology*, 26(13), 3330–3342. <https://doi.org/10.1111/mec.14126>
- terSteege, H., Pitman, N. C. A., Sabatier, D., Baraloto, C., Salomão, R. P., Guevara, J. E., Phillips, O. L., Castilho, C. V., Magnusson, W. E., Molino, J.-F., Monteagudo, A., Vargas, P. N., Montero, J. C., Feldpausch, T. R., Coronado, E. N. H., Killeen, T. J., Mostacedo, B., Vasquez, R., Assis, R. L., ... Silman, M. R. (2013). Hyperdominance in the Amazonian Tree Flora. *Science*, 342(6156), 1243092–1243092.
<https://doi.org/10.1126/science.1243092>
- Thomé, M. T. C., Carstens, B. C., Rodrigues, M. T., Jr, P. M. G., Alexandrino, J. & Haddad, C. F. B. (2021). A role of asynchrony of seasons in explaining genetic differentiation in a Neotropical toad. *Heredity* 2021, 127(4), 1–10. <https://doi.org/10.1038/s41437-021-00460-7>
- Turner, A. J., Köhler, P., Magney, T. S., Frankenberg, C., Fung, I. & Cohen, R. C. (2020). A double peak in the seasonality of California’s photosynthesis as observed from space. *Biogeosciences*, 17(2), 405–422. <https://doi.org/10.5194/bg-17-405-2020>
- Vermote, E. & Wolfe, R. (2015a). *MOD09GA MODIS/Terra Surface Reflectance Daily L2G Global 1kmand 500m SIN Grid V006 [Data set]*. NASA EOSDIS Land Processes DAAC. <https://doi.org/10.5067/MODIS/MOD09GA.006>

- Vermote, E. & Wolfe, R. (2015b). *MYD09GA MODIS/Aqua Surface Reflectance Daily L2G Global 1kmand 500m SIN Grid V006 [Data set]*. NASA EOSDIS Land Processes DAAC. <https://doi.org/10.5067/MODIS/MYD09GA.006>
- Walther, S., Voigt, M., Thum, T., Gonsamo, A., Zhang, Y., Köhler, P., Jung, M., Varlagin, A. & Guanter, L. (2016). Satellite chlorophyll fluorescence measurements reveal large-scale decoupling of photosynthesis and greenness dynamics in boreal evergreen forests. *Global Change Biology*, 22(9), 2979–2996. <https://doi.org/10.1111/gcb.13200>
- Wang, I. J. (2013). Examining the full effects of landscape heterogeneity on spatial genetic variation: a multiple matrix regression approach for quantifying geographic and ecological isolation. *Evolution*, 67(12), 3403–3411. <https://doi.org/10.1111/evo.12134>
- Wang, S., Zhang, Y., Ju, W., Wu, M., Liu, L., He, W. & Peñuelas, J. (2022). Temporally corrected long-term satellite solar-induced fluorescence leads to improved estimation of global trends in vegetation photosynthesis during 1995–2018. *ISPRS Journal of Photogrammetry and Remote Sensing*, 194, 222–234. <https://doi.org/10.1016/j.isprsjprs.2022.10.018>
- Whittaker, R. H. (1970). *Communities and Ecosystems*. In *Communities and Ecosystems* (Vol. 42, p. 154). <https://doi.org/10.2307/4511758>
- Wright, M. N. & Ziegler, A. (2017). ranger: A Fast Implementation of Random Forests for High Dimensional Data in C++ and R. *Journal of Statistical Software*, 77(1), 1–17. <https://doi.org/10.18637/jss.v077.i01>
- Yang, J. C., Magney, T. S., Albert, L. P., Richardson, A. D., Frankenberg, C., Stutz, J., Grossmann, K., Burns, S. P., Seyednasrollah, B., Blanken, P. D. & Bowling, D. R. (2022). Gross primary production (GPP) and red solar induced fluorescence (SIF) respond differently to light and seasonal environmental conditions in a subalpine conifer forest. *Agricultural and Forest Meteorology*, 317, 108904. <https://doi.org/10.1016/j.agrformet.2022.108904>
- Yu, L., Wen, J., Chang, C. Y., Frankenberg, C. & Sun, Y. (2019). High-Resolution Global Contiguous SIF of OCO-2. *Geophysical Research Letters*, 46(3), 1449–1458. <https://doi.org/10.1029/2018gl081109>
- Zeng, L., Wardlow, B. D., Xiang, D., Hu, S. & Li, D. (2020). A review of vegetation phenological metrics extraction using time-series, multispectral satellite data. *Remote Sensing of Environment*, 237, 111511. <https://doi.org/10.1016/j.rse.2019.111511>
- Zeng, Y., Hao, D., Huete, A., Dechant, B., Berry, J., Chen, J. M., Joiner, J., Frankenberg, C., Bond-Lamberty, B., Ryu, Y., Xiao, J., Asrar, G. R. & Chen, M. (2022). Optical vegetation indices for monitoring terrestrial ecosystems globally. *Nature Reviews Earth & Environment*, 1–17. <https://doi.org/10.1038/s43017-022-00298-5>
- Zhang, C., Brodylo, D., Sirianni, M. J., Li, T., Comas, X., Douglas, T. A. & Starr, G. (2021). Mapping CO2 fluxes of cypress swamp and marshes in the Greater Everglades using eddy covariance measurements and Landsat data. *Remote Sensing of Environment*, 262, 112523. <https://doi.org/10.1016/j.rse.2021.112523>
- Zhang, J., Xiao, J., Tong, X., Zhang, J., Meng, P., Li, J., Liu, P. & Yu, P. (2022). NIRv and SIF better estimate phenology than NDVI and EVI: Effects of spring and autumn phenology on ecosystem production of planted forests. *Agricultural and Forest Meteorology*, 315, 108819. <https://doi.org/10.1016/j.agrformet.2022.108819>

- Zhang, M., Berry, J., Shiga, Y. P., Doughty, R. B., Li, X., Xiao, J., Wen, J., Sun, Y. & Miller, S. (n.d.). *Solar-induced fluorescence products show variable skill in constraining global patterns in biospheric CO₂ fluxes*. <https://doi.org/10.1002/essoar.10511426.1>
- Zhisheng, A., Guoxiong, W., Jianping, L., Youbin, S., Yimin, L., Weijian, Z., Yanjun, C., Anmin, D., Li, L., Jiangyu, M., Hai, C., Zhengguo, S., Liangcheng, T., Hong, Y., Hong, A., Hong, C. & Juan, F. (2015). Global Monsoon Dynamics and Climate Change. *Annual Review of Earth and Planetary Sciences*, 43(1), 1–49. <https://doi.org/10.1146/annurev-earth-060313-054623>
- Zhong, L., Yu, L., Li, X., Hu, L. & Gong, P. (2016). Rapid corn and soybean mapping in US Corn Belt and neighboring areas. *Scientific Reports*, 6(1), 36240. <https://doi.org/10.1038/srep36240>

Supplemental Tables

var	neigh_comp	r2
NIRv	50 100	0.899428419489627
	50 150	0.815051543652447
	100 150	0.95337020709664
SIF	50 100	0.922104746726498
	50 150	0.855458768735066
	100 150	0.963384263699549
tmp.min	50 100	0.916903137228962
	50 150	0.829598693399889
	100 150	0.964465064925608
tmp.max	50 100	0.908327519666317
	50 150	0.810494307011837
	100 150	0.962091739192981
ppt	50 100	0.905029141661629
	50 150	0.805820312505186
	100 150	0.958061440484181
def	50 100	0.885057581983968
	50 150	0.768739569782019
	100 150	0.952231915095714
cld	50 100	0.921166001980917
	50 150	0.836354336807822
	100 150	0.964853126234789

Table S1: R²s of inter-neighborhood asynchrony map comparisons for all asynchrony variables produced The left column lists all variables for which we calculated asynchrony maps (abbreviations: tmp.min = TerraClimate minimum temperature; tmp.max = maximum temperature; ppt = TerraClimate precipitation; def = TerraClimate climatic water deficit; cld =

MODIS cloud cover). The middle column shows the two neighborhood radii for which the R^2 is calculated.

covar	description	data source
<i>tmp.min.asy</i>	spatial asynchrony in the seasonality of monthly minimum temperature, based on 73 years (01/01/1958 - 12/01/2021) of TerraClimate monthly means	https://developers.google.com/earth-engine/datasets/catalog/IDAHO_EPSCOR_TERRACLIMATE
<i>tmp.max.asy</i>	spatial asynchrony in the seasonality of monthly maximum temperature, based on 73 years (01/01/1958 - 12/01/2021) of TerraClimate monthly means	https://developers.google.com/earth-engine/datasets/catalog/IDAHO_EPSCOR_TERRACLIMATE
<i>ppt.asy</i>	spatial asynchrony in the seasonality of monthly precipitation, based on 73 years (01/01/1958 - 12/01/2021) of TerraClimate monthly means	https://developers.google.com/earth-engine/datasets/catalog/IDAHO_EPSCOR_TERRACLIMATE
<i>def.asy</i>	spatial asynchrony in the seasonality of monthly climatic water deficit, based on 73 years (01/01/1958 - 12/01/2021) of TerraClimate monthly means	https://developers.google.com/earth-engine/datasets/catalog/IDAHO_EPSCOR_TERRACLIMATE
<i>cld.asy</i>	spatial asynchrony in the seasonality of cloud cover, based on 10 years (01/01/2010 - 01/01/2020) of MODIS Aqua (MYD09GA v.006) and Terra (MOD09GA v.006) daily cloud band data	https://developers.google.com/earth-engine/datasets/catalog/MODIS_006_MOD09GA https://developers.google.com/earth-engine/datasets/catalog/MODIS_006_MYD09GA
<i>vrn.med</i>	median of the vector ruggedness metric (which an index of topographic complexity), calculated from SRTM DEM data and calculated at within ~50 km-resolution pixels	http://www.earthenv.org/topography
<i>veg.ent</i>	entropy, in 100 km-radius circular neighborhoods, of vegetation structural types, reclassified from the MODIS v.006 annual land cover time series	https://developers.google.com/earth-engine/datasets/catalog/MODIS_006_MCD12Q1
<i>x</i>	pixel-center longitude	n/a
<i>y</i>	pixel-center latitude	n/a

Table S2: Synopsis of covariates for phenological asynchrony random forest model Table includes the abbreviated code for each covariate (left column; these codes also appear in Figure 4), a detailed description of the metric calculated (center column), and the URL of the original data source (right column).

