


Phenotypic Convergence Is Not Mirrored at the Protein Level in a Lizard Adaptive Radiation

Russell B. Corbett-Detig^{*,1,2} Shelbi L. Russell,³ Rasmus Nielsen ^{1,4} and Jonathan Losos⁵

¹Department of Integrative Biology, University of California, Berkeley, Berkeley, CA

²Department of Biomolecular Engineering and Genomics Institute, University of California, Santa Cruz, Santa Cruz, CA

³Department of Molecular, Cellular and Developