

Reviewed Preprint

v1 • July 1, 2026

Not revised

✉ For correspondence:

ard.louis@physics.ox.ac.uk

Competing interests: No

competing interests declared

Funding: See [page 23](#)

Reviewing editor: Arvind Murugan, University of Chicago, United States

© 2026, Malone et al. This article is distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use and redistribution provided that the original author and source are credited.

Developmental bias explains the evolutionary trend towards simple leaf shapes

James S Malone^{1,2}, Nora S Martin³, Samuel HA von der Dunk¹, Liliana M Dávalos⁴, Ard A Louis¹ ✉

¹Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford, United Kingdom • ²Department of Biology, University of Oxford, Oxford, United Kingdom • ³CRG (Barcelona Collaboratorium for Modelling and Predictive Biology), Barcelona, Spain • ⁴Department of Ecology and Evolution and Consortium for Inter-Disciplinary Environmental Research, Stony Brook University, Stony Brook, United States

eLife Assessment

This paper presents **important** findings on how the shapes of leaves might be biased towards simpler shapes due to biases in how variation is generated by developmental processes rather than selection. The authors present **solid** evidence that combines image analysis of a herbarium dataset and computational analysis of a model of leaf development. The paper should be of interest to diverse researchers, ranging from plant development to the evolution of complexity more broadly.

<https://doi.org/10.7554/eLife.110692.1.sa3>

Abstract

The relative influence of developmental bias and natural selection on evolutionary outcomes remains a central, unresolved question in evolutionary biology. Here, we combine large-scale phylogenetic trees with herbarium data to confirm that, across the evolutionary history of angiosperms, transitions from complex to simple leaf shapes have occurred at a significantly higher rate than the reverse. To investigate whether this directional trend reflects an underlying developmental bias, we simulate evolutionary transitions by introducing random mutations into a well-established computational model of leaf development. Despite excluding selection, the transition rates produced by the model closely align with those observed in the phylogenetic data. This agreement indicates that a strong developmental bias, reflecting the relative ease with which gene regulatory networks generate simple over complex leaves, has played a major role in shaping the evolution of leaf morphology. Our findings echo recent work showing that similar biases towards simplicity influence the evolution of molecular structures, including RNA secondary structures and protein quaternary assemblies. Together, these results suggest that a developmental bias favouring simplicity may be a widespread and general feature of evolutionary systems.

I. Introduction

Much of the phenotypic diversity across the tree of life can be explained by adaptation to ecological demands via natural selection. However, this process is constrained by the mechanisms that generate phenotypic variation, including development [1, 2], which translates genetic information into form. While this idea is widely acknowledged in evolutionary developmental biology [1–3], only recently have advances in methodology and our understanding of development allowed researchers to examine its predictive power for evolutionary outcomes [4–11]. Yet it remains unclear how widespread and strong this developmental bias is across phylogeny [12, 13], particularly for complex traits like morphology. Here, we test whether developmental bias can account for macroevolutionary patterns in angiosperm leaf shape evolution.

Angiosperm leaves are an attractive system to study developmental bias. They display vast morphological diversity, much of which can be effectively represented in only two spatial dimensions due to minimal growth along the adaxial-abaxial (upper-lower surface) axis [14]. This diversity arises from variations in growth patterns regulated by gene regulatory networks (GRNs) that orchestrate both local and global growth dynamics during leaf development, allowing complex forms to emerge from anisotropic expansion. These processes include deeply conserved genes and developmental mechanisms [15]. Variation in leaf morphology is driven not so much by changes in gene content, but rather by differences in the spatial and temporal expression of shared developmental regulators. This helps explain why closely related species can display wide variation in leaf shape. For instance, the genus *Pelargonium* includes species with unlobed, lobed, dissected, and compound leaves [16]. Similarly, significant variation can occur within a single organism, as seen in *Morus alba* (mulberry), which produces both unlobed and lobed leaves, often with varying degrees of lobe symmetry [14, 17]. The conservation of basic growth mechanisms also suggests that some developmental biases will be, to first order, shared across many different angiosperm species, enabling large-scale studies of the role of development in determining evolutionary outcomes.

At the same time, there is evidence that leaf shape can be influenced by ecological demands such as thermoregulation [18, 19], hydraulic efficiency [20], light interception [21–23] and herbivory [24, 25]. However, it is still unclear how strong and general these effects are in natural environments [26, 27]. Size is often found as a stronger predictor of ecology than shape *per se* [27–30] (see section S3 for further discussion).

In a notable study, Geeta *et al.* [36] analysed a phylogeny of 560 angiosperms, showing that most leaf shapes (76% in their sample - table II [36]) are unlobed, and that transition rates between leaf shapes are highly asymmetric. Specifically, transitions from more complex forms - such as lobed, dissected, or compound leaves - to simpler unlobed forms occurred far more frequently than transitions in the opposite direction. What explains this global bias towards simpler leaves? Is it primarily development or is it primarily adaptation?

In order to answer this question, we proceed with two parallel lines of work. In the first, we extend the phylogenetic analysis of Geeta *et al.* [36] with two much larger, recent phylogenies [37, 38] and our own shape classification dataset made from 8,911 herbarium images [31]. These confirm that transitions to unlobed leaves are much more likely than transitions away from unlobed leaves.

In the second line of work, we examine the role of random mutations using a well-established computational model of angiosperm leaf development by Runions *et al.* [14] (fig. 5 [14]). This model captures key conserved developmental mechanisms thought to be broadly shared across angiosperms (Methods section IV A), allowing us to explore the generic effects of neutral evolution in the absence of direct selection on leaf shape. Remarkably, we find that random mutation alone is sufficient to reproduce the relative phylogenetic rates toward unlobed leaf forms. This suggests that developmental bias may be the first-order explanation of this macroevolutionary trend.

The basic mechanism driving this bias towards simplicity is straightforward [8]: complex leaf shapes need more specificity in their GRNs than the simpler unlobed leaf shape [40]. As a result, random genetic variation is more likely to produce simple shapes rather than complex ones that need more finely tuned parameter combinations.

II. Results

A. Transition rates from simple leaves to complex leaves are greater than vice-versa across angiosperm phylogeny

To establish whether there is a phylogenetic bias towards simple, unlobed leaves, we inferred the rate of evolutionary transitions between different leaf shapes across angiosperm phylogeny. We first classified 8,911 herbarium images (fig. 1a [31]) into 4 discrete categories - unlobed, lobed, dissected and compound. We used these data, as well as another shape classification dataset from

Geeta *et al.* [36], to label the tips of 5 phylogenetic trees representing angiosperm evolution [36–38]. Of these trees, 3 were resolved at the level of genus (Zuntini *et al.* [38] genus, Janssens *et al.* [37] genus and Geeta *et al.* [36]). The other two were at the level of species (Zuntini *et al.* [38] species and Janssens *et al.* [37] species) enabling us to examine how altering the amount of species-level diversity affects inferred transition rates. In agreement with Geeta *et al.* [36], we found that for these new phylogenies the vast majority have unlobed leaves (table II), already suggesting a strong bias towards simple, unlobed leaves.

To infer the rate of evolution between each of these 4 categories while accounting for tree topology, we fitted a continuous-time Markov chain (CTMC) model to the labelled trees (Methods section IV E), using a Markov chain Monte Carlo (MCMC) approach implemented in BayesTraits V4.1.2 [41]. In agreement with Geeta *et al.* [36], we detected a significantly higher rate of transitions to the simplest shape (unlobed) compared to transitions away to more complex shapes (figs. 1a and 7a).

The bias towards unlobed shapes was consistent across the five phylogenetic trees (figs. 7a and 7b). We found no cases where the direction of two statistically significant (black arrows - where 0 falls outside a 95% credible interval) net rates disagreed between phylogenies (fig. 7a).

B. A computational model of leaf development recapitulates *in-vivo* experiments demonstrating developmental bias towards unlobed leaf shape

To investigate how much of the preference for unlobed leaves can be explained by developmental bias, we use a popular model from Runions *et al.* [14] that combines gene-regulatory networks and key growth mechanisms for leaf shape (Methods section IV A). As a first validation test, we use the model to analyse *in-vivo* experiments by Kierzkowski *et al.* [40] that identified a bias towards simpler leaf shapes in *Arabidopsis thaliana* and *Cardamine hirsuta*.

In particular, this study found that expression of the genes *SHOOTMERISTEMLESS* (*STM*) and *REDUCED COMPLEXITY* (*RCO*) together was necessary to generate the compound shape in *C. hirsuta* and sufficient to induce compound shape in *A. thaliana* (fig. 2a). Therefore, in this system, compound shape requires a more specific combination of gene expression than simpler shapes (e.g. lobed, unlobed).

While the Runions *et al.* [14] model was not designed to model this system specifically, we tested whether it could nonetheless reproduce the result from Kierzkowski *et al.* [40] by selecting two model parameters that control similar growth properties to the genes *STM* and *RCO*, total growth duration and morphogen repression strength respectively (fig. 2). In agreement with the experiments, we found that activating both parameters together was sufficient to lead to a compound shape while activating each in isolation was not (figs. 2a and 2b).

Having reproduced the experimental phenotypes, we next used the model to explore the *STM*/*RCO* parameter space in more detail, systematically varying morphogen repression strength and duration of growth within their biologically plausible ranges (fig. 2c). We found that the majority of this space (20 of 25 parameter combinations) was occupied by simple, unlobed shapes, with only a small fraction representing more complex shapes (1/25 dissected, 4/25 compound - see Methods section IV D). This agrees with Kierzkowski *et al.* [40] who also found that a similar fine-tuning is required for compound leaves.

C. Mutation in a computational model of leaf development is sufficient to reproduce phylogenetic bias towards simple leaves

The mechanism derived from the experiments on varying *STM*/*RCO*, shown in fig. 2, is one specific driver of developmental bias. To investigate a wider range of possibilities, we study the effect of random mutations for a broad set of Runions *et al.* [14] model parameters (table S1). We first generated a sample of initial leaves representing a diverse morphospace (figs. S1 and 1b). We mutated the initial sample using a random walk algorithm with two different mutational schemes,

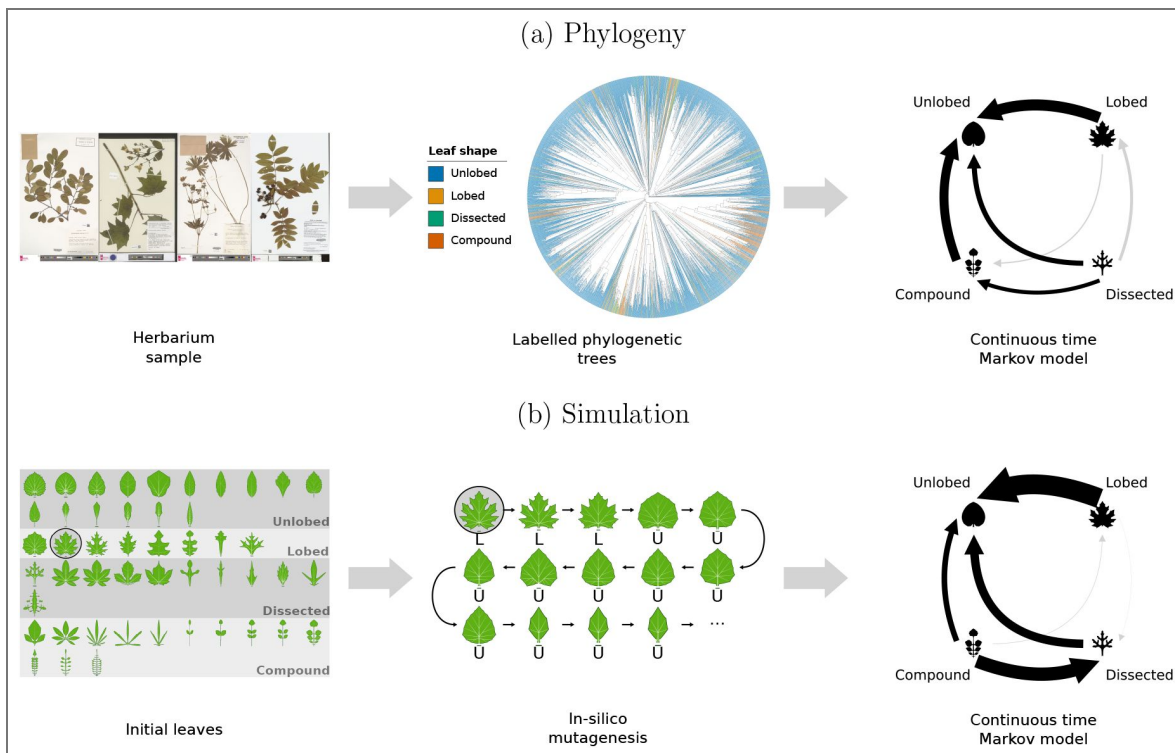


FIG. 1. A summary of the methods and main results.