		SHAP IMPORTANCE						PERMUT IMPORTANCE						MODEL SUMMARY									
		NIRv			SIF			NIRv			SIF			NIRv			SIF						
		neigh rad	50	100	150	50	100	150	neigh rad	50	100	150	50	100	150	neigh rad	50	100	150	50	100	150	
COORDS	cld.asy		9.225.32	8.217.89	7.873.60	12.343.97	10.134.40	9.928.15		0.37	0.35	0.34	0.40	0.40	0.38								
	def.asy		4.951.81	4.720.25	4.921.52	5.703.22	5.397.73	6.183.77		0.24	0.22	0.24	0.24	0.26	0.25		RMSE	0.50	0.44	0.42	0.49	0.46	0.43
	ppt.asy		16.513.37	16.396.75	17.087.65	13.152.29	13.405.43	14.979.50		0.62	0.59	0.67	0.47	0.56	0.60		R2	0.75	0.81	0.83	0.76	0.79	0.82
	tmp.max.asy		7.224.10	7.892.92	6.192.93	5.959.35	6.943.61	9.332.95		0.38	0.42	0.48	0.31	0.39	0.45								
	tmp.min.asy		15.727.16	15.276.42	14.528.57	11.116.95	11.315.15	11.964.51		0.62	0.64	0.69	0.42	0.44	0.58								
	veg.ent		10.325.14	10.363.99	9.847.93	9.699.29	9.217.36	8.512.24		0.33	0.32	0.30	0.40	0.38	0.33								
	vrn.med		8.994.35	7.442.87	7.012.78	9.950.51	7.303.11	7.096.19		0.30	0.27	0.24	0.33	0.27	0.27								
	x		24.238.15	25.073.43	24.491.64	23.845.28	24.103.61	23.611.99		1.04	1.08	1.11	1.09	1.10	1.08								
	y		29.366.17	30.370.73	30.889.52	32.332.98	33.378.70	37.174.29		1.54	1.65	1.82	1.46	1.54	1.64								
	NO COORDS	cld.asy		14,503.21	14,598.44	15,258.10	18,123.76	16,514.77	16,109.31		0.51	0.52	0.58	0.56	0.60	0.61							
def.asy			11,491.19	11,817.77	11,170.20	12,734.40	11,864.31	12,504.65		0.35	0.41	0.42	0.39	0.43	0.47		RMSE	0.70	0.63	0.58	0.69	0.63	0.57
ppt.asy			31,976.58	29,810.09	29,997.22	24,386.77	25,836.93	25,294.61		0.81	0.87	0.94	0.71	0.87	0.91		R2	0.51	0.61	0.56	0.53	0.61	0.67
tmp.max.asy			14,996.67	16,522.85	18,222.31	16,999.85	17,866.98	19,921.62		0.80	0.85	0.88	0.71	0.74	0.83								
tmp.min.asy			28,478.01	29,506.27	27,805.39	26,775.94	25,771.47	23,638.32		1.02	1.08	1.08	0.87	0.88	0.88								
veg.ent			16,385.70	15,957.21	14,703.94	16,149.20	16,021.87	15,063.03		0.38	0.40	0.41	0.52	0.56	0.53								
vrn.med			14,369.22	12,958.12	12,884.54	14,374.72	12,673.35	11,885.61		0.37	0.40	0.42	0.41	0.41	0.42								
x																							
y																							

Table S3: Summary results of all phenological asynchrony random forest models Results include variable importance (both SHAP- and permutation-based; left and center columns) and overall model performance (R^2 and mean squared error (MSE)); right column), and are shown for models either including or excluding geographic coordinates as covariates (rows) and for models using both phenology metrics and using all three neighborhood radii (columns and sub-columns within each sub-table).

Supplemental Figures

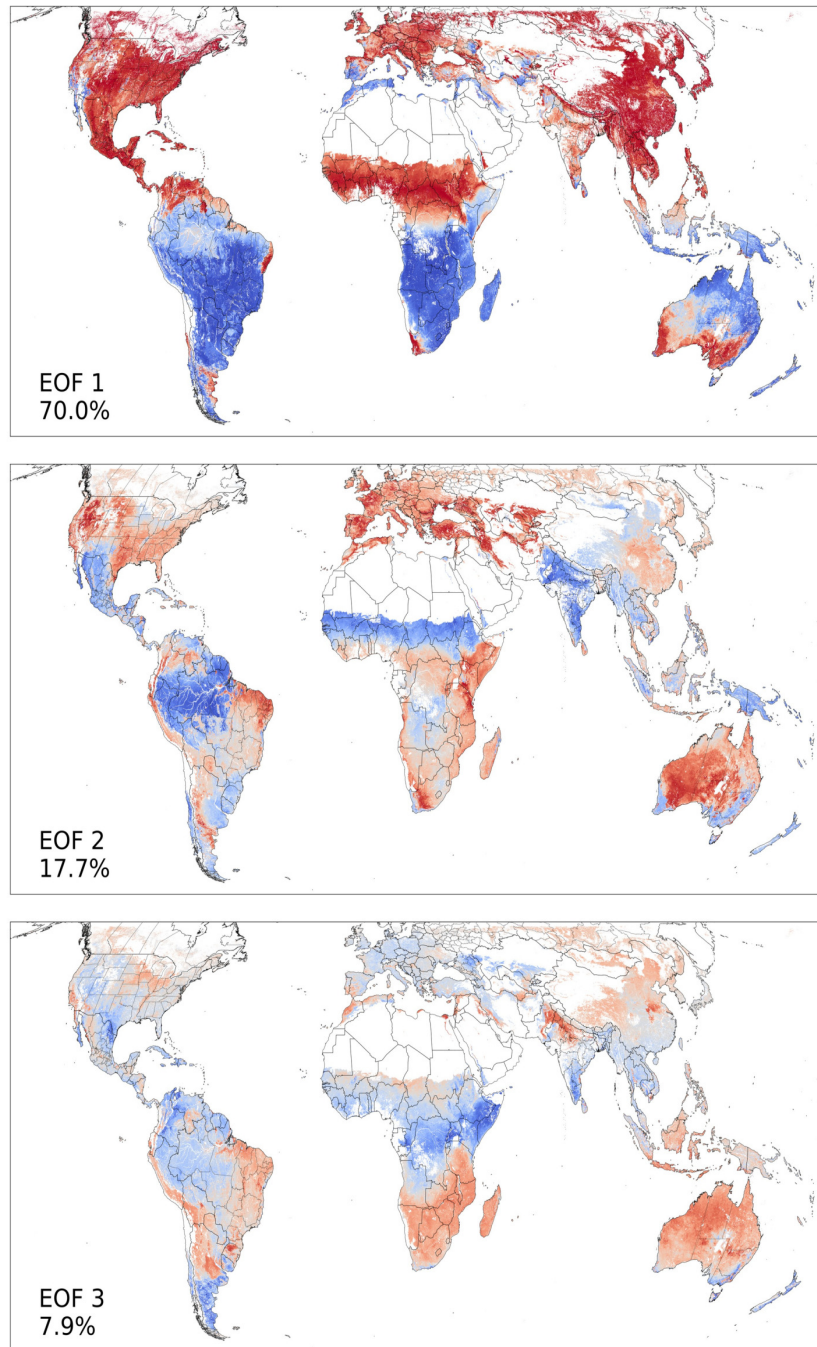


Fig S1: Raw EOF maps Each map shows one of the three predominant modes of characteristic annual LSP spatiotemporal variability resulting from our EOF analysis. EOF values are standardized and centered on zero, and maps are ordered by decreasing percent total variance explained, from top to bottom.