We inferred transition rates between leaf shapes in two different ways: **(a)** by fitting a continuous time Markov chain (CTMC) model of trait evolution to phylogenetic trees and **(b)** by simulating mutation in a computational model of leaf development. **(a)** (Left) A sample of images from the Naturalis Biodiversity Center [31][32–35] herbarium are manually classified as unlobed (u), lobed (l), dissected (d) or compound (c) (Methods section IV F). This dataset, as well as an analogous dataset by Geeta *et al.* [36], are used to label the tips of 5 different phylogenetic trees representing angiosperm evolution [36–38]. (Centre) the Zuntini *et al.* [38] species tree is shown in full. Rates are inferred by fitting a CTMC model to the labelled trees using a Markov chain Monte Carlo (MCMC) approach in BayesTraits V4.1.2 [39]. (Right) Posterior net rates for the Zuntini *et al.* [38] species tree are shown as arrows. Black (grey) arrows depict rates for which 0 falls outside (within) a 95% credible interval. **(b)** (Left) An initial sample of 48 leaves representing a wide morphospace is generated using the Runions *et al.* [14] model (Methods sections IV A and IV B). Leaves in this sample are mutated in parallel through a random walk algorithm, with no selection besides rejecting parameter combinations that give physically implausible leaves (table I [4]), generating a sequence of leaves that evolve away from the initial leaf (Methods section IV C). (Middle) Every leaf is classified as unlobed (u), lobed (l), dissected (d) or compound (c) using an automated morphometric approach (Methods section IV D), giving a sequence of character states for each walk. To this trajectory, a CTMC model is fitted using MCMC (Methods section IV E). (Right) Arrows depict posterior net rates for the MUT2 scheme (alg. 1 [4]).




FIG. 2. An experimentally motivated illustration showing that simple leaves are easier to generate than complex leaves.

(a) Leaf silhouettes of wild-type *A. thaliana* and three mutants, which result in lobed (+*STM*, +*RCO*) and compound shape (+*STM* /+*RCO*) (reproduced from Kierzkowski *et al.* [40] with permission). Ectopic *STM* and *RCO* expression individually create a lobed shape, but together induce compound leaflets. **(b)** A reproduction of the Kierzkowski *et al.* [40] result using the Runions *et al.* [14] computational model of leaf development. +*STM* is approximated by increasing the duration of growth and +*RCO* by increasing the growth repression strength of one morphogen on the margin of the leaf. Morphogen presence is represented by the purple/blue lines outside the margin. **(c)** A morphospace showing the resulting shapes from different levels of the two parameters we use to approximate *STM* and *RCO*. The range shown for *morphogen repression strength* is the range in which biologically plausible shapes can be produced. The range for *duration of growth* is the range in which leaves can be generated in a reasonable timeframe. In the parameter files of the Runions *et al.* [14] model, the parameters we refer to as *morphogen repression strength* and *duration of growth* are given by FAIRING4 and FINALFRAME respectively.

MUT2 (alg. 1 [↗](#)) and MUT5 (alg. S2), and automatically classified the resulting shapes into the same four shape categories (unlobed, lobed, dissected and compound) using a morphometric classifier (fig. 6 [↗](#)). Of the 76,800 total leaves generated per simulation, we found that the proportion of unlobed, lobed, dissected and compound were 73.4%, 9.6%, 11.4% and 5.6% respectively for MUT2 and 76.1%, 7.2%, 11.4% and 5.4% respectively for MUT5. This shows that our results are robust to the two different mutational schemes. Both reflect the fact that unlobed leaves occupy a significantly greater proportion of model parameter space that produces viable leaves than more complex shapes do (see also Results section II D).

Regardless of the initial leaf shape, over the course of these random walks, the proportion of walks in an unlobed state increased to an equilibrium just below 80% while most other shape categories stabilise to a significantly lower proportion (figs. 3 [↗](#) and 8 [↗](#)). These predictions lie within the margins of error of the phylogeny-derived CTMC model predictions (figs. 3 [↗](#) and 8 [↗](#)), suggesting that mutation and development alone are, to first order, sufficient to predict the relative frequency of different leaf shapes over evolutionary time.

To enable a more detailed comparison between simulation and phylogeny, we inferred the rates of shape transitions from our simulations. By estimating rates for all 12 possible (non-self) transitions, we gain deeper insight into the dynamics of shape change over time than by examining the proportions of the four shape categories alone. To do this, we fitted the same CTMC model used for the phylogenetic inference to the random walk data, to estimate the evolutionary rates between shape categories in the absence of everything but mutation, using MCMC [42, 43]. In agreement with our phylogenetic inference, we found that transitions from unlobed occurred at a significantly lower rate than transitions to unlobed, giving a consistent net transition rate towards unlobed (figs. 1b [↗](#) and 7a [↗](#)). This robust agreement between relative rates derived independently from simulations and phylogenies suggests that the evolutionary preference for simple leaves observed across angiosperm phylogeny may primarily be driven by developmental bias instead of selection.

By contrast, for transition rates between complex leaf shapes (lobed, dissected and compound) (figs. 1 [↗](#) and 7a [↗](#)) there is less agreement between both the different phylogenies, and between the phylogenies and the simulations. These differences could be because selection has played a more prominent role in the transitions among more complex leaf shapes than it has for transitions from complex to simple. The model may also be less reliable at resolving relative rates of complex leaves, for which more detail needs to be specified. Fully resolving the discrepancies for transitions between more complex leaf shapes would require reducing the error associated with the phylogenetic rate estimates, which in turn would necessitate significantly more leaf shape data.

D. Complex leaves occupy less volume in parameter space and are closer to simple leaves, but not vice versa

To further investigate the mechanisms behind the bias towards simple leaves, we explored – much as we did for the experimentally inspired study in fig. 2c [↗](#) – the relative volumes in parameter space that different shapes occupy. Projecting the parameter values of all random walk outputs into a 2D space via principle component analysis (PCA), reveals several possible factors affecting relative rates (fig. 4 [↗](#)). First, it shows that each of the four shape categories occupy a different sized area (fig. 4a [↗](#)). The simplest category, unlobed, occupies by far the greatest area (fig. 4b [↗](#)), suggesting that the bias towards simple forms in our simulations can be partially explained because, as argued before (fig. 2 [↗](#)), more parameter combinations map to simple leaves than complex leaves. We find that this pattern also holds for PC3-PC4 (fig. S16e [↗](#)) and PC5-PC6 (fig. S16f [↗](#)).

Moreover, the vast majority of lobed, dissected and compound leaves are contained within a broader area occupied by unlobed in the 2D parameter space (fig. 4a [↗](#)), suggesting that more complex leaves are close to simpler leaves in parameter space but not always vice versa. This

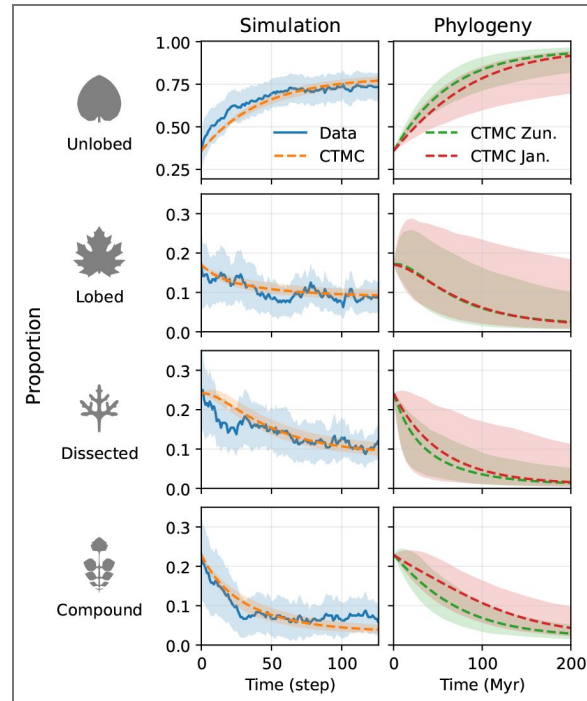


FIG. 3. The proportions of each leaf shape over time as predicted by *in-silico* mutagenesis and phylogeny.

(Simulation) The solid blue line represents the mean proportion of random walk chains in each state at each time step ($\pm 1.96 \times$ standard error) in our simulation (MUT2 alg. 1 [38]). To this data, we fit a CTMC via maximum likelihood estimation (MLE) (dashed line) and MCMC with errors inferred from MCMC posteriors via Monte-Carlo error propagation. Since step size is arbitrary, for comparison with the phylogenetic model, we adjust the x-limits for the simulation such that the CTMC predictions for both the simulation and phylogeny reach halfway from the initial distribution to the stationary distribution by approximately the same distance along the axis. **(Phylogeny)** We predict how the proportions of each leaf shape evolve from the same initial proportions used in the simulations, using the CTMC inferred from each phylogeny (green - Zuntini *et al.* [38] genus (fig. S3 [38]), red - Janssens *et al.* genus (fig. S5 [38])). Dashed lines represent the prediction using MLE transition rate values. Errors for all CTMC predictions were inferred from the transition rate MCMC posteriors using Monte-Carlo error propagation ($N = 500$), and represent a 95% credible interval for each time step. We show these extrapolations for a time period equal to the approximate age of the angiosperm clade (200 Myr) [38].

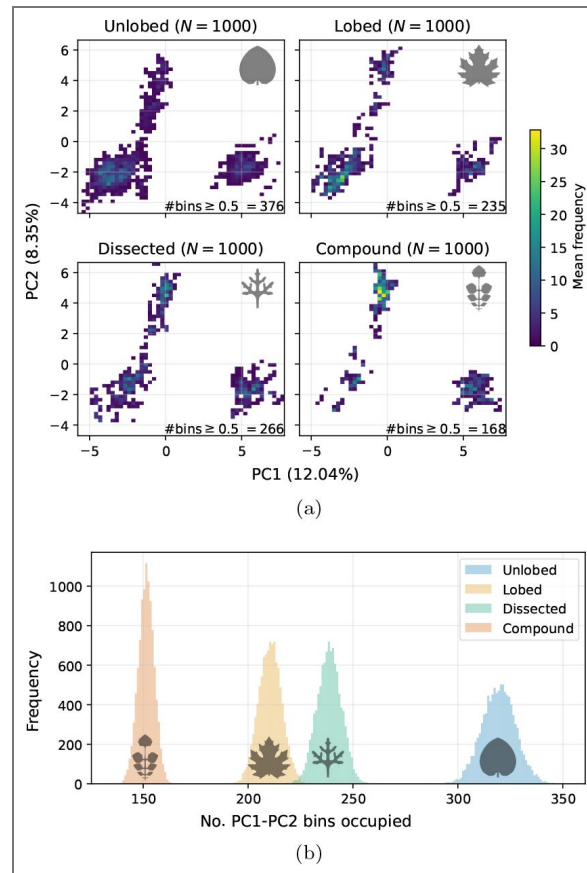


FIG. 4. Simpler leaves occupy a greater volume than complex leaves in model parameter space.

(a) A 2D histogram of model parameter space given by the 2 principle axes from a principle components analysis (PCA) of a sample of 4000 leaves from each shape category from one of the simulations (MUT2). Mean bin occupancy is calculated via a bootstrapping approach, by sub-sampling with replacement $N = 1000$ leaves from each shape category for 10000 iterations. Bins are shaded if the mean occupancy across these iterations is ≥ 0.5 . This reveals that unlobed leaves occupy a larger region of model parameter space than more complex shapes and that this larger space also contains the majority of more complex leaves. Moreover, more complex leaves are concentrated within a much smaller area than unlobed leaves. (b) The no. PC1-PC2 2D histogram bins occupied by each leaf shape in 10000 bootstrap samples ($N = 1000$) of a balanced sample of MUT2 (4000 per shape category). Bins with at least 1 leaf present are counted. This reveals that unlobed occupies by far the largest area in this 2D projection of parameter space. We find the same is true for other PCs (fig. S16).

FIG. 5. A description of the leaf model from Runions *et al.* [14], highlighting the three key components of leaf development and their interactions.

(a) Model components illustrated in leaf form (inspired by Runions *et al.* [14] fig. 3). Margin in black, veins in blue, lamina in green, convergence points in white with black outlines. (b) The three components of the model with arrows indicating interactions that affect growth. (1) Existing morphogens and convergence points influence where new ones will form. (2) Convergence points and morphogens influence local growth rate. (3) Convergence points induce new veins. (4) Veins determine local growth direction. (5) Existing veins influence the growth of new veins. (c) Zoomed view of the margin, showing how growth of the vasculature and patterning on the margin interact to produce protrusions and sinuses. Veins (in blue) extend, pushing the margin out, however a morphogen (in orange) inhibits webbing between the veins, causing a sinus to form. This increases the distance between the two convergence points, leading to the insertion of another convergence point between the two that excludes the morphogen from its vicinity. Existing vasculature connects to the convergence point and specifies a new axis of growth. (d) Different growth stages of a model leaf. The dark blue morphogen inhibits webbing and convergence point formation, the light blue morphogen delineates the petiole. The value of t denotes the iteration of the simulation.


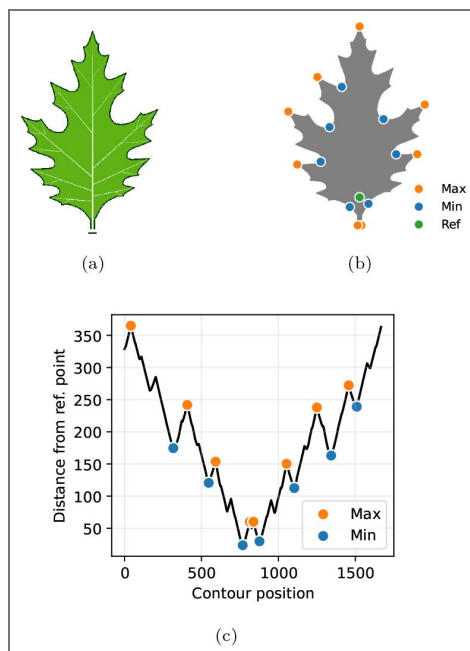


FIG. 6. A summary of the morphometric classifier.

(a) Example lobed leaf. (b) Silhouette with local maxima in orange, minima in blue and reference point in green. (c) Graph of distances of each contour point to the reference point. Local maxima shown in orange, minima in blue.



provides another mechanism making transitions from complex leaves to simple easier than the reverse. We also observe this same pattern for PC3-PC4 (fig. S16b) and PC5-PC6 (fig. S16c). Finally, the parameter space of unlobed leaves appears to be much more connected than the space of more complex leaves. Constructing a nearest neighbour graph (no. neighbours = 6) from the PC1-PC2 parameter space reveals that unlobed separates into far fewer connected components (3) than lobed (131), dissected (96) and compound (235) (fig. 9). This may further contribute to the bias in transition rates towards unlobed, as less connected spaces would be more frequently exited by random walks than more connected spaces.

Our results show that the finding from experiments [40] and manual parameter searches (fig. 2c)—that simpler leaves occupy a larger volume within the parameter space of the developmental programme—extends to the Runions *et al.* [14] model more generally. This volume disparity, in addition to other properties including the degree of overlap between shape categories and the connectedness, may account for much of the bias towards unlobed we see in our simulations,

III. Discussion

We set out to test the hypothesis that mutation alone was sufficient to produce the higher transition rates towards simple leaves compared to complex leaves observed in angiosperm phylogeny [36]. By performing *in-silico* mutagenesis, we found that non-adaptive processes alone produce greater transition rates to simple unlobed leaves than to complex leaf shapes such as lobed or compound (figs. 1b and 3). We also reproduced the observation from angiosperm phylogeny that unlobed leaves have evolved at a faster rate from complex leaves than vice versa [36], using larger and more recent trees [37, 38] and new classification data (figs. 1a and 3). Together, this suggests that a developmental bias towards simple, unlobed leaves can explain this macroevolutionary pattern in flowering plants.

Our finding is consistent with the growing body of evidence that developmental bias can play a significant role in evolutionary outcomes for morphological traits [4–1]. The particular bias towards simplicity is consistent with recent work arguing for a generic bias towards descriptiveness in genotype-phenotype maps [8]. Johnston *et al.* [8] provided examples of this bias strongly affecting RNA secondary structure, where simpler structures are much more likely to be observed in nature than complex ones (see also [44]) and for self-assembling protein quaternary structure, where this bias helps explain the prevalence of highly symmetric forms in the protein data bank (PDB). Furthermore, this bias has been predicted to hold more widely for input-output maps [45]. This generic argument implies that the simplicity bias observed in molecular phenotypes may extend to the level of macroscopic phenotypes, which we show here for leaf shape, but also potentially to a much wider set of morphological phenotypes. More specifically, our results suggest that this bias toward simplicity may hold generically when development uses conserved mechanisms and GRNs where complex phenotypes need tighter parameter specification than simpler phenotypes.

Our study also opens up a number of future directions for studying the evolution of leaf shape. First, it would be valuable to repeat similar simulations in alternative models of leaf development [40, 46, 47]. The Runions *et al.* [14] model is uniquely suited to macroevolutionary studies such as this one, as it captures key conserved elements of angiosperm leaf development while generating a large morphospace, but, if the same bias towards simplicity was observed in other models, this would provide further evidence that it is a general property of leaf development. Second, it would be illuminating to conduct similar *in-silico* mutagenesis, but classify the resulting leaves using a continuous classification system instead of discrete categories. For example directly measuring margin complexity using a compression algorithm on the leaf silhouette [48]. This could allow for more subtle biases between different leaf forms to be identified. Third, when inferring transition rates from phylogeny, we were limited by the availability of herbarium data, the resolution of phylogenetic trees and the rate at which herbarium leaves could be manually classified. To more accurately resolve transition rates between more complex shape categories, it would be useful to repeat this study with larger phylogenies. This would likely require an automated method for

herbarium shape classification. Finally, the methodology presented in this study – combining phylogenetic analysis with evolutionary and developmental computational models – could be used to study the evolution of other traits where a reliable developmental model exists, and phylogenetic data is available (e.g. [5]). This would greatly expand our understanding of how well development can predict trait evolution across phylogeny.

IV. Methods

A. Simulating leaf development

We use a well-established model by Runions *et al.* [14] to approximate the developmental programme underpinning angiosperm leaf shapes. In this model, leaves are represented in 2D as three components (figs. 5a [4] and 5b [4]). First, an open tree structure with growing edges to which new edges are added, representing the vasculature. Second, a polygon with morphogen presence or absence at each vertex that deforms according to the growing vasculature and instructs the initiation of new veins, representing the margin. Third, a passive triangular mesh representing the lamina that fills the space inside the margin.

Many similar models capture a narrow range of shape diversity—for example, variation within *A. thaliana* [46, 47]. On the other hand, the Runions *et al.* [14] model captures a broad spectrum of leaf shapes within a single parameter space. This enables the inference of transition rates between highly divergent shapes such as unlobed and compound. However, it also introduces challenges as the genetic basis of leaf shape differs among lineages and remains poorly resolved for many non-model species. To address this, the Runions *et al.* [14] model is formulated at the level of developmental mechanism rather than specific genes. It is important to understand therefore, that this model is limited in what it can tell us about leaf development and evolution at the level of specific genes. Growth in the Runions *et al.* [14] framework is a product of two key developmental mechanisms. The first is morphogen patterning on the leaf margin. The margin—the tissue at the adaxial-abaxial boundary—has long been recognised as a critical site for patterning during leaf development. For example, auxin patterns marginal outgrowths in *Solanum* [49], *A. thaliana* [46] and *C. hirsuta* [50]. The products of genes such as *CUC2* in *A. thaliana* [46] or *RCO* in *C. hirsuta* [40] also localise to the margin and determine whether these auxin-driven outgrowths form serrations, lobes or leaflets. The closely related *CUC3* and *NO APICAL MERISTEM (NAM)* genes localise to leaflet primordia boundaries, and their down-regulation reduces serration and lobing in diverse lineages, namely *Aquilegia caerulea*, *Pisum sativum*, *Cardamine hirsuta*, *Solanum lycopersicum* and *Solanum tuberosum* [51]. Although the role of auxin in patterning protrusions is highly conserved, many other developmental regulators are not conserved across phylogeny. For example, *RCO* is only present in the Bras-sicaceae family [52] and therefore plays no role in determining leaf shape in other compound lineages such as *Solanum*. Furthermore, there are likely many other genes regulating leaf shape in non-model species that are currently unknown. The model accounts for this variation and uncertainty by not assuming a conserved identity for the morphogens. The number of morphogens present can vary as well as their effect on growth according to their associated parameter values. What is assumed to be conserved is the broader principle that growth-modifying morphogens form patterns along the margin.

The second key mechanism that the model incorporates is the alignment of local growth direction with veins. This is consistent with the telome theory [53], according to which leaves first evolved by the webbing of tissue between shoots, as well as a number of lines of empirical evidence. At the tissue level, deeply dissected leaves (e.g. *Pelargonium carnosum*) necessarily align growth directions with main veins and variegated leaves also show growth orientations that track major veins [14]. At the molecular level, auxin has been shown to pattern veins [54, 55] as well as marginal outgrowths [46, 49, 50] in model species, consistent with the model assumption that vein growth drives margin growth. In the Runions *et al.* [14] model, consistent with observed patterns of auxin efflux carriers in *A. thaliana* [56], new veins are inserted between the current vasculature and special points on the margin known as “convergence points”. New convergence points form between old convergence points depending on the distance between old points and the effect of any morphogens present (fig. 5c [4]). This is an extension of Hofmeister’s rule [57] for leaf

primordium emergence in the shoot apical meristem, where new primordia emerge in the largest space between existing primordia. Some properties of veins are not captured by the model, specifically closed reticulate and parallel vein architectures. This is partly because there is some uncertainty surrounding the mechanisms by which these architectures form [55, 58], but we acknowledge this as a limitation of the model.

More broadly, much of our understanding of leaf development derives from a small number of model species. These have revealed some conserved genes and mechanisms [59], but many details are lineage-specific [52]. Consequently, some unavoidable uncertainty remains in how well the Runions *et al.* [14] model captures global trends.

B. Generating the initial sample

Since a range of different forms are contained within each shape category, we began our simulations from leaves drawn from an initial sample representing a wide variety of morphologies (16 unlobed, 8 lobed, 11 dissected, 13 compound) (figs. S1 and 1b). This sample was generated using a variety of methods. Some were taken from Runions *et al.* [14], others were generated by tuning parameters by hand and some were taken from the output of random walks on leaves from Runions *et al.* [14].

C. *In-silico* mutagenesis

To simulate mutation to our sample, we used a constrained random walk algorithm. The algorithm samples the parameter space of biologically plausible leaf shapes surrounding the starting leaf by making a small change to one of the parameters, checking the validity of the resultant leaf (table S1), and then moving onto another parameter (see fig. S2 for examples of rejected leaves). We chose this approach over other methods, such as uniform sampling or latin hypercube sampling, as the biologically relevant ranges of each parameter were unknown. Random walks have been shown to cause some biases towards high-robustness regions of parameter space in discrete parameter spaces represented as networks [60], which in this case would result in the under-representation of invalid leaves. We have no reason to assume that this would systematically bias the rates between valid leaves but this cannot be fully excluded. For each of the 48 leaves in our initial sample (fig. S1), we ran 5 random walks for 320 steps, generating a total of $5 \times 48 \times 320 = 76,800$ per simulation. We allow 100/122 model parameters to vary throughout the simulations (table S1). Not all parameters are included, as some are not relevant to leaf shape (e.g. control how the image is rendered).

To test the robustness of our results, we repeated the simulation with 2 different mutational schemes. MUT2 (alg. 1) iterates through the parameters in a random order, and attempts to change the parameter by a value selected at random from an array of numbers randomly generated at 3 different orders of magnitude. MUT5 (alg. S2) is the same as MUT2 except the value each parameter is multiplied by 10% of the range of that parameter within the initial leaves (fig. S1). The aim here was to provide some way of accounting for the biologically relevant sampling range.

Check	Condition
<i>Weight difference</i>	The difference in leaf area either side of the midvein must not exceed 6% of the total leaf area.
<i>Leaf width</i>	Final leaf must be wider than 110% of the width of the leaf at the first iteration of the growth simulation.
<i>Veins outside lamina</i>	Less than 3.7% of the total vein area can be non-overlapping with the lamina.
<i>Veins width</i>	The maximum width of the vasculature must be greater than 5% of the total image width.
<i>Veins area</i>	The proportion of total leaf area that is occupied by veins must be greater than 5%.
<i>Overlapping margin</i>	The proportion of the lamina (interior area of the leaf) that is occupied by leaf margin must be less than 0.1%.
<i>Error or timeout</i>	Simulation exits unexpectedly or lasts longer than 160 seconds.

TABLE I. Random walk constraints.

The only type of selection we apply during the walk is stabilising selection for physically plausible leaves. See [fig. S2](#) for examples of rejected leaves. *Weight difference* ([fig. S2a](#)) ensures that the midvein and petiole are roughly in line with the centre of mass of the leaf. *Leaf width* ([fig. S2b](#)) ensures that leaves which stop growing after the first few stages of the model simulation are rejected. *Veins outside lamina* ([fig. S2c](#)) prevents leaves forming whose veins grow outside the lamina. *Veins width* ([fig. S2d](#)) and *Veins area* ([fig. S2e](#)) ensure that leaves that generate without adequate vasculature are rejected. *Overlapping margin* ([fig. S2f](#)) rejects leaves with ingrown margins. *Error or timeout* – some parameter combinations result in the simulation exiting unexpectedly or taking a very long time to complete, posing an additional constraint.

Algorithm 1.

Random Walk MUT2

```

 $\theta \leftarrow$  initial parameter values
for step  $\leftarrow$  0 to  $N_{maxGenerated} - 1$  do
  for attempt  $\leftarrow$  0 to  $N_{maxAttempts} - 1$  do
     $\theta' \leftarrow \theta$  // Copy parameters for this step
    mutation  $\leftarrow$  random draw from {
      uniform[0.1, 1),
      uniform[1, 10),
      uniform[10, 100)
    }
    sign  $\leftarrow$  random choice from  $\{-1, 1\}$ 
     $i \leftarrow$  random draw from  $\{0, \dots, \text{length}(\theta) - 1\}$ 
     $\theta'[i] \leftarrow \theta[i] + (\text{mutation} \times \text{sign})$  // Mutate
    Run simulation with  $\theta'$ 
    if leaf passes sanity checks then
       $\theta \leftarrow \theta'$  // Save new parameters
      break

```

D. Shape classification

To sort the leaves into each of the four shape categories (unlobed, lobed, dissected, compound), we developed a computer-vision program (fig. 6a), to reduce error from subjectivity. It finds the coordinates of key landmarks in the leaf image, extracts morphometric measurements and assigns a category based on threshold values of these measurements (eq. (1)). These morphometric measurements were also used to check if leaves were valid (table I) in the random walk algorithm.