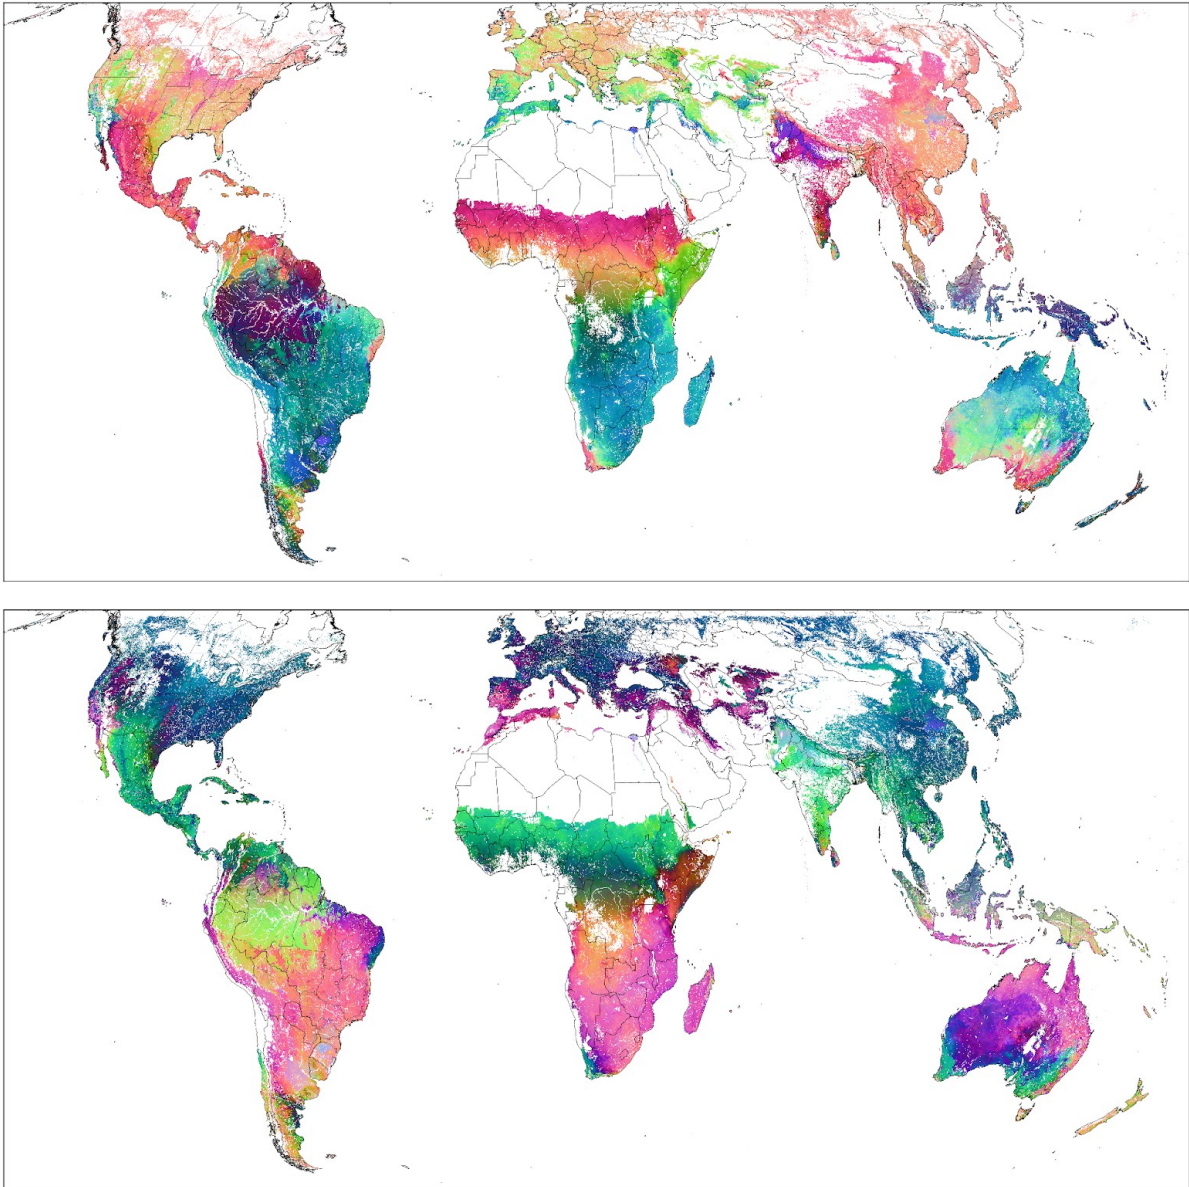


Fig S2: Untransformed RGB LSP maps Above: The raw values of all three top EOF modes depicted as red, green, and blue. Below: The three top modes depicted as red, green, and blue, but with modes 1 and 2 transformed to 1-mode (as is done in Figure 1 for the portions of the map south of the ITCZ).

validation test sites, plotted within orbital gap areas

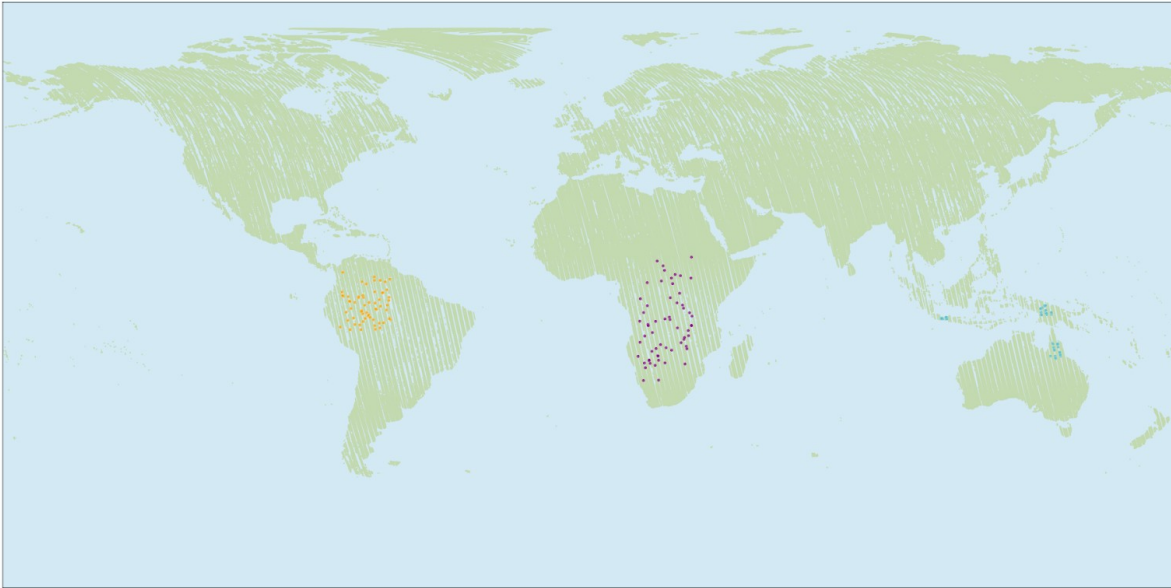


Figure S3: Orbital-gap validation points for SIF dataset: Map of random points chosen within Orbiting Carbon Observatory 2 (OCO-2) orbital gaps, used for the validation of our main SIF dataset against TROPospheric Monitoring Instrument (TROPOMI) SIF data (Figure S3).

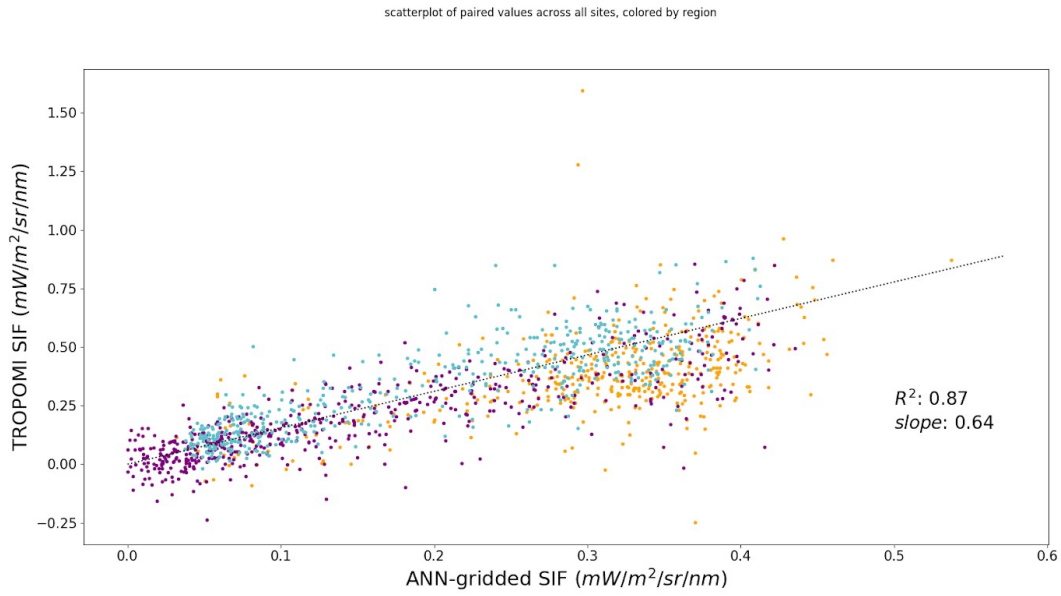


Figure S4: Validation of OCO-2 orbital-gap SIF time series against TROPOMI SIF time series For each point sampled in each of the three regions depicted in Figure S2 (orange: South America; purple: Africa; blue: Indo-Pacific and Australia), all contemporaneous estimates from the ANN-gridded SIF dataset used in our asynchrony maps and from the TROPOMI validation dataset are scattered. A simple linear regression depicts the level of agreement between these two independent sets of measurements.