First, the main visual features are extracted from the leaf image, including silhouette, margin, lamina and veins. The coordinates of each pixel along the edge of the silhouette (the contour) are extracted. A reference point is defined in the centre of the mid-vein at the point where it first branches outwards into lateral veins (fig. 6b). Then, the distance between each point on the contour and this reference point is calculated (fig. 6c). The local maxima and minima in this distance array that pass a certain prominence threshold are determined using `find_peaks` from the `scipy.signal` package with `distance=10` and `prominence=25`. These pertinent extrema represent estimates for the coordinates of the tips and sinuses of any protrusions in the leaf margin (fig. 6b shown in red and blue). Because of the prominence threshold, small protrusions such as serrations are excluded.

From the coordinates of the local minima, maxima and the reference point, we calculate a number of morphometric measurements which are then used to classify the leaf. 3 measurements are used: e , the total no. pertinent local maxima and minima; α , the average angle bound by each local maxima and its two adjacent minima; and ω , the average horizontal distance between each pair of local minima, closest to the midvein, that share the same vertical coordinate (measured in pixels). The definitions for each of the shape categories are as follows:

$$\left. \begin{aligned}
 &(\epsilon > 5 \wedge \alpha < 8^\circ) \vee \\
 &(\epsilon > 5 \wedge \omega < 25 \text{ pixels}) \} \Rightarrow \text{Compound (c),} \\
 &\epsilon > 5 \wedge \alpha < 30^\circ \Rightarrow \text{Dissected (d),} \\
 &\epsilon > 5 \wedge \alpha < 60^\circ \Rightarrow \text{Lobed (l),} \\
 &\epsilon \leq 5 \vee \alpha > 60^\circ \Rightarrow \text{Unlobed (u).}
 \end{aligned} \right\} \tag{1}$$

Any leaf that fails to satisfy any of these conditions is labelled as “unclassified”, however there were no unclassified leaves in both the MUT2 and MUT5 simulations.

Leaf shape is a continuous trait, thus the boundaries between discrete categories, such as those in eq. (1), are arbitrary. This makes it difficult to assess the accuracy of the classifier algorithm. Nevertheless, we wanted to have some measure of how close it matches classification by eye. Therefore, we manually classified 200 leaves randomly selected from one of the simulations, with an equal number from the output of each initial leaf. 78.5% of the manually ascribed labels matched those assigned by the classifier. For comparison, we fit a machine learning algorithm (Histogram-based Gradient Boosting Regression Tree from the scikit-learn python package) to a sample of 500 manually labelled leaves, sampled in the same way as the set of 200. The algorithm was trained on the distance of 200 equally spaced points along the leaf contour to the reference point for each leaf (fig. 6c). When used to predict shape for the same sample for 200 leaves used for the classical classifier, it returned an accuracy of 73.0%. We then repeated the machine learning approach using morphometric data rather than contour data, which returned an accuracy of 79.5%. This suggests that the threshold classifier is a good enough match with subjective shape classification compared to other classification methods. It also demonstrates that the morphometrics used to construct the original classifier are useful features when classifying shape as unlobed, lobed, dissected or compound. The two machine-learning methods agreed on 77.0% of the labels.

E. Inferring evolutionary rates from random walk data

To test the model for bias in leaf shape transition rates upon mutation, we ran 5 random walks on each of the 48 initial leaves, for 320 steps, generating a total of 76,800 leaves for both MUT2 and MUT5. Using the classifier (fig. 6), we assigned each leaf the category unlobed, lobed, dissected or compound (eq. (1)). Evolutionary rates between shape categories were inferred by fitting a continuous-time Markov chain (CTMC) to the labelled walks.

This model is defined by a matrix of transition rates Q between each state - unlobed (u), lobed (l), dissected (d) and compound (c)

$$Q = \begin{bmatrix} \sum_{\theta \in \{l,d,c\}} -q_{u\theta} & q_{ul} & q_{ud} & q_{uc} \\ q_{lu} & \sum_{\theta \in \{u,d,c\}} -q_{l\theta} & q_{ld} & q_{lc} \\ q_{du} & q_{dl} & \sum_{\theta \in \{u,l,c\}} -q_{d\theta} & q_{dc} \\ q_{cu} & q_{cl} & q_{cd} & \sum_{\theta \in \{u,l,d\}} -q_{c\theta} \end{bmatrix} \tag{2}$$

- where the diagonal elements must be equal to the negative sum of all other elements in the same row [39].

The model assumes that at any given instant, over a continuous time interval, the system is in one of these four states. The probability $p(i, j, t)$ of transitioning between state i and j , depends only on the current state, the transition rate matrix Q and the time t until the end of the transition. The matrix of all possible transition probabilities $P(t)$ relates to Q and t like so [39, 42]

$$P(t) = \exp(Qt). \tag{3}$$

Therefore, the probability of a transition from state i to j over time t is

$$p(i, j, t) = \exp(\mathbf{Q}t)_{ij}. \quad (4)$$

The probability of a set of observed walks $\{x = x_0, x_1, \dots, x_W\}$, where each walk contains a sequence of labelled steps $x_w = \{x_{w0}, x_{w1}, \dots, x_{wn}\}$, given a particular set of rate parameters \mathbf{Q} is

$$P(x | \mathbf{Q}) = \prod_{w=0}^W \prod_{n=0}^{N-1} p(x_{wn}, x_{w(n+1)}, t), \quad (5)$$

where w and n are walk and step index respectively. We assume $t = 1$ for every step, as over the course of the entire random walk, the average step size in parameter space will be approximately constant.

Rates were inferred from the random walk simulation data by finding the maximum of this function, using MCMC with the emcee package for python [61]. We ran 24 chains for 5,000 steps with uniform priors between 0 and 0.1 for all transition rates. Chains were thinned by removing all but every 10th step (figs. S9 and S10) and the first 2,500 were discarded to obtain the posterior distribution. We took the mean of the posterior for each transition as the estimate for each transition rate in fig. 1. As an additional check, we also inferred the rates via maximum likelihood estimation (MLE) and found that they converged with the MCMC estimates (figs. 7a and 7b).

Because the inferred rates depend on the unit of time, we needed to normalise the transition rates to compare with the rates inferred from phylogeny. We did this for each mutational scheme by dividing by the mean of all transition rate estimates q_{ij} . We define a normalised rate as

$$\hat{q}_{ij} = \frac{q_{ij}}{N}, \quad (6)$$

$$N = \frac{1}{12} \sum_{\substack{i, j \in \{u, l, d, c\} \\ i \neq j}} q_{ij}.$$

F. Inferring evolutionary rates from phylogeny

Although Geeta *et al.* [36] had already inferred rates between these 4 categories with a phylogeny of 561 taxa, we wanted to test how robust these results are. Those results used a phylogeny that is now 25 years old [62], and there are now much larger (e.g. Zuntini *et al.* [38] 9,506 species) and more reliable phylogenies available. To do this analysis, we combined 2 phylogenetic trees [37, 38] with leaf images from the Naturalis herbarium [31]. We first filtered all digitised records to just those from angiosperm families as per APG IV [63]. We then removed all but the first record for each species.

To see how the rates are affected by looking at the species level vs the genus level, we generated 2 trees for each original tree, one with species at the tips and one with all but one tip per genera removed. For the species level trees, we took the intersect between the filtered herbarium data and the full trees. If more than 3,100 tips remained, we sampled at random 3,100 tips, to allow for manual classification within reasonable time limits. This gave a combined set 5,552 leaf images to classify for both species trees. Notably, the intersect between the herbarium data and the Zuntini *et al.* [38] species tree gave 3057 taxa. Since this was less than 3,100, no further sub-sampling was required, and so unlike all the other trees in this analysis, the Zuntini *et al.* [38] species tree represents the raw intersect between the original tree and the herbarium. For this phylogeny, our genus tree and our species tree should both have similar levels of reliability. For the Janssens *et al.* [37] phylogeny, we trust the genus level tree more than the species level tree, because the overlap with the herbarium data was not as good at the species level. For example, taxa were more heavily

concentrated in a smaller number of angiosperm families, compared to the other trees (fig. S8). Moreover, we sub-sampled down to 3,100 taxa from 16,436 in the raw intersect to make manual classification feasible - something not required for the other trees.

For the genus level trees, we further sub-sampled the angiosperm filtered herbarium data to 200 records per family, or the maximum if fewer than 200 records were present. The purpose of this step was to sample more evenly over phylogeny, accounting for potential biases in the herbarium. We then took the intersect between genera in this sample and each of our phylogenies and selected at random one species from within each genus that was present in the herbarium to use as the label. This gave a combined set of 3,359 leaf images to classify for both genus trees.

These images were then used to classify each species as one of the 4 categories (unlobed, lobed, dissected or compound) or ambiguous. This was done by eye, relying on a sample of model leaves and their morphometric classifier output for reference (figs. S1 and 1b). After classification, all species labelled as ambiguous were removed, leaving an unambiguously labelled phylogeny (table II and figs. S3 to S7).

We used 3 different phylogenies from the literature to infer the evolutionary rates. First, to compare our predictions directly to the Geeta *et al.* [36] study we used their phylogeny. Second, we used a tree by Janssens *et al.* [37] for a large sample, as this tree contained over of 36,101 species, representing 8,399 genera, 426 families and all orders of flowering plants (available in Supplementary material 5 of Janssens *et al.* [37]). Finally, we used a recent phylogeny by Zuntini *et al.* [38], containing 9,506 species (7,923 genera and all angiosperm families and orders), inferred using 353 nuclear genes and calibrated with 200 fossil records (available at <https://zenodo.org/records/10778207> we used 4-young-tree-smoothing-10-pruned-for-diversification-analyses). We show results from the Geeta *et al.* [36] phylogeny, labelled with their own classification data and the Janssens *et al.* [37] and Zuntini *et al.* [38] trees labelled with our own classification data from Naturalis Biodiversity Center [31].

Rates were inferred using the MultiState method from BayesTraits V4.1.2 [39]. We used MCMC for 1,010,000 iterations with a sample period of 1000 and removed the first 10000 iterations to obtain posterior distributions. We used a uniform prior between 0 and 0.1 for the Zuntini *et al.* [38] trees (figs. S11 and S12) and Janssens *et al.* [37] trees (figs. S13 and S14), however, as the unit of branch length was substitutions per site rather than millions of years, for Geeta *et al.* [36] we instead used a uniform prior between 0 and 100 (fig. S15). We took the mean of each posterior as the transition rate estimate. In addition, we estimated the rates using MLE for comparison (figs. 7a and 7b). As the units of time differ between our phylogenies and the simulations, we needed to normalise the rates to allow for comparison. We did this using the same method as for the simulations (eq. (6)): by dividing each transition rate estimate by the mean transition rate estimate for each tree.

Tree	Unlobed (u)	Lobed (l)	Dissected (d)	Compound (c)	ld	lc	ldc	No. tips
Geeta <i>et al.</i> [36]	428 (76.3%)	38 (6.8%)	17 (3.0%)	68 (12.1%)	7 (1.2%)	1 (0.2%)	2 (0.4%)	561
Zuntini <i>et al.</i> [38] species	2445 (86.0%)	102 (3.6%)	83 (2.9%)	212 (7.5%)	-	-	-	2842
Zuntini <i>et al.</i> [38] genus	1910 (88.3%)	82 (3.8%)	47 (2.2%)	123 (5.7%)	-	-	-	2162
Janssens <i>et al.</i> [37] species	2606 (87.2%)	125 (4.2%)	99 (3.3%)	160 (5.4%)	-	-	-	2990
Janssens <i>et al.</i> [37] genus	2178 (88.6%)	97 (3.9%)	55 (2.2%)	128 (5.2%)	-	-	-	2458

TABLE II. The number of taxa in each shape category for each labelled tree used in our inference.

Columns with multiple shape categories (e.g. ld) represent taxa which could fit into any of the listed categories.

V. Extended data

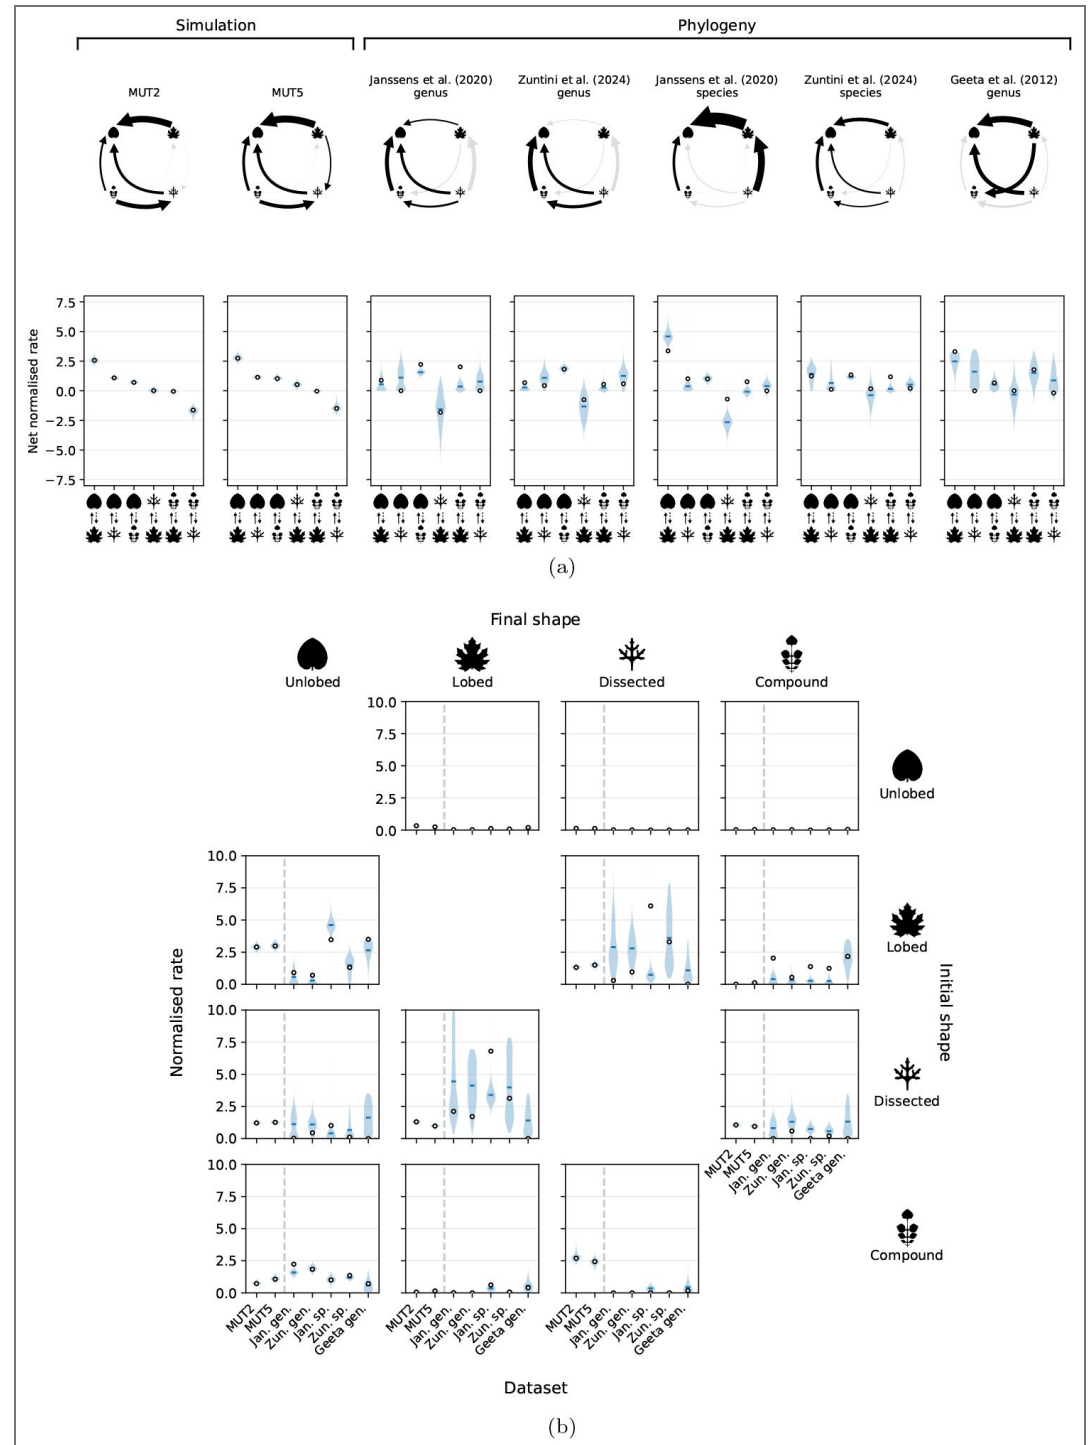


FIG. 7. Further details of transition rates from phylogenies and simulations. (a) Net normalised transition rates between the 4 leaf shape categories, inferred from both simulation and phylogeny. (Top row) Arrow thickness and direction represents the mean net normalised transition rate between each pair of leaf shape categories. Black arrows represent transitions where the 0 does not fall within an equal-tailed 95% credible interval for the net normalised posterior, and grey where it does. (Bottom row) Posterior distributions for the net normalised rate of all transitions. Solid bars indicate the mean. The net rate is calculated by subtracting the transition given by the dashed arrow from the transition given by the solid arrow in the x-axis labels. **(b)** Normalised transition rates between different leaf shapes, inferred from both *in-silico* mutagenesis (left of dashed

line) and phylogeny (right of dashed line). Each subplot represents a different transition, from the shape shown on the vertical figure axis to the shape shown on the horizontal figure axis. These rates were normalised by dividing by the mean transition rate estimate within each dataset (eq. (6) [↗](#)).

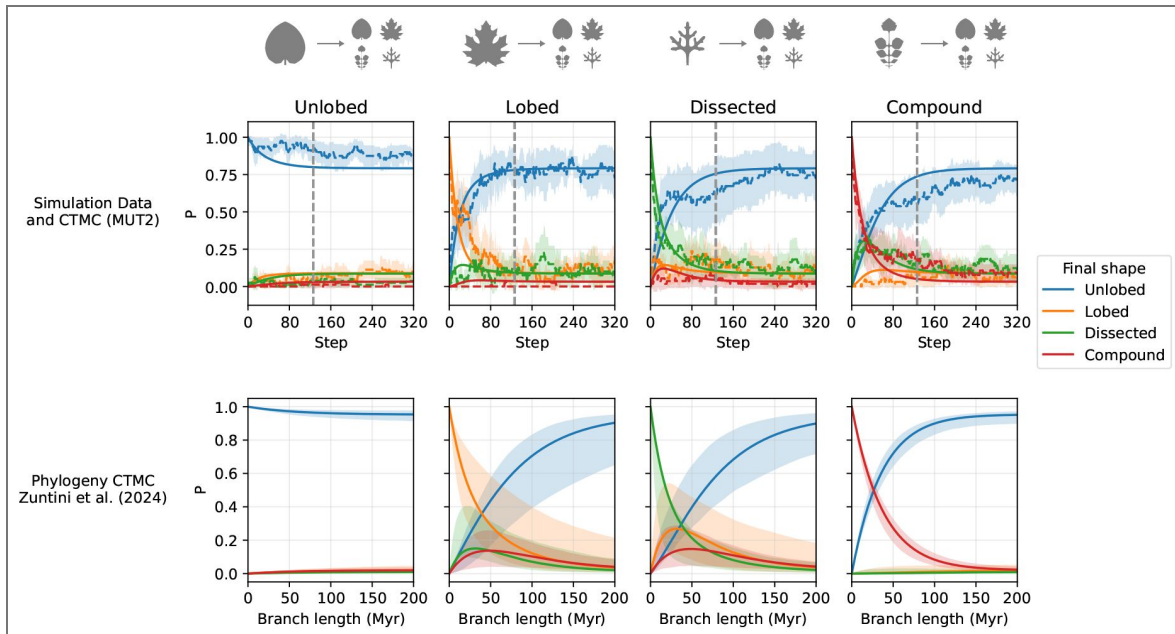


FIG. 8. Comparing *in-silico* mutagenesis and phylogeny for the probability of each shape transition over time, separated out by starting shape.

The column header gives the initial shape and the line colour gives the final shape. **(Top row)** The dashed line represents the proportion of walks from each unique initial leaf (fig. S1), averaged over all initial leaves in each state. The error bands represents $\pm 1.96\times$ standard error for this average. The solid lines represents the CTMC model fit (MLE) to the output of the random walks. The vertical, grey, dashed line represents the estimated simulation step that is equivalent to $t = 200$ Myr for the phylogeny CTMC. This is calculated as described in fig. 3, and is also the simulation x-limit in fig. 3. We show all 320 simulation steps here for a more detailed summary of the simulation trajectories. **(Bottom row)** The predicted transition probabilities given by the CTMC model fit to the labelled Zuntini *et al.* [38] genus tree (fig. S3). The solid line represents the MLE fit and the errors are inferred from MCMC posteriors via Monte-Carlo error propagation ($N = 500$) as per fig. 3. Note the good qualitative agreement between the dynamics between the phylogenies and the simulations, especially for the (blue) data related to transitions to/from unlobed leaves.

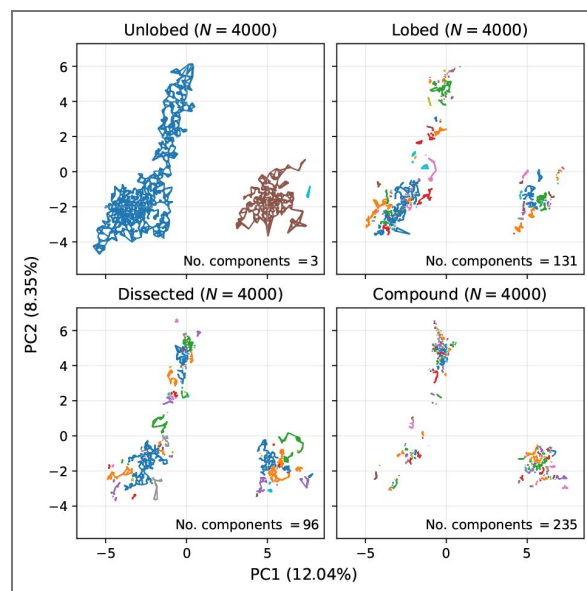


FIG. 9. The nearest neighbour graph of parameter PCA space for each shape category highlights how the unlobed space is the most connected, followed by dissected, lobed and then compound.

Complex leaves subdivide into many more connected components than unlobed. The number of nearest neighbours used is 6. Colours highlight different connected components, however, since the number of unique colours is only 10, this under represents the true number of connected components for lobed, dissected and compound. Thus, the colours here are intended only as a visual guide.

Data availability

The code for the analyses presented in this paper is openly accessible at <https://github.com/malonej501/leaf-shape-evolution>.

Acknowledgements

JSM was supported by funding from the Biotechnology and Biological Sciences Research Council (UKRI-BBSRC) [grant number BB/T008784/1]. LMD was supported, in part, by National Science Foundation OISE 2020577 and a Stony Brook University Presidential Innovation and Excellence. SvD and NSM were supported by the Issachar fund. NSM acknowledges support of the Spanish Ministry of Science and Innovation through the Centro de Excelencia Severo Ochoa (CEX2020-001049-S, MCIN/AEI/10.13039/501100011033), the EMBL partnership and the Generalitat de Catalunya through the CERCA programme. Research for this publication has been partially carried out in the Barcelona Collaboratorium for Modelling and Predictive Biology. This research is part of grant JDC2022-049526-I to NSM, funded by MCIU/AEI/10.13039/501100011033 and by European Union NextGenerationEU/PRTR.

Additional files

[Supplementary Information](#)

Additional information

Funding

Funder	Grant reference number	Author
UKRI Biotechnology and Biological Sciences Research Council (AFRC)	BB/T008784/1	James S Malone
NSF OD Office of International Science and Engineering (ISE)	2020577	Liliana M Davalos
SUNY Stony Brook University (SBU)		Liliana M Davalos
Issachar fund		Samuel HA von der Dunk
Issachar fund		Nora S Martin
European Molecular Biology Laboratory (EMBL)		Nora S Martin
Generalitat de Catalunya (Government of Catalonia)		Nora S Martin
Ministerio de Ciencia, Innovación y Universidades (MCIU)		Nora S Martin

Author ORCID iDs

James S Malone: <https://orcid.org/0009-0002-6243-3218>

Samuel HA von der Dunk: <https://orcid.org/0000-0002-5741-7558>

References

- [1] Smith JM, Burian R, Kauffman S, Alberch P, Campbell J, Goodwin B, Lande R, Raup D, Wolpert L (1985) Developmental Constraints and Evolution: A Perspective from the Mountain Lake Conference on Development and Evolution. *The Quarterly Review of Biology* **60**:265 <https://doi.org/10.1086/414425>
- [2] Alberch P (1980) Ontogenesis and morphological diversification. *Am Zool* **20**:653