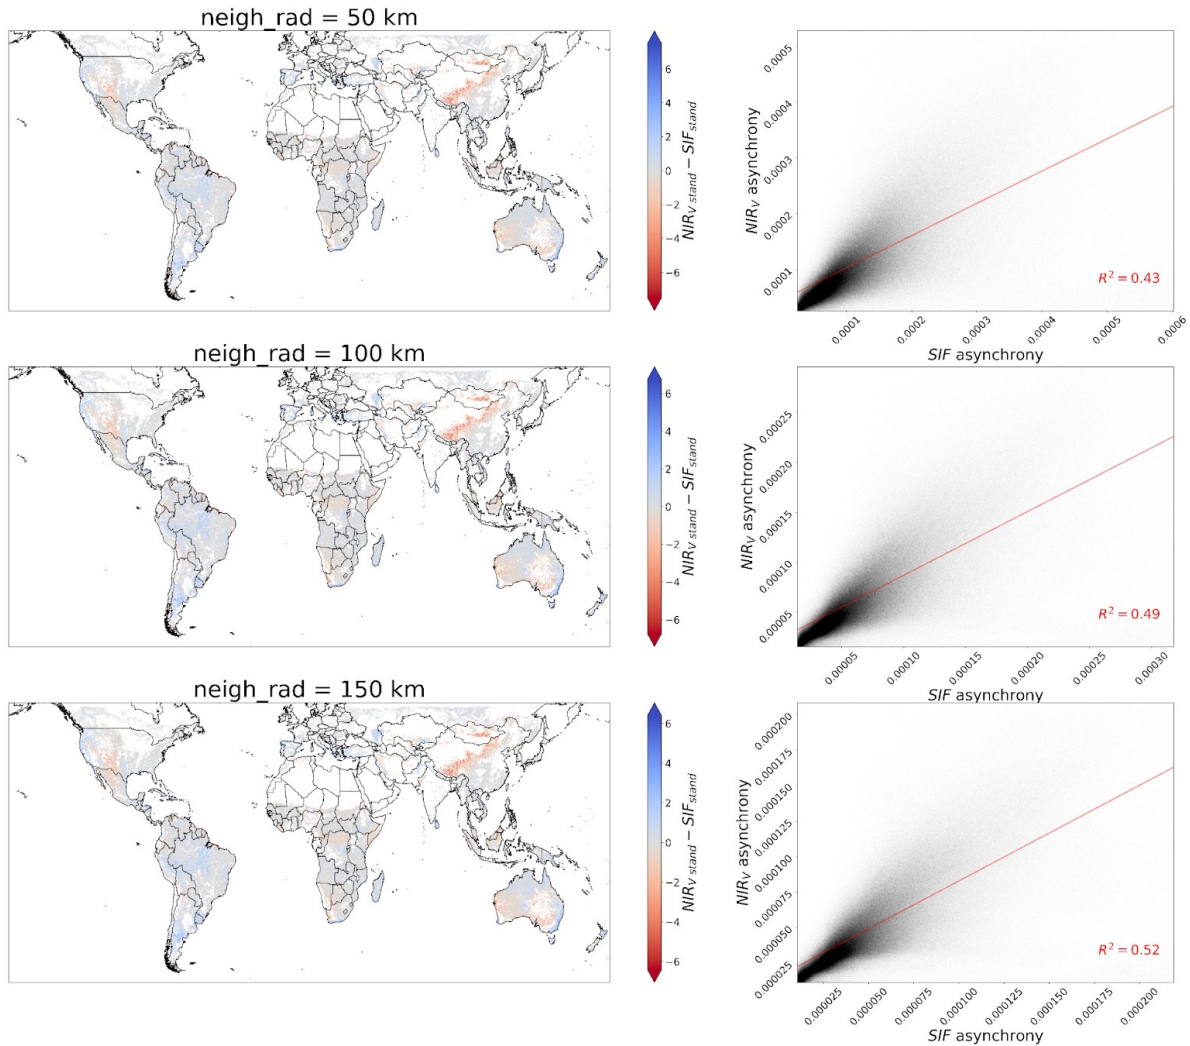


Figure S5: Correlation between NIR_v and SIF phenological asynchrony maps, across 3 neighborhood radii Maps show the raster difference between standardized NIR_v and SIF-derived phenological asynchrony maps for a given asynchrony neighborhood radius (top: 50 km; middle: 100 km; bottom: 150 km). For each neighborhood radius, original asynchrony map pixel values are scattered together at the right, and a simple linear regression depicts their relationship and correlation (R^2).