- [3] Uller T, Moczek AP, Watson RA, Brakefield PM, Laland KN (2018) Developmental Bias and Evolution: A Regulatory Network Perspective. *Genetics* **209**:949 <https://doi.org/10.1534/genetics.118.300995> | PubMed
- [4] Alberch P, Gale EA (1985) A developmental analysis of an evolutionary trend: Digital reduction in amphibians. *Evolution* **39**:8 <https://doi.org/10.2307/2408513>
- [5] Prusinkiewicz P, Erasmus Y, Lane B, Harder LD, Coen E (2007) Evolution and Development of Inflorescence Architectures. *Science* **316**:1452 <https://doi.org/10.1126/science.1140429> | PubMed
- [6] Kavanagh KD, Shoval O, Winslow BB, Alon U, Leary BP, Kan A, Tabin CJ (2013) Developmental bias in the evolution of phalanges. *Proc. Natl. Acad. Sci. U.S.A* **110**:18190 <https://doi.org/10.1073/pnas.1315213110> | PubMed
- [7] Hagolani PF, Zimm R, Vroomans R, Salazar-Ciudad I (2021) On the evolution and development of morphological complexity: A view from gene regulatory networks. *PLoS Comput Biol* **17**:e1008570 <https://doi.org/10.1371/journal.pcbi.1008570> | PubMed
- [8] Johnston IG, Dingle K, Greenbury SF, Ca-margo CQ, Doye JPK, Ahnert SE, Louis AA (2022) Symmetry and simplicity spontaneously emerge from the algorithmic nature of evolution. *Proc. Natl. Acad. Sci. U.S.A* **119**:e2113883119 <https://doi.org/10.1073/pnas.2113883119> | PubMed
- [9] Zimm R, Berio F, Debais-Thibaud M, Goudemand N (2023) A shark-inspired general model of tooth morphogenesis unveils developmental asymmetries in phenotype transitions. *Proc. Natl. Acad. Sci. U.S.A* **120**:e2216959120 <https://doi.org/10.1073/pnas.2216959120> | PubMed
- [10] Staps M, Miller PW, Tarnita CE, Mallarino R (2023) Development shapes the evolutionary diversification of rodent stripe patterns. *Proc. Natl. Acad. Sci. U.S.A* **120**:e2312077120 <https://doi.org/10.1073/pnas.2312077120> | PubMed
- [11] Rohner PT, Berger D (2023) Developmental bias predicts 60 million years of wing shape evolution. *Proc. Natl. Acad. Sci. U.S.A* **120**:e2211210120 <https://doi.org/10.1073/pnas.2211210120> | PubMed
- [12] Laland K, Uller T, Feldman M, Sterelny K, Müller GB, Moczek A, Jablonka E, Odling-Smee J, Wray GA, Hoekstra HE, *et al.* (2014) Does evolutionary theory need a rethink?. *Nature* **514**:161 <https://doi.org/10.1038/514161a> | PubMed
- [13] Welch JJ (2017) What's wrong with evolutionary biology?. *Biol Philos* **32**:263 <https://doi.org/10.1007/s10539-016-9557-8> | PubMed
- [14] Runions A, Tsiantis M, Prusinkiewicz P (2017) A common developmental program can produce diverse leaf CERCA programme. *New Phytologist* **216**:401 <https://doi.org/10.1111/nph.14449> | PubMed
- [15] Nikolov LA, Runions A, Gupta M Das, Tsiantis M (2019) Leaf development and evolution. In: Grossniklaus U (Ed). *Current Topics in Developmental Biology* **131** Elsevier. pp. 109-139 <https://doi.org/10.1016/bs.ctdb.2018.11.006> | PubMed
- [16] Nicotra AB, Leigh A, Boyce CK, Jones CS, Niklas KJ, Royer DL, Tsukaya H (2011) The evolution and functional significance of leaf shape in the angiosperms. *Functional Plant Biol* **38**:535 <https://doi.org/10.1071/fp11057> | PubMed
- [17] Kidner CA, Umbreen S (2010) Why is leaf shape so variable?. *International Journal of Plant Developmental Biology* **4**:64
- [18] Vogel S (1970) Convective cooling at low airspeeds and the shapes of broad leaves. *J Exp Bot* **21**:91 <https://doi.org/10.1093/jxb/21.1.91>
- [19] Givnish TJ (1987) Comparative studies of leaf form: Assessing the relative roles of selective pressures and phylogenetic constraints. *New Phytologist* **106**:131 <https://doi.org/10.1111/j.1469-8137.1987.tb04687.x>
- [20] Sisó S, Camarero J, Gil-Pelegrín E (2001) Relationship between hydraulic resistance and leaf morphology in broadleaf Quercus species: A new interpretation of leaf lobation. *Trees* **15**:341 <https://doi.org/10.1007/s004680100110>
- [21] Roden JS (2003) Modeling the light interception and carbon gain of individual fluttering aspen (*Populus tremuloides* Michx) leaves. *Trees* **17**:117 <https://doi.org/10.1007/s00468-002-0213-3>

- [22] Perez RPA, Dauzat J, Pallas B, Lamour J, Verley P, Caliman J.-P, Costes E, Faivre R (2018) Designing oil palm architectural ideotypes for optimal light interception and carbon assimilation through a sensitivity analysis of leaf traits. *Annals of Botany* **121**:909 <https://doi.org/10.1093/aob/mcx161> | PubMed
- [23] Otto D, Munz S, Memic E, Hartung J, Graeff-Hönninger S (2025) A computer vision approach for quantifying leaf shape of maize (*Zea mays* L.) and simulating its impact on light interception. *Front. Plant Sci* **16**:1521242 <https://doi.org/10.3389/fpls.2025.1521242> | PubMed
- [24] Dell'Aglio DD, Losada ME, Jiggins CD (2016) Butterfly learning and the diversification of plant leaf shape. *Front. Ecol. Evol* **4**:81 <https://doi.org/10.3389/fevo.2016.00081>
- [25] Higuchi Y, Kawakita A (2019) Leaf shape deters plant processing by an herbivorous weevil. *Nat. Plants* **5**:959 <https://doi.org/10.1038/s41477-019-0505-x> | PubMed
- [26] Campitelli BE, Stinchcombe JR (2013) Testing potential selective agents acting on leaf shape in *Ipomoea hederacea*: Predictions based on an adaptive leaf shape cline. *Ecology and Evolution* **3**:2409 <https://doi.org/10.1002/ece3.641>
- [27] Leigh A, Sevanto S, Close J, Nicotra A (2017) The influence of leaf size and shape on leaf thermal dynamics: Does theory hold up under natural conditions?. *Plant Cell & Environment* **40**:237 <https://doi.org/10.1111/pce.12857> | PubMed
- [28] Sack L, Frole K (2006) Leaf structural diversity is related to hydraulic capacity in tropical rain forest trees. *Ecology* **87**:483 <https://doi.org/10.1890/05-0710> | PubMed
- [29] Scoffoni C, Rawls M, McKown A, Cochard H, Sack L (2011) Decline of leaf hydraulic conductance with dehydration: Relationship to leaf size and venation architecture. *Plant Physiology* **156**:832 <https://doi.org/10.1104/pp.111.173856> | PubMed
- [30] Donoghue MJ, Eaton DAR, Maya-Lastra CA, Landis MJ, Sweeney PW, Olson ME, Cacho NI, Moeglein MK, Gardner JR, Heaphy NM, *et al.* (2022) Replicated radiation of a plant clade along a cloud forest archipelago. *Nat Ecol Evol* **6**:1318 <https://doi.org/10.1038/s41559-022-01823-x> | PubMed
- [31] (2024) Naturalis Biodiversity Center, Naturalis Bioportal.
- [32] Plowman TC (1981) *Erythroxyllum confusum* Britton. <https://data.biodiversitydata.nl/naturalis/specimen/U.1284618>
- [33] de Bruijn J (1981) *Kirengeshoma palmata* Yatabe. <https://data.biodiversitydata.nl/naturalis/specimen/WAG.1457187>
- [34] van Royen P (1973) *Geranium caespitosum* E.James. <https://data.biodiversitydata.nl/naturalis/specimen/L.2011772>
- [35] Bhikhi CR (2011) *Crepidospermum goudotianum* (Tul.) Triana & Planch. <https://data.biodiversitydata.nl/naturalis/specimen/U.1615275>
- [36] Geeta R, Dávalos LM, Levy A, Bohs L, Lavin M, Mummenhoff K, Sinha N, Wojciechowski MF (2012) Keeping it simple: Flowering plants tend to retain, and revert to, simple leaves. *New Phytologist* **193**:481 <https://doi.org/10.1111/j.1469-8137.2011.03951.x> | PubMed
- [37] Janssens SB, Couvreur TL, Mertens A, Dauby G, Dagallier L.-PM, Abeele S Vanden, Vandeloek F, Mascarello M, Beeckman H, Sosef M, *et al.* (2020) A largescale species level dated angiosperm phylogeny for evolutionary and ecological analyses. *Bdj* **8**:e39677 <https://doi.org/10.3897/bdj.8.e39677> | PubMed
- [38] Zuntini AR, Carruthers T, Maurin O, Bailey PC, Leempoel K, Brewer GE, Epitawalage N, Françoso E, Gallego-Paramo B, McGinnie C, *et al.* (2024) Phylogenomics and the rise of the angiosperms. *Nature* **629**:843 <https://doi.org/10.1038/s41586-024-07324-0> | PubMed
- [39] Pagel M (1994) Detecting correlated evolution on phylogenies: A general method for the comparative analysis of discrete characters. *Proc. R. Soc. Lond. B* **255**:37 <https://doi.org/10.1098/rspb.1994.0006>

- [40] Kierzkowski D, Runions A, Vuolo F, Strauss S, Lymbouridou R, Routier-Kierzkowska A.-L, Wilson-Sánchez D, Jenke H, Galinha C, Mosca G, *et al.* (2019) A Growth-Based Framework for Leaf Shape Development and Diversity. *Cell* **177**:1405 <https://doi.org/10.1016/j.cell.2019.05.011> | PubMed
- [41] Pagel M, Meade A, Barker D (2004) Bayesian Estimation of Ancestral Character States on Phylogenies. *Systematic Biology* **53**:673 <https://doi.org/10.1080/10635150490522232> | PubMed
- [42] Kalbfleisch JD, Lawless JF (1985) The Analysis of Panel Data under a Markov Assumption. *Journal of the American Statistical Association* **80**:863 <https://doi.org/10.1080/01621459.1985.10478195>
- [43] McGibbon RT, Pande VS (2015) Efficient maximum likelihood parameterization of continuous-time Markov processes. *The Journal of Chemical Physics* **143**:034109 <https://doi.org/10.1063/1.4926516> | PubMed
- [44] von der Dunk SHA, Martin NS, Dingle K, Louis AA (no date) RNA secondary structures are conserved but random.
- [45] Dingle K, Camargo CQ, Louis AA (2018) Input– output maps are strongly biased towards simple outputs. *Nat Commun* **9**:761 <https://doi.org/10.1038/s41467-018-03101-6> | PubMed
- [46] Bilsborough GD, Runions A, Barkoulas M, Jenkins HW, Hasson A, Galinha C, Laufs P, Hay A, Prusinkiewicz P, Tsiantis M (2011) Model for the regulation of Arabidopsis thaliana leaf margin development. *Proc. Natl. Acad. Sci. U.S.A* **108**:3424 <https://doi.org/10.1073/pnas.1015162108> | PubMed
- [47] Kuchen EE, Fox S, De Reuille P Barbier, Kennaway R, Bensmihen S, Avondo J, Calder GM, Southam P, Robinson S, Bangham A, *et al.* (2012) Generation of Leaf Shape Through Early Patterns of Growth and Tissue Polarity. *Science* **335**:1092 <https://doi.org/10.1126/science.1214678> | PubMed
- [48] Martin NS, Camargo CQ, Louis AA (2024) Bias in the arrival of variation can dominate over natural selection in Richard Dawkins’s biomorphs. *PLoS Comput Biol* **20**:e1011893 <https://doi.org/10.1371/journal.pcbi.1011893> | PubMed
- [49] Koenig D, Bayer E, Kang J, Kuhlemeier C, Sinha N (2009) Auxin patterns Solanum lycopersicum leaf morphogenesis. *Development* **136**:2997 <https://doi.org/10.1242/dev.033811> | PubMed
- [50] Barkoulas M, Hay A, Kougioumoutzi E, Tsiantis M (2008) A developmental framework for dissected leaf formation in the Arabidopsis relative Cardamine hirsuta. *Nat Genet* **40**:1136 <https://doi.org/10.1038/ng.189> | PubMed
- [51] Blein T, Pulido A, Vialette-Guiraud A, Nikovics K, Morin H, Hay A, Johansen IE, Tsiantis M, Laufs P (2008) A Conserved Molecular Framework for Compound Leaf Development. *Science* **322**:1835 <https://doi.org/10.1126/science.1166168> | PubMed
- [52] Vlad D, Kierzkowski D, Rast MI, Vuolo F, Ioio R Dello, Galinha C, Gan X, Hajheidari M, Hay A, Smith RS, *et al.* (2014) Leaf shape evolution through duplication, regulatory diversification, and loss of a homeobox gene. *Science* **343**:780 <https://doi.org/10.1126/science.1248384> | PubMed
- [53] Zimmerman Walter (1952) Main results of the “Telome Theory”. *The Palaeobotanist* **1**:456 <https://doi.org/10.54991/jop.1952.423>
- [54] Scarpella E, Barkoulas M, Tsiantis M (2010) Control of Leaf and Vein Development by Auxin. *Cold Spring Harbor Perspectives in Biology* **2**:a001511 <https://doi.org/10.1101/cshperspect.a001511> | PubMed
- [55] Scarpella E (2024) Leaf Vein Patterning. *Annual Review of Plant Biology* **75**:377 <https://doi.org/10.1146/annurev-arplant-062923-030348> | PubMed
- [56] Scarpella E, Marcos D, Friml J, Berleth T (2006) Control of leaf vascular patterning by polar auxin transport. *Genes Dev* **20**:1015 <https://doi.org/10.1101/gad.1402406> | PubMed
- [57] Kirchoff BK (2003) Shape Matters: Hofmeister’s Rule, Primordium Shape, and Flower Orientation. *International Journal of Plant Sciences* **164**:505 <https://doi.org/10.1086/375421>
- [58] Perico C, Tan S, Langdale JA (2022) Developmental regulation of leaf venation patterns: Monocot versus eudicots and the role of auxin. *New Phytologist* **234**:783 <https://doi.org/10.1111/nph.17955> | PubMed

- [59] Bar M, Ori N (2014) Leaf development and morphogenesis. *Development* **141**:4219 <https://doi.org/10.1242/dev.106195> | PubMed
- [60] Van Nimwegen E, Crutchfield JP, Huynen M (1999) Neutral evolution of mutational robustness. *Proc. Natl. Acad. Sci. U.S.A* **96**:9716 <https://doi.org/10.1073/pnas.96.17.9716> | PubMed
- [61] Foreman-Mackey D, Hogg DW, Lang D, Goodman J (2013) Emcee: The MCMC Hammer. *Publications of the Astronomical Society of the Pacific* **125**:306 <https://doi.org/10.1086/670067>
- [62] Soltis D (2000) Angiosperm phylogeny inferred from 18S rDNA, rbcL, and atpB sequences. *Botanical Journal of the Linnean Society* **133**:381 <https://doi.org/10.1111/j.1095-8339.2000.tb01588.x>
- [63] The Angiosperm Phylogeny Group (2016) An update of the Angiosperm Phylogeny Group classification for the orders and families of flowering plants: APG IV. *Bot J Linn Soc* **181**:1 <https://doi.org/10.1111/boj.12385>
- Zuntini AR, et al (2024) Phylogenomics and the rise of the angiosperms. Zenodo. <https://doi.org/10.5281/zenodo.10778206>

Peer reviews

Reviewer #1 (Public review):

Summary

The authors aim to understand, in the context of leaf shape, how the constraints imposed by development inform evolution. Leaf shape is a good place to study the influence of development on evolution because it is a trait that exhibits a lot of diversity, and the developmental mechanisms that give rise to leaf shapes are apparently rather conserved across angiosperms.