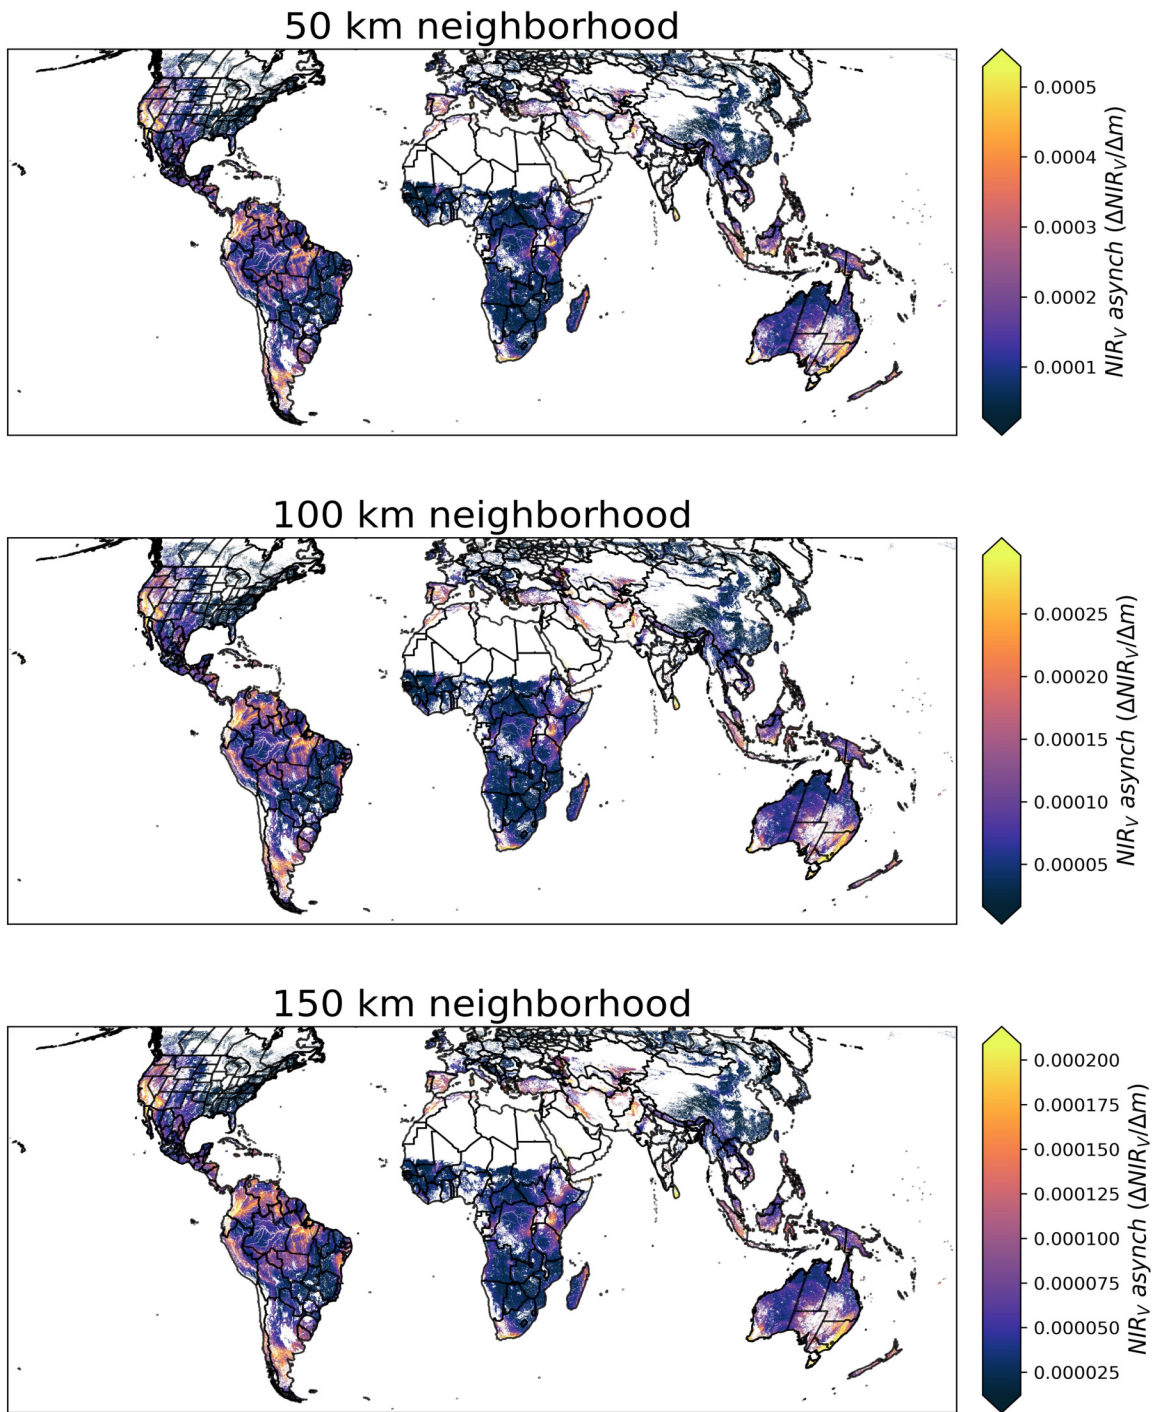


Figure S6: NIR_v asynchrony maps, for all 3 neighborhood radii

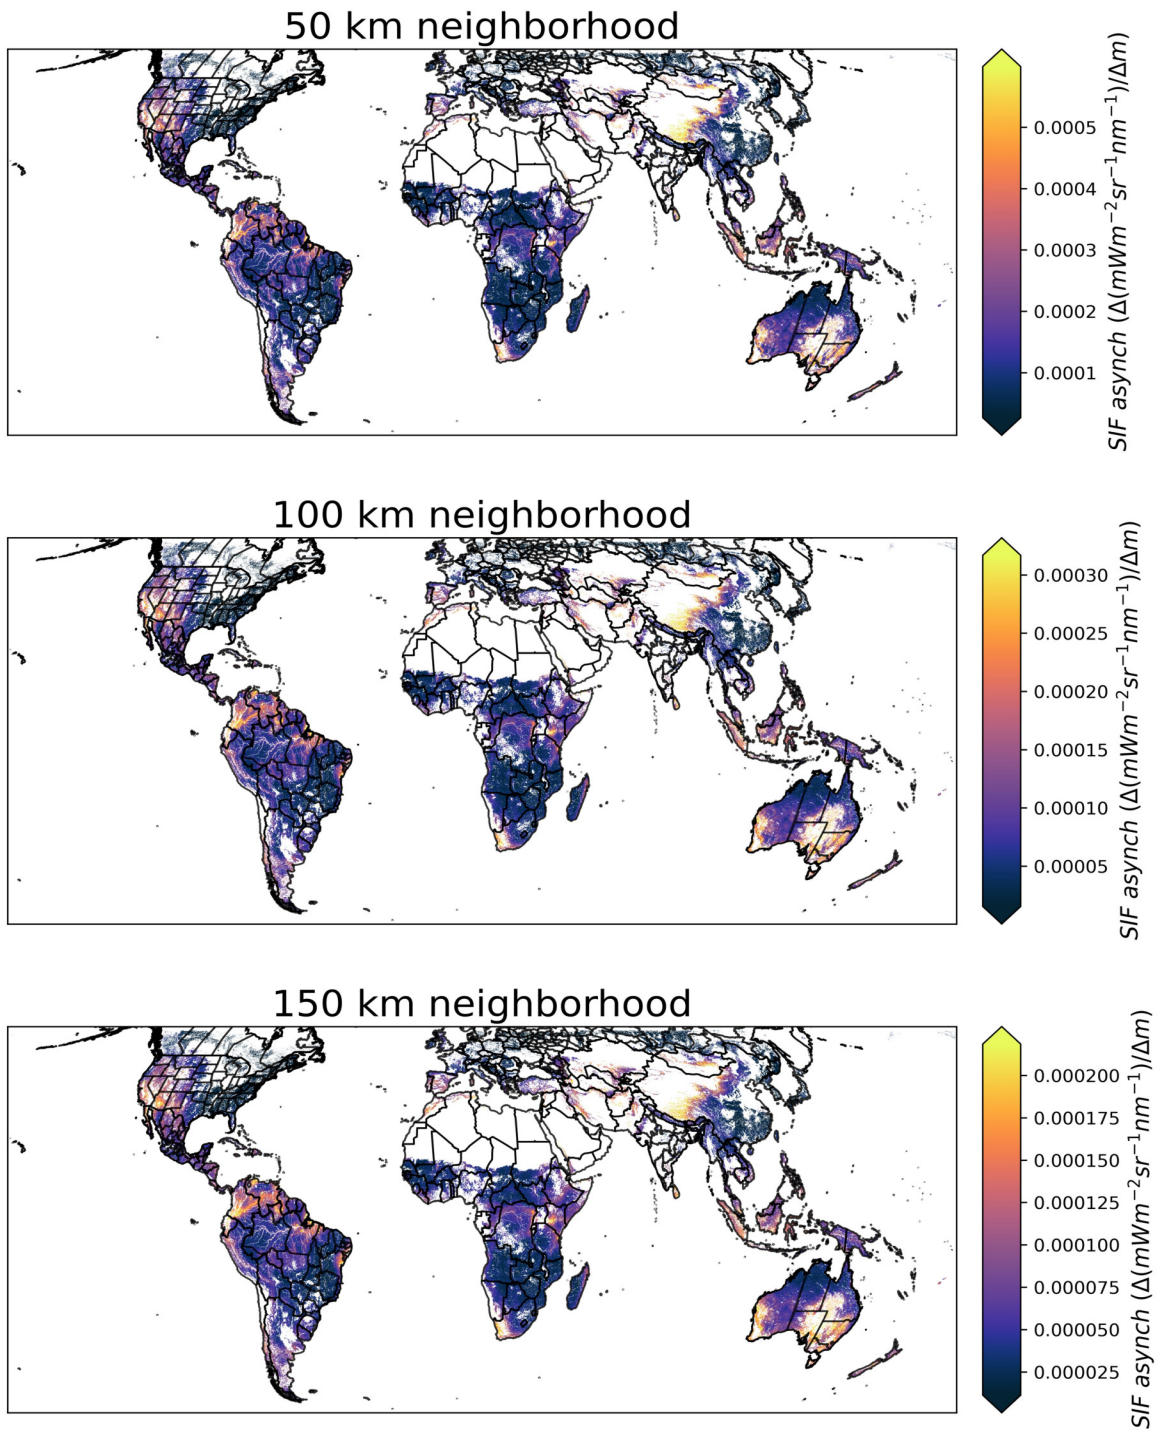


Figure S7: SIF asynchrony maps, for all 3 neighborhood radii

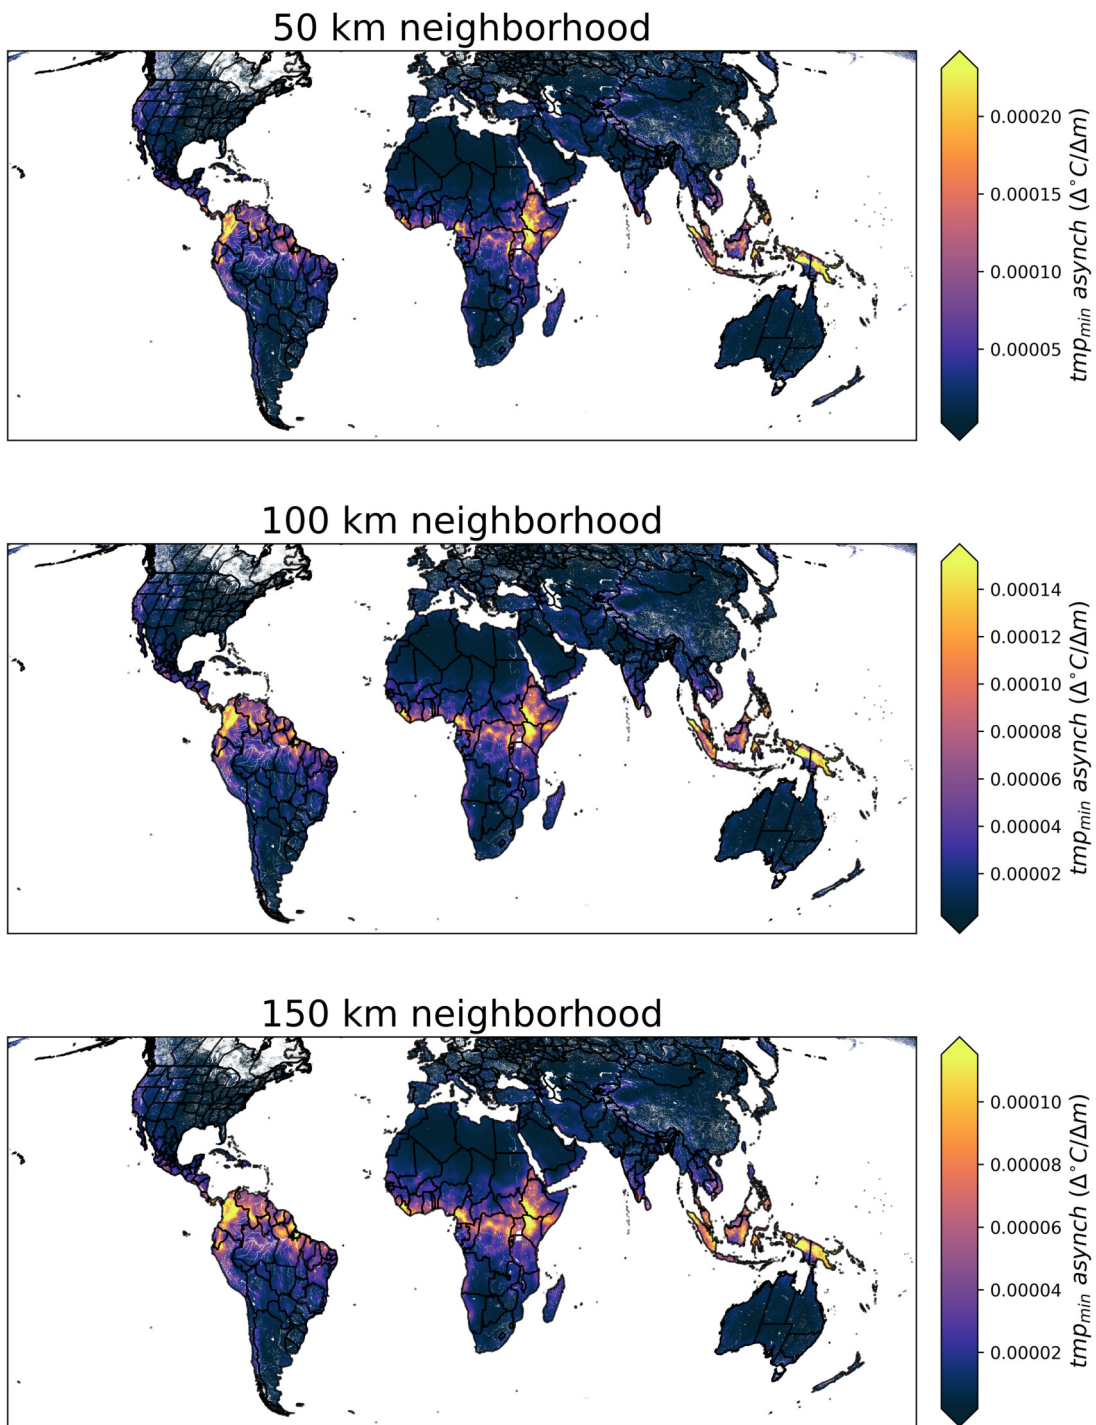


Figure S8: Minimum temperature asynchrony maps, for all 3 neighborhood radii

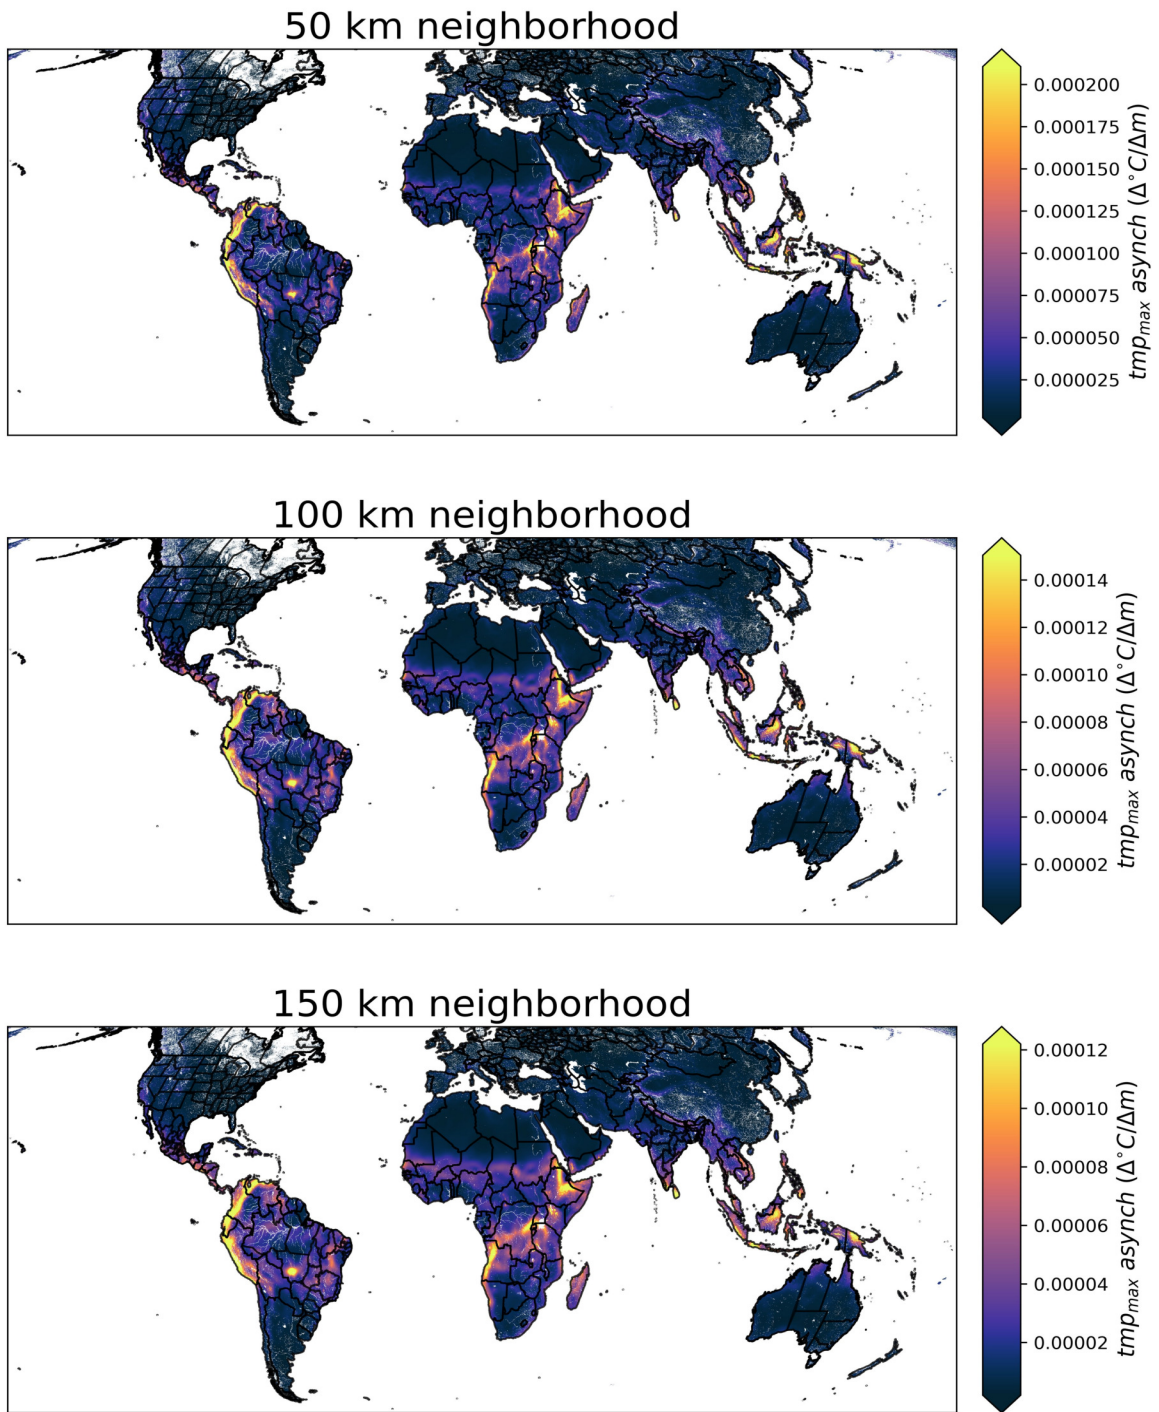


Figure S9: Maximum temperature asynchrony maps, for all 3 neighborhood radii

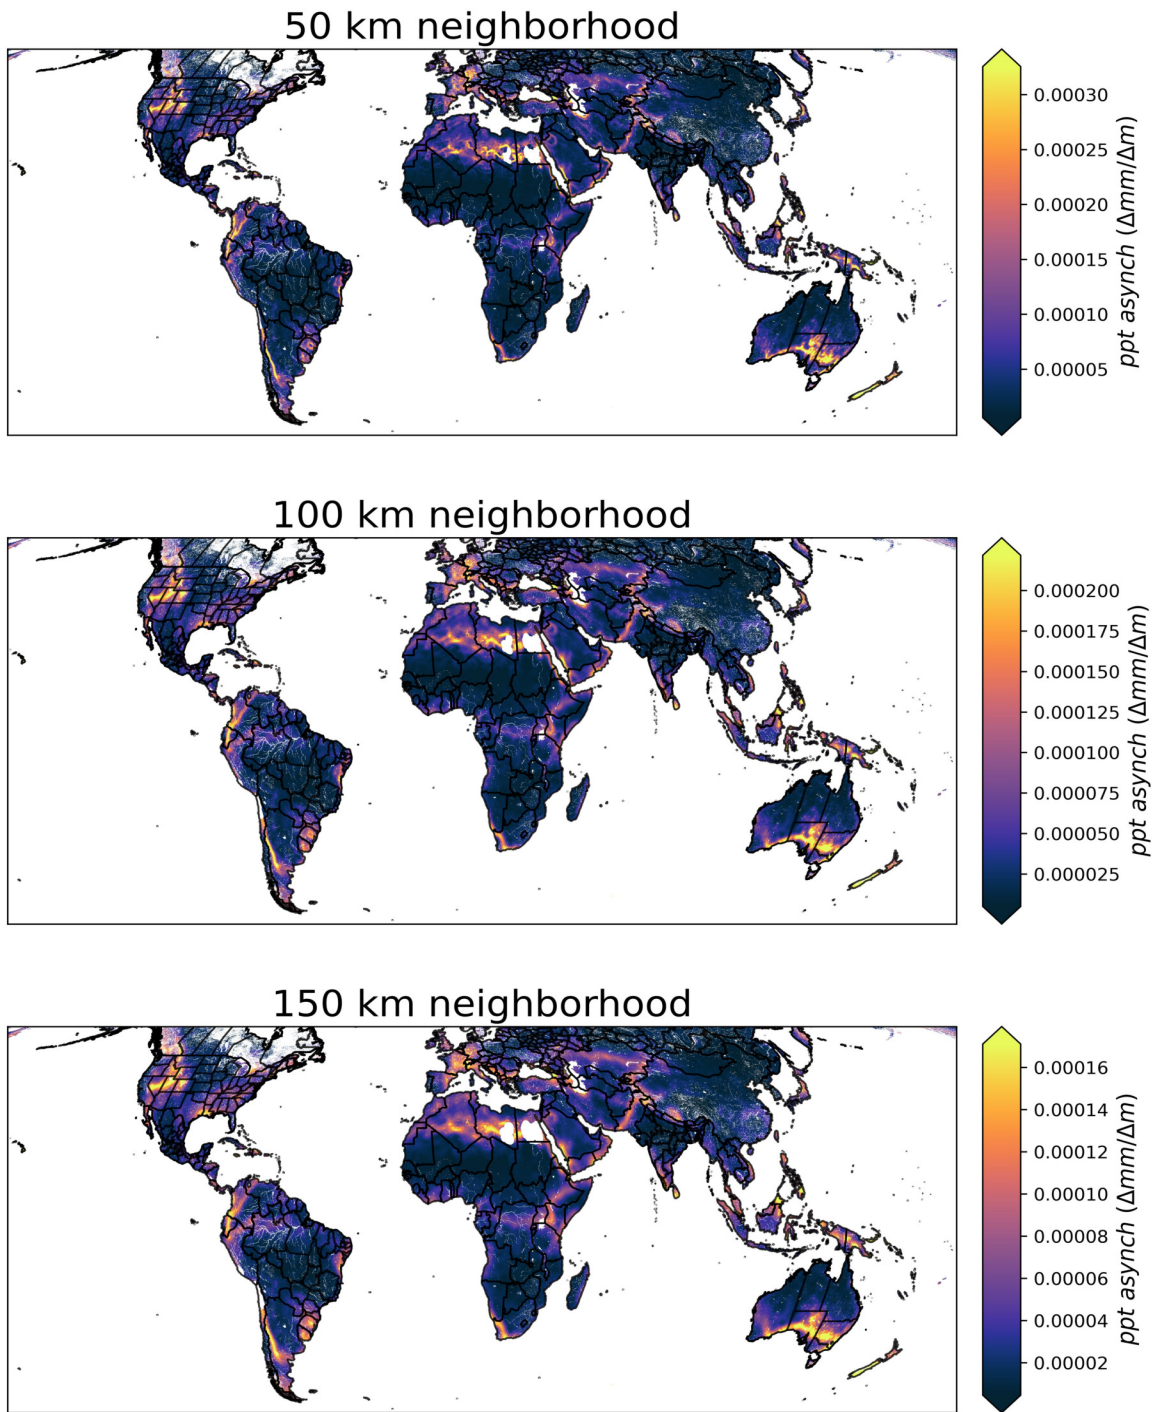


Figure S10: Precipitation asynchrony maps, for all 3 neighborhood radii

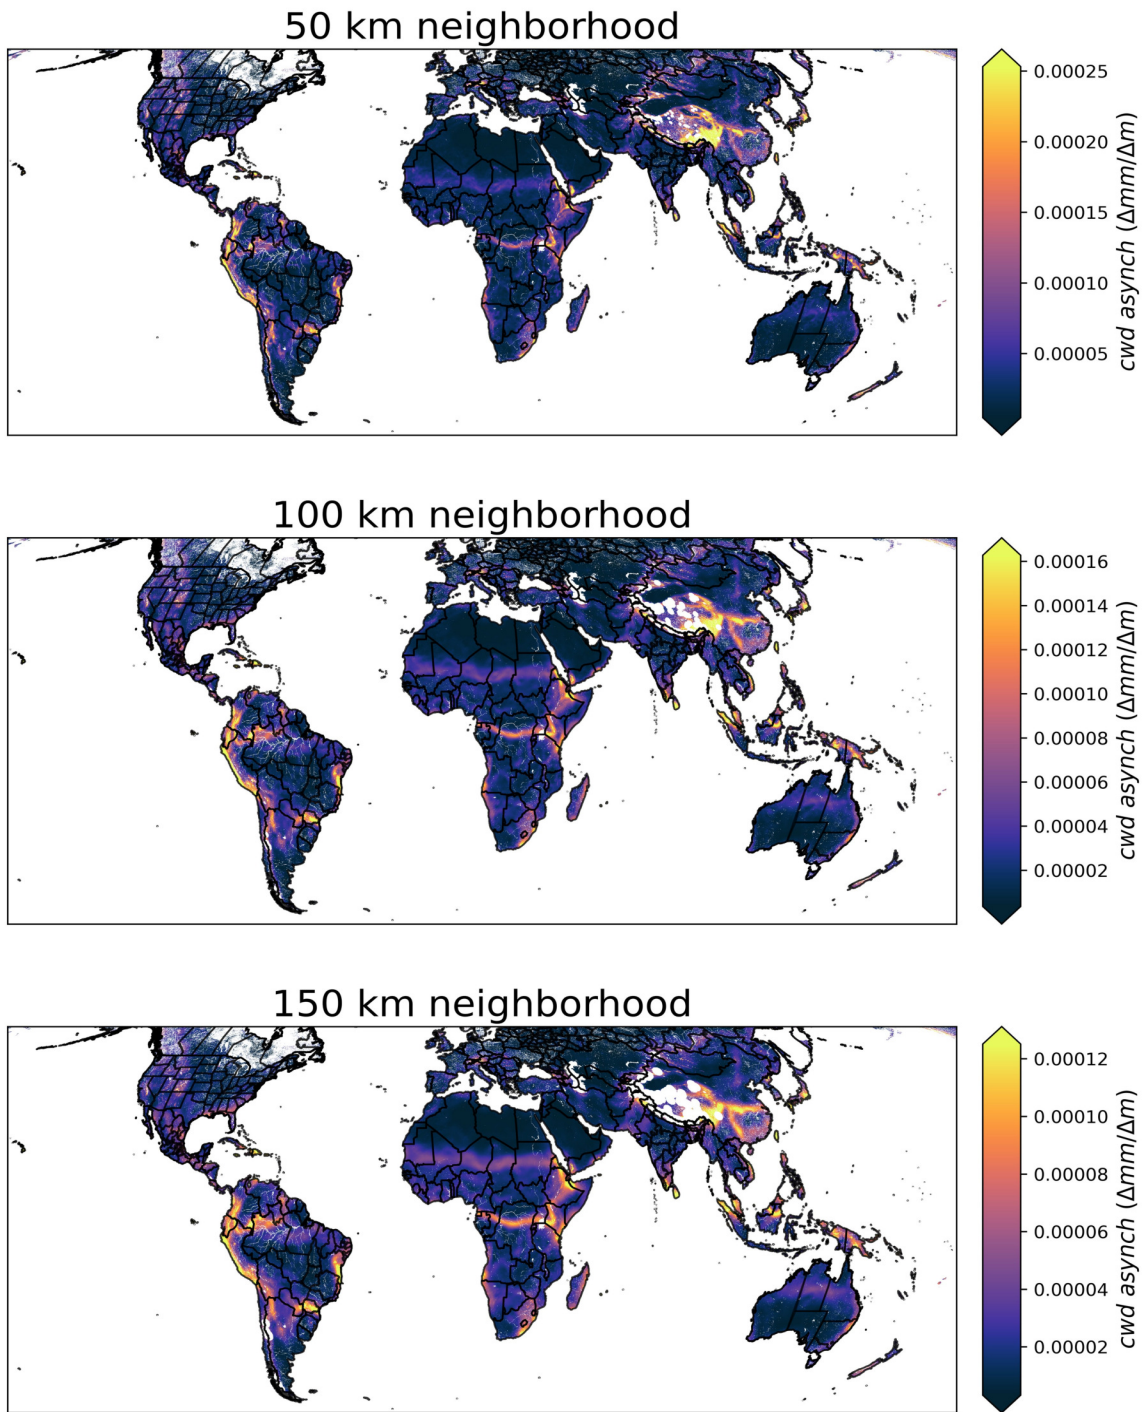


Figure S11: Climatic water deficit asynchrony maps, for all 3 neighborhood radii

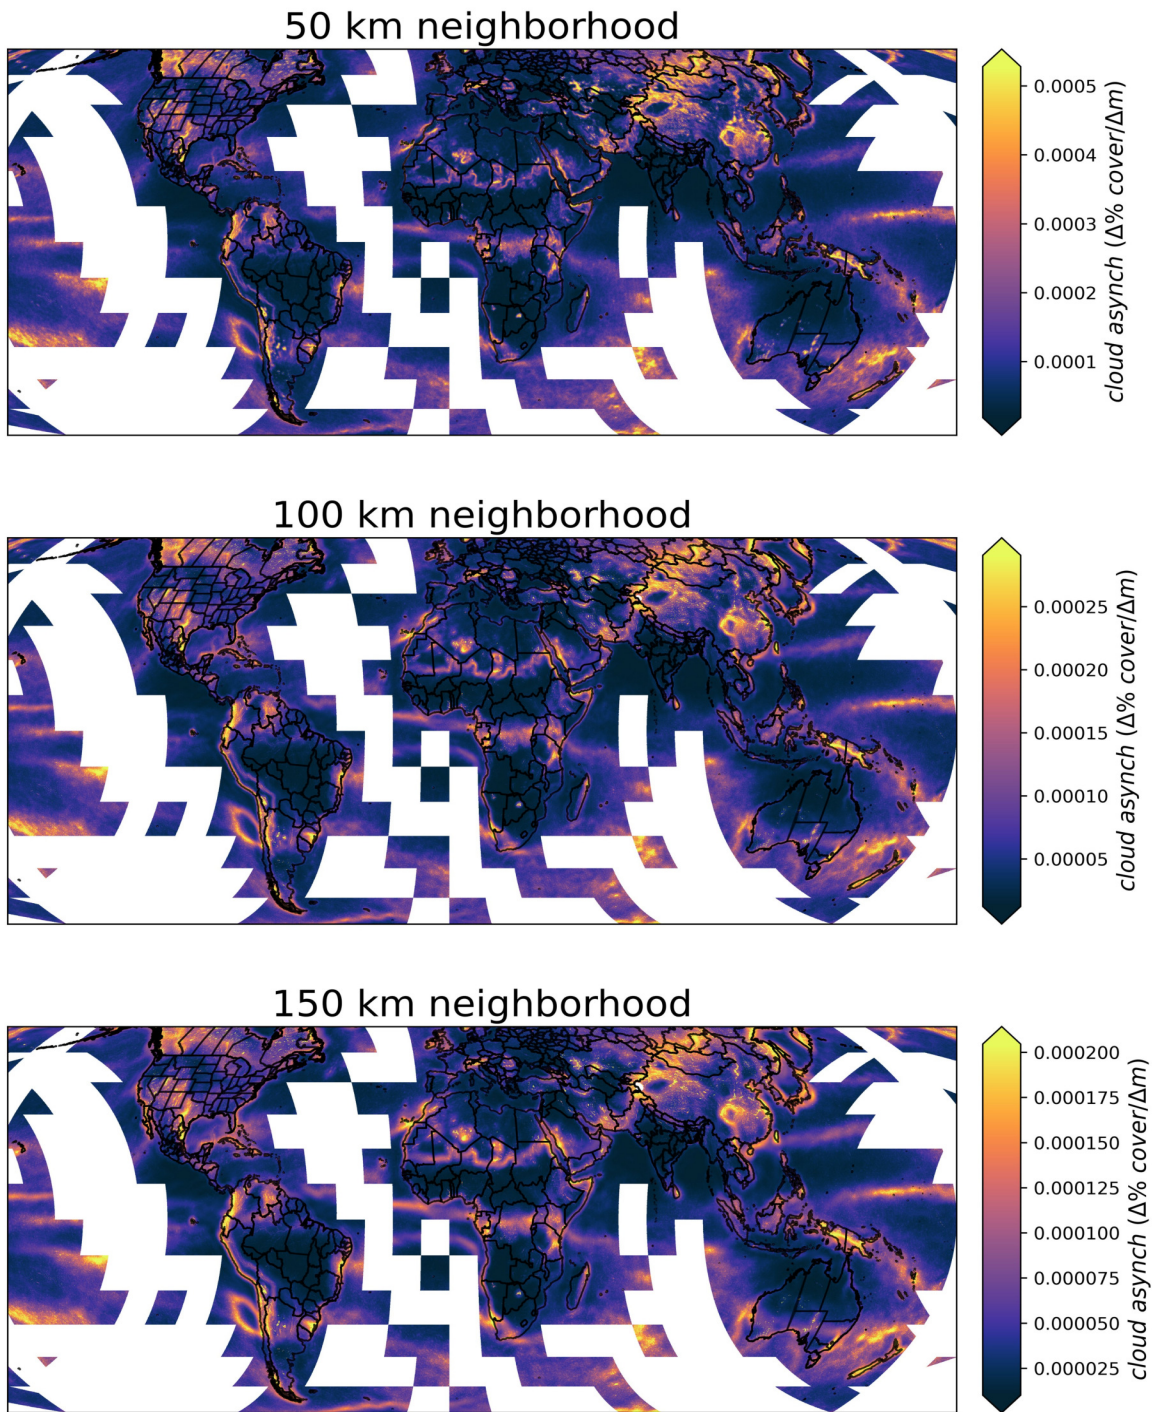


Figure S12: Cloud-cover asynchrony maps, for all 3 neighborhood radii

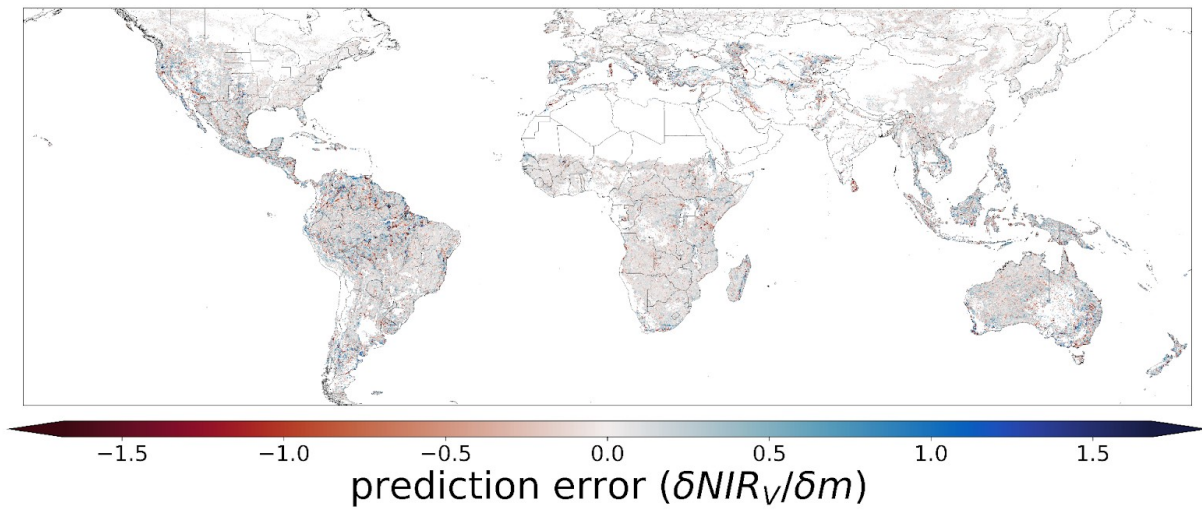


Figure S13: Map of standardized random forest prediction errors Values show the difference between the LSP asynchrony predicted by our random forest model and the actual LSP asynchrony value. Values shown are for our main model, using NIR_v-based LSP data and a 100 km neighborhood radius and including geographic coordinates as covariates.

Conclusion

In this final section I conclude the body of work detailed above. For each chapter, I first summarize key research findings, then articulate their importance, limitations, and implications for future work. It is my hope that the key points outlined below will serve as stepping stones, both for myself and for other researchers, in the endless effort to understand how evolution unfolds in the space and time of real-world environments.

In chapter 1, I produce and validate Geonomics, a robust Python scripting framework within which arbitrarily complex landscape genomic simulations can be built and explored. All validation tests are successful, demonstrating the Geonomics successfully recreates results expected under classical population genetics theory. And as the demonstration simulations in that chapter make clear, the broad capacity, flexibility, and ease of use of Geonomics can make it of great use to future theoretical and empirical research efforts. This contribution is important, as previously available comparable tools can be prohibitively complicated to learn and use, limiting the accessibility and usability of the state-of-the-art simulation resources that are crucial to progress in landscape genetics.

Geonomics, like all tools, nonetheless has its many limitations – it is a strong and convenient resource for the purposes for which it was designed, but would be a cumbersome choice of tool for purposes outside that. These outside purposes can be grouped into three categories: The first category is purposes for which less complex simulators are a better fit (e.g., population-based simulations would be best run in a more standard population-based framework). The second consists of purposes for which other complex simulators are a better fit (e.g., simulations requiring codon-specific genomic representation or allowing simulation of tandem-repeat markers should seek out simulators that integrate these biological facets). Third, however, is the category of purposes that are currently outside the purview of Geonomics but that have high potential for future expansion. At time of writing, examples of this third category about which interested researchers have contacted me include: simulation of a population's geographic expansion from a point introduction (currently feasible, but not supported by built-in functionality), calculation of both local and global effective population sizes (plausible, but would require development and validation of a methodology for spatial estimation of this normally aspatial statistic), and coevolutionary simulation of interacting species (would require substantial investment in the development and optimization of new modules of code, despite having been considered as a possible expansion since the early design of the Geonomics codebase). Ultimately, these requests for support, as well the recent publication of the first peer-reviewed use of Geonomics for empirical research (Maier *et al.* 2022), demonstrate the interest and utility that will help Geonomics find a life beyond the dissertation.

Chapter 2 not only provides the first theoretical study to rely on Geonomics, but also makes significant contributions to our understanding of evolutionary responses to global change. First, rather than using a simplistic notion of climate change as the uniform shift of climates across geographic space, I instead adopt an alternative landscape model of climate change as the basis for my simulations – one in which the environmental gradients subtending a species' local adaptation decouple as climate changes, causing the emergence of novel environments and consequent shifts in the adaptive landscape experienced by local subpopulations. This is a significant conceptual contribution, given that this scenario is rarely used as the basis for

evolutionary simulations, yet is a closer match to the consensus understanding of the landscape effects of climate change.