As part of the motivation for their work, the authors cite a previous study (Geeta et al), which found that in angiosperm phylogenies, transitions from complex to simple leaf shapes occur through evolution more often than transitions in the opposite direction. Is this due to developmental constraints or adaptation?

The authors undertake two parallel lines of work:

- (1) Extending the study of Geeta et al with more data, consisting of both phylogenies and a shape classification dataset. The conclusion from this line of inquiry is that transitions from lobed to unlobed leaves are more common than transitions away from unlobed leaves.
- (2) The authors conduct evolution simulations in a computational model of leaf development. Here, they look at neutral mutations and whether simply neutral evolution is sufficient to drive the observed trend.

The conclusion of the second part of the work is that the driver of the evolution toward simple leaf shape is entropy: there are more ways to make unlobed leaves than to make lobed leaves (at least in terms of gene regulation parameters that will produce the two leaf types). The argument is that random gene regulatory networks are more likely to produce unlobed leaves than lobed leaves; therefore, neutral evolution drives this trend.

Data Analysis

Roughly 9000 images of leaves were classified into 4 categories: unlobed, lobed, dissected, and compound. These labels were applied to the tips of 5 phylogenetic trees of angiosperms (3 resolved at the genus level and 2 at the species level). By fitting a continuous-time Markov chain to the labelled trees, the authors claim that there is a significantly higher rate of transition to the unlobed leaf shape compared to transitions to more complex shapes.

Simulation

First, the authors validate a computational model (Runions et al) for leaf growth on an experimental dataset. By changing parameters in the model, they can recapitulate the morphological changes in the shapes of Arabidopsis leaves engendered by expression of two particular genes.

Then the authors run an evolutionary model (without selection, just random mutations) on top of the computational leaf development model. As the random walk in parameter space reaches a stationary distribution, they look at both the proportions of the leaf categories in the steady state as well as the transition rates between different categories. The result is that transitions to unlobed leaves are more common than from unlobed leaves.

General Comments

The authors use angiosperm phylogenies from other works as the basis for the data analysis part of their work. Given the centrality of these phylogenies for their conclusions, more information is needed about how these phylogenies were constructed and what they mean. What is the timescale that they span? What method is used to infer them? What regions of DNA were sequenced in order to build the phylogenies? Also, maybe some more discussion of angiosperm evolution (e.g., when was the most recent common ancestor of all angiosperms?) would help put the study in context.

We also need a more in-depth discussion of the computational model. What are all the >100 parameters doing, and what informs the seemingly strange mutational model that changes parameters by 3 orders of magnitude?

I am confused about how the rates of transitions were inferred from the phylogeny. Here, one has a phylogeny inferred by some method (which needs to be described in more detail), and just the leaves are labelled. It is stated in the methods that BayesTraits was used to infer the transition rates. I realize this method is probably documented elsewhere, but a bit of a summary of how it works and how to interpret its results would (1) make the paper more self-contained and (2) if the algorithm is credible, make the results firmer.

I am a bit skeptical of the authors' interpretation of the biological trend (of complex to simple leaf shapes) as being driven by neutral evolution. Why does one expect that the mutations generated by the random walk models described in the work are in fact neutral mutations?

- If the entropy of simple leaf shapes is higher than that of complex leaf shapes, why did we have complex leaves at all? I suspect the authors might argue that this is due to selection. In that case, what allows these complex shapes to become simpler? Wouldn't they be losing the selective advantage that drove them to be more complex in the first place? Or maybe the idea is that the rates are inferred assuming some steady state that generates the phylogeny? I did not understand this point.

Are the rates of transitions between leaf types inferred for the phylogeny assuming that the phylogeny is generated by the steady state of some Markov process? (I think the answer is no: in that case, how does one explain the initial condition?) If I take the mutation model (random walk) seriously, then shouldn't I expect that this steady state obeys detailed balance? In that case I should have $p_{i \rightarrow j} = p_{j \rightarrow i}$ for each of the occupancies p_i and transition rates $r_{i \rightarrow j}$ for the shape categories. How close are the rates inferred from the phylogenies to obeying detailed balance? Presumably, the Markov chain fitted to the simulation data obeys detailed balance because the mutation model itself does?

I find it hard to take the discussion of development seriously without some consideration of mechanics. Presumably, the mechanics are hidden in the computational leaf development model, but this model is not discussed in enough detail for the reader to know. It seems to me

that the interesting question is: what are the mechanical constraints on development that drive the apparent trend in evolution towards simpler leaf shapes? Maybe it is something about the type of differential growth needed to make complex leaf shapes less robust to mutation. But in this case, I would assume that selection plays a role in the complexity of shape. In any case, a better understanding (or explanation) of the computational model is needed to make this interpretation.

Some discussion of timescales is needed, especially when invoking neutral evolutionary arguments. If a neutral mutation occurs, its time to fix in a population of size N is $\sim N$ generations. What are the relevant angiosperm population sizes and the number of mutations that separate branches on the tree? Are timescales remotely consistent with e.g., the age of angiosperms on Earth?

<https://doi.org/10.7554/eLife.110692.1.sa2>

Reviewer #2 (Public review):

Strengths:

The paper's underlying question is interesting, extending the authors' prior work on RNA along similar conceptual lines. The paper combines both image analysis of leaves and a computational analysis of a simple model of leaf development.

Weaknesses:

The entire paper is based on the Runion model. More intuition about the Runion model would be useful for a broader readership that cares about the evolutionary aspect of this, but may not know the developmental model in question. Obviously, this is prior well-established work, but 2 - 3 sentences highlighting the key structural aspects of such a model would be great. Currently, that intuition is found implicitly in a sentence on page 2 ("complex leaf shapes need more specificity in their GRNs than their simpler unlobed leaf shape"), but the reader is left wondering - is the Runion model a detailed mechanistic one with multiple interacting genes/proteins? If so, how many? Or is it just 2 - 3 genes but with complexity entirely in how long they are each expressed/when they are turned off, etc.

The Runions model has nearly 100 free parameters. Random walks in 100-dimensional spaces have generic properties like a tendency to move toward regions of larger volume that have nothing to do with leaf biology. How do you disentangle the geometry of high-dimensional random walks from genuinely biological developmental bias? Would a toy model with 100 parameters and arbitrary phenotype categories also show "bias toward simplicity" if "simple" phenotypes occupy more volume?

The discussion of Figure 4 (PCA of parameter space) uses "area" loosely when what's actually being measured is bin count in a 2D projection of a high-dimensional space. I would think that, in general, PCA projections can be misleading about volume in the full parameter space, but I can't tell if that's an issue in this case. Some comments/thoughts here would be useful.

The classifier validation section is in the Methods section, but it seems critical to the whole story. The < 80% agreement with manual classification could propagate to the rest of the estimates in the paper. Again, some comments/thoughts here would be useful.

The authors should explain Mut2 and Mut5 in the main paper with a sentence or two, at least schematically, because how you mutate is obviously very relevant to interpreting a paper about biases in variation.

The two mutational schemes use additive perturbations to individual parameters. Real mutations presumably affect regulatory networks in more structured ways (e.g., changing

binding affinities that affect multiple parameters simultaneously). How sensitive are the results to the assumption of independent single-parameter mutations?

The connectedness argument is made using a 2D PCA projection. Is there a way to check this statement in the full parameter space or perhaps in higher-dimensional projections to test the robustness of this result? Connected components can merge/split under different projections.

<https://doi.org/10.7554/eLife.110692.1.sa1>

Author response:

Reviewer #1 (Public review):

Summary:

The authors aim to understand, in the context of leaf shape, how the constraints imposed by development inform evolution. Leaf shape is a good place to study the influence of development on evolution because it is a trait that exhibits a lot of diversity, and the developmental mechanisms that give rise to leaf shapes are apparently rather conserved across angiosperms.

As part of the motivation for their work, the authors cite a previous study (Geeta et al), which found that in angiosperm phylogenies, transitions from complex to simple leaf shapes occur through evolution more often than transitions in the opposite direction. Is this due to developmental constraints or adaptation?

The authors undertake two parallel lines of work:

(1) Extending the study of Geeta et al with more data, consisting of both phylogenies and a shape classification dataset. The conclusion from this line of inquiry is that transitions from lobed to unlobed leaves are more common than transitions away from unlobed leaves.

(2) The authors conduct evolution simulations in a computational model of leaf development. Here, they look at neutral mutations and whether simply neutral evolution is sufficient to drive the observed trend.

The conclusion of the second part of the work is that the driver of the evolution toward simple leaf shape is entropy: there are more ways to make unlobed leaves than to make lobed leaves (at least in terms of gene regulation parameters that will produce the two leaf types). The argument is that random gene regulatory networks are more likely to produce unlobed leaves than lobed leaves; therefore, neutral evolution drives this trend.

Data Analysis

Roughly 9000 images of leaves were classified into 4 categories: unlobed, lobed, dissected, and compound. These labels were applied to the tips of 5 phylogenetic trees of angiosperms (3 resolved at the genus level and 2 at the species level). By fitting a continuous-time Markov chain to the labelled trees, the authors claim that there is a significantly higher rate of transition to the unlobed leaf shape compared to transitions to more complex shapes.

Simulation

First, the authors validate a computational model (Runions et al) for leaf growth on an experimental dataset. By changing parameters in the model, they can recapitulate the morphological changes in the shapes of Arabidopsis leaves engendered by expression of two particular genes.

Then the authors run an evolutionary model (without selection, just random mutations) on top of the computational leaf development model. As the random walk in parameter space reaches a stationary distribution, they look at both the proportions of the leaf categories in the steady state as well as the transition rates between different categories. The result is that transitions to unlobed leaves are more common than from unlobed leaves.

We thank the reviewer for the helpful and clear summary of our work.

General Comments

The authors use angiosperm phylogenies from other works as the basis for the data analysis part of their work. Given the centrality of these phylogenies for their conclusions, more information is needed about how these phylogenies were constructed and what they mean. What is the timescale that they span? What method is used to infer them? What regions of DNA were sequenced in order to build the phylogenies? Also, maybe some more discussion of angiosperm evolution (e.g., when was the most recent common ancestor of all angiosperms?) would help put the study in context.

We also need a more in-depth discussion of the computational model. What are all the >100 parameters doing, and what informs the seemingly strange mutational model that changes parameters by 3 orders of magnitude?

I am confused about how the rates of transitions were inferred from the phylogeny. Here, one has a phylogeny inferred by some method (which needs to be described in more detail), and just the leaves are labelled. It is stated in the methods that BayesTraits was used to infer the transition rates. I realize this method is probably documented elsewhere, but a bit of a summary of how it works and how to interpret its results would (1) make the paper more self-contained and (2) if the algorithm is credible, make the results firmer.

We thank the referee for the suggestion to make the paper more accessible. The tool we use to infer transition rates from the phylogenies, BayesTraits, is standard in the field. However, the referee is right that for an interdisciplinary journal, it may be helpful to more fully flesh out how these methods work. To that end, we have added an additional section "Phylogenetic rate inference" in the supplementary information that includes a longer description of how BayesTraits works, and how we used it to infer transition rates from phylogenies.

All trees are shown in the supplementary information section "Phylogenetic trees" with scale-bars showing the amount of time or genetic change that the trees span. For a broader discussion of angiosperm evolution, there is supplementary information section "The adaptive significance of leaf shape review".

Regarding the more in-depth discussion of the computational model, we have added supplementary information section S1 "Leaf model details" to give a more detailed description of the leaf model.

I am a bit skeptical of the authors' interpretation of the biological trend (of complex to simple leaf shapes) as being driven by neutral evolution. Why does one expect that the mutations generated by the random walk models described in the work are in fact neutral mutations?

A random walk is a well-established way of modelling the dynamics of neutral evolution in the monomorphic regime, where the population has a narrow diversity of different genotypes. In the higher mutation rate polymorphic regime, where the diversity of genotypes in the population is larger, we also expect that a random walk should still recapitulate the

correct average transition rates. The purpose of the simulations is not to model every aspect of population genetics, but to ask whether developmental bias alone is sufficient to generate the observed directional asymmetry. By assigning equal fitness to all viable leaves, we isolate the contribution of development from that of selection. The agreement with the phylogenetic transition rates therefore demonstrates sufficiency rather than exclusivity: selection may also contribute, but it is not required to explain the observed bias. We discuss the evidence for the role adaptation in leaf shape further in supplementary information section "The adaptive significance of leaf shape review".