Second, when I employ this model across a suite of simulation scenarios of varying genomic architecture I produce results that offer novel and valuable insights, at two levels. At a detailed level, my analysis suggests particular areas of genomic architecture space in which adaptation to climate change is more or less likely to proceed – most notably, I find that architectures characterized by high polygenicity and low genotypic redundancy are likely to have especially low adaptive capacity, increasing the risk of local maladaptation and extinction. These specific findings are compelling but they are, of course, quite limited: No one simulation study could make hard claims as to the precise and generalizable influence of specific components of genomic architecture on all evolutionary outcomes under climate change. Rather, I hope that a growing body of work can: 1.) demonstrate the replicability of the findings I present here; 2.) expand them to incorporate other factors not considered here (e.g., population size, environmental complexity, distribution of allelic effect sizes) and characterize their degree of sensitivity to these factors; and 3.) validate these and future findings within empirical systems. This stream of research effort would help move the field toward a more generalized understanding of the evolutionary effects of environmental change. Nonetheless, at a high level, this study affirms that genomic architecture has the potential to influence, and even predominantly control, evolutionary dynamics and outcomes under climate change. This, in and of itself, flags this as an important but underappreciated area of future work for landscape geneticists.

The results presented in chapter 3 help highlight another underappreciated area of interest for landscape genetics: The potential for spatially asynchronous phenology to cause allochronic isolation, heightened population divergence, and perhaps even elevated speciation rates, as proposed by the Asynchrony of Seasons Hypothesis (Martin *et al.* 2009). However, despite the ASH having been introduced over a decade ago and having been the motivation for a few species-specific studies since then, some of its main implications – that seasonal asynchrony leads to phenological asynchrony predominantly in the tropics, and that seasonal asynchrony can be decoupled from climatic difference there – have gone largely unassessed. My work in chapter 3 inspects these implications and finds broad support: I find that global hotspots of phenological asynchrony are concentrated in tropical montane regions, but also in Mediterranean and neighboring semi-arid climate regions; that phenological asynchrony appears to be driven by precipitation asynchrony and minimum temperature asynchrony across many of those regions; and that phenological asynchrony between climatically similar sites tends to be greater on average in the tropics than at higher latitudes, though this varies regionally. These results provide, to my knowledge, the first global analysis of phenological asynchrony, and they not only help highlight phenological allochrony as a promising area for future landscape genetics research but they also provide a global, moderate-resolution (0.05°) map that can help identify locations for future field-based work.

On the way to documenting the aforementioned results in chapter 3, my work also incidentally produces a global map of land surface phenology (LSP) using a novel, multivariate LSP-mapping approach. I provide and interpret this map, showing that it uncovers a wealth of compelling and ecologically coherent patterns, including both regional heterogeneity that appears to reflect well-known environmental gradients, and intercontinental convergence in some

patterns of regional heterogeneity that recapitulates broad biogeographic and climatic patterns. My results also provide useful, globally-consistent performance assessment for the two LSP datasets I use (near-infrared reflectance of vegetation, or NIR_v; and sun-induced chlorophyll fluorescence, or SIF). Both of these are recently-derived remote sensing proxies of terrestrial productivity and are in a period of active development and improvement. Thus, my NIR_v-SIF cross-validation and my validation of both datasets against global FLUXNET eddy-covariance GPP estimates make useful contributions to the NIR_v and SIF remote sensing literature. However, although my results identify clear regions of reduced NIR_v-SIF correlation and NIR_v- or SIF-GPP correlation – namely, global semi-arid regions – they are not conclusive as to whether poorer validation in those regions is a result of the inadequacy of these metrics as proxies for photosynthesis, of the inadequacy of the resolution of my analysis as it compares to the finer spatial scale of flux-tower data, or of some other cause. This is a major limitation of this study, and it calls for future work, likely at higher resolution in select focal regions or localities, and paired with field-sampling, to fully elucidate the relationship and divergence between these remote sensing proxies and actual photosynthetic activity on the ground.

References

- Maier, Paul A., Amy G. Vandergast, and Andrew J. Bohonak. "Using landscape genomics to delineate future adaptive potential for climate change in the Yosemite toad (*Anaxyrus canorus*)." *Evolutionary Applications* (2022).
- Martin, P., Bonier, F., Moore, I. & Tewksbury, J. (2009). Latitudinal variation in the asynchrony of seasons: implications for higher rates of population differentiation and speciation in the tropics. *Ideas in Ecology and Evolution*, 2. <https://doi.org/10.4033/iee.2009.2.3.n>