If the entropy of simple leaf shapes is higher than that of complex leaf shapes, why did we have complex leaves at all? I suspect the authors might argue that this is due to selection. In that case, what allows these complex shapes to become simpler? Wouldn't they be losing the selective advantage that drove them to be more complex in the first place? Or maybe the idea is that the rates are inferred assuming some steady state that generates the phylogeny? I did not understand this point.

The entropy language is a useful framing. Within that framework, one can view our study as showing that the entropy (defined here as the logarithm of the volume of parameter space mapping to a phenotype) of simple leaf shapes is higher than that of complex leaf shapes. If this entropy were to be ignored, then all states would be equally likely in our simulations, where we do not take fitness differences into account. What we show is that the differences in entropy -- related to differences in volumes of the parameter space that map to different phenotypes -- also affects the rates. The inferred transition rates for both simulation and phylogeny from unlobed to more complex shapes are lower than vice versa but not zero. Therefore, complex leaf shapes arise stochastically through mutation and in this model would eventually reach a steady state proportion, even in the absence of selection.

Are the rates of transitions between leaf types inferred for the phylogeny assuming that the phylogeny is generated by the steady state of some Markov process? (I think the answer is no: in that case, how does one explain the initial condition?)

The tool we use to infer transition rates from phylogenies—BayesTraits—allows the initial state at the root of the tree to vary during the numerical optimisation (Pagel, 1994). Therefore, it is not assumed that the initial state is generated by the steady state of the Markov process.

If I take the mutation model (random walk) seriously, then shouldn't I expect that this steady state obeys detailed balance? In that case I should have $p_i r_{i \rightarrow j} = p_j r_{j \rightarrow i}$ for each of the occupancies p_i and transition rates $r_{i \rightarrow j}$ for the shape categories. How close are the rates inferred from the phylogenies to obeying detailed balance? Presumably, the Markov chain fitted to the simulation data obeys detailed balance because the mutation model itself does?

BayesTraits allows off-diagonal transition rates of the rate matrix to vary freely during numerical optimisation (Pagel, 1994). Therefore, there is no requirement for the detailed balance to hold for the inferred rate matrix. For our simulations, the mutations are symmetric at the parameter level, therefore at this level, the process would be expected to obey the detailed balance.

I find it hard to take the discussion of development seriously without some consideration of mechanics. Presumably, the mechanics are hidden in the computational leaf development model, but this model is not discussed in enough detail for the reader to know. It seems to me that the interesting question is: what are the (not mechanical) constraints on development that drive the apparent trend in evolution towards simpler leaf shapes? Maybe it is something about the type of differential growth needed to make complex leaf shapes less robust to mutation. But in this case, I would assume that

selection plays a role in the complexity of shape. In any case, a better understanding (or explanation) of the computational model is needed to make this interpretation.

We thank the referee for the suggestion to make the paper more accessible. We have added a more detailed and pedagogical description of the model from (Runions, Tsiantis and Prusinkiewicz, 2017) in the supplementary information section S1 "Leaf model details". We also note that Fig. 5 in the methods that gives an overview of how the model works, including some mechanical aspects of development and growth.

More generally, mechanics is one component of the developmental map that determines which parameter combinations produce viable leaf morphologies. Our analysis concerns the geometry of this complete developmental map, irrespective of whether its constraints arise from gene regulation, tissue mechanics, or their interaction.

On the interesting question of what is causal, perhaps the example in figure 2 is helpful. We focus on two parameters, a morphogen repression strength, and a duration of growth. A key physical process here is called webbing, where cellular growth fills in the gaps between branching veins. This process flattens the leaf structure and creates a continuous, solid leaf blade (lamina). Strong webbing, characterized by a significant resistance to stretching and bending, results in a smoother margin (Runions, Tsiantis and Prusinkiewicz, 2017). The morphogen repression strength affects the physical parameters that determine how strong the webbing is. The duration of growth determines how long the leaf has to grow. Varying these two parameters varies the physical processes that determine leaf shape. The mechanics of growth operate downstream of these parameters that we vary in our evolutionary simulations according to the details of the leaf developmental model.

Some discussion of timescales is needed, especially when invoking neutral evolutionary arguments. If a neutral mutation occurs, its time to fix in a population of size N is $\sim N$ generations. What are the relevant angiosperm population sizes and the number of mutations that separate branches on the tree? Are timescales remotely consistent with e.g., the age of angiosperms on Earth?

Neutral processes have a well-established role in key aspects of angiosperm evolution, for example genome complexity (Lynch and Conery, 2003). This would suggest that the relevant time scales and generation times are not completely prohibitive of neutral processes also playing a role in the evolution of angiosperm leaf shape. Effective population sizes in plants are highly variable but estimates span 10^3 - 10^6 . Assuming diploidy (and therefore average fixation time of $4N_e$) and generation times of 1-10 years, this gives fixation timescales of 10^3 - 10^7 years. This is within the timescales of the trees we analyse, which span >150 million years.

Reviewer #2 (Public review):

Strengths:

The paper's underlying question is interesting, extending the authors' prior work on RNA along similar conceptual lines. The paper combines both image analysis of leaves and a computational analysis of a simple model of leaf development.

Weaknesses:

The entire paper is based on the Runion model. More intuition about the Runion model would be useful for a broader readership that cares about the evolutionary aspect of this, but may not know the developmental model in question. Obviously, this is prior well-established work, but 2 - 3 sentences highlighting the key structural aspects of such a model would be great. Currently, that intuition is found implicitly in a sentence on page 2 ("complex leaf shapes need more specificity in their GRNs than their simpler unlobed leaf

shape"), but the reader is left wondering - is the Runion model a detailed mechanistic one with multiple interacting genes/proteins? If so, how many? Or is it just 2 - 3 genes but with complexity entirely in how long they are each expressed/when they are turned off, etc.

We thank the referee for the suggestion to make the paper more useful for a broader readership. To that end, we have added a more detailed description of the (Runions, Tsiantis and Prusinkiewicz, 2017) model in supplementary information section S1 "Leaf model details".

The Runions model has nearly 100 free parameters. Random walks in 100dimensional spaces have generic properties like a tendency to move toward regions of larger volume that have nothing to do with leaf biology. How do you disentangle the geometry of high-dimensional random walks from genuinely biological developmental bias? Would a toy model with 100 parameters and arbitrary phenotype categories also show "bias toward simplicity" if "simple" phenotypes occupy more volume?

Our argument is largely independent of the number of parameters. While it is true that most of the volume is near the surface in a high-dimensional space, our argument is about the relative volumes of the sets of parameters that map to each of the four phenotypes, an entropic argument if you wish. The basic intuition is that a simple phenotype needs fewer parameters to be fine-tuned, and so a larger volume of parameter space will map to a simpler phenotype.

The question about a toy-model with arbitrary phenotypes is helpful, because it allows us to clarify that what we are illustrating here with the biologically realistic example of leaf shapes is a much more generic principle. We can say with confidence that if the toy-model generates a many to one set of outputs (phenotypes) through an algorithmic process whose description length does not grow faster than logarithmically with the size of the genotype space, then it should produce a bias towards simplicity regardless of the number of dimensions, see for example Johnston et al. (2022) and Dingle, Camargo and Louis (2018) for a longer discussion of this more general point which is based on arguments from algorithmic information theory (AIT). We don't use that framing in the current paper because the basic intuition for GRNs that more complex phenotypes need more parameters fine-tuned, and so have relatively smaller volumes, is more straightforward to understand than the more abstract AIT arguments. Our general prediction that this principle should hold more widely for GRNs can be made both by the more formal AIT route, or via the more heuristic fine-tuned parameter route.

The discussion of Figure 4 (PCA of parameter space) uses "area" loosely when what's actually being measured is bin count in a 2D projection of a highdimensional space. I would think that, in general, PCA projections can be misleading about volume in the full parameter space, but I can't tell if that's an issue in this case. Some comments/thoughts here would be useful.

The quantitative estimate of phenotype frequencies is computed directly in the full parameter space and does not depend on PCA. I.e. We estimate that the total volume of viable leaves maps to simple unlobed leaves about 80% of the time. However, the volume is extremely high-dimensional, and so hard to visualise. PCA is used solely to provide an interpretable visualization of this otherwise high-dimensional structure. The PCA plots in Fig 4 and Fig S16 are there to be illustrative, not quantitative. Because the volume differences are large, we do not think that the projections of the main PCA components would be misleading on at least the ordering of the sizes of the parameter space components that map to each leaf shape. We provided a similar analysis for other projections -- PC1-PC6 (supplementary information section "PCA occupancy for higher dimensions"), finding the same trend. To make this point clearer, we have now changed the sentence in the Fig. 4 caption slightly "This

(reveals that --> illustrates how) unlobed leaves occupy a larger region of model parameter space than more complex shapes and that this larger space also contains the majority of more complex leaves.”

The classifier validation section is in the Methods section, but it seems critical to the whole story. The < 80% agreement with manual classification could propagate to the rest of the estimates in the paper. Again, some comments/thoughts here would be useful.

We have repeated the analysis of the agreement between by-eye and automatic morphometric classification. Generating a confusion matrix for the two classification methods shows that the agreement is high for unlobed, dissected and compound, with the main source of disagreement being leaves that were classified as lobed by-eye being classified as either unlobed or dissected by the automatic-morphometric method. The proportion of by-eye lobed leaves classified by the automatic morphometric method as either unlobed (27%) or dissected (23%) is relatively balanced, which we think will help cancel out some error as well. Moreover, we find that the agreement between the automatic-morphometric method and by-eye classification increases to 90.0% when using the categories unlobed and all other categories grouped into one. This is the most important classification for our finding that development and phylogeny are both biased towards unlobed.

The authors should explain Mut2 and Mut5 in the main paper with a sentence or two, at least schematically, because how you mutate is obviously very relevant to interpreting a paper about biases in variation.

In the results section we have added a sentence for more detail on the random walk.

"[We mutated the initial sample using a random walk algorithm with two different mutational schemes, MUT2 (alg. 1) and MUT5 (alg. S2).] These algorithms work by iterating through model parameters one by one and perturbing the value by a small amount. We then [automatically classified the resulting shapes...]"

Moreover, in methods section C there is already a more detailed description of both algorithms.

"MUT2 (alg. 1) iterates through the parameters in a random order, and attempts to change the parameter by a value selected at random from an array of numbers randomly generated at 3 different orders of magnitude. MUT5 (alg. S2) is the same as MUT2 except the value each parameter is multiplied by 10% of the range of that parameter within the initial leaves (fig. S1). The aim here was to provide some way of accounting for the biologically relevant sampling range. "

Moreover, the MUT2 algorithm is described in pseudocode in Algorithm 1 in the main text, and the pseudocode for MUT5 is in supplementary information section S1 C, as algorithm S2.

The two mutational schemes use additive perturbations to individual parameters. Real mutations presumably affect regulatory networks in more structured ways (e.g., changing binding affinities that affect multiple parameters simultaneously). How sensitive are the results to the assumption of independent single-parameter mutations?






The referee raises an interesting and well-known issue concerning this widely studied class of GRN models. Without a detailed understanding of how individual genetic mutations map onto model parameters, it is difficult to determine with confidence whether a mutation would produce correlated changes in certain sets of parameters. Our main argument, however, is that the primary source of the observed bias is geometric: the volume of parameter space (or equivalently, the entropy) corresponding to simple leaf morphologies is substantially larger than that corresponding to complex morphologies. As long as mutations explore parameter space approximately symmetrically, even if they involve correlated

changes in multiple parameters, larger phenotype regions will tend to be encountered more frequently and retained for longer than smaller regions. We therefore expect the observed bias to be robust to many alternative mutation models, although quantifying this robustness is an interesting direction for future work.

The connectedness argument is made using a 2D PCA projection. Is there a way to check this statement in the full parameter space or perhaps in higher dimensional projections to test the robustness of this result? Connected components can merge/split under different projections.

Constructing the nearest neighbour graph for the full dimensional data results in the following no. connected components: unlobed-146, lobed-274, dissected-255, compound-315. This follows the same pattern identified for the PC1-PC2 projection, that unlobed splits into fewer connected components than other leaf shape categories.

References:

- Dingle, K., Camargo, C.Q. and Louis, A.A. (2018) 'Input–output maps are strongly biased towards simple outputs', *Nature Communications*, 9(1), p. 761. Available at: <https://doi.org/10.1038/s41467-018-03101-6> .
- Johnston, I.G. et al. (2022) 'Symmetry and simplicity spontaneously emerge from the algorithmic nature of evolution', *Proceedings of the National Academy of Sciences*, 119(11), p. e2113883119. Available at: <https://doi.org/10.1073/pnas.2113883119> .
- Lynch, M. and Conery, J.S. (2003) 'The Origins of Genome Complexity', *Science*, 302(5649), pp. 1401–1404. Available at: <https://doi.org/10.1126/science.1089370> .
- Pagel, M. (1994) 'Detecting correlated evolution on phylogenies: a general method for the comparative analysis of discrete characters', *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 255(1342), pp. 37–45. Available at: <https://doi.org/10.1098/rspb.1994.0006> .
- Runions, A., Tsiantis, M. and Prusinkiewicz, P. (2017) 'A common developmental program can produce diverse leaf shapes', *New Phytologist*, 216(2), pp. 401–418. Available at: <https://doi.org/10.1111/nph.14449> .
- <https://doi.org/10.7554/eLife.110692.1.sa0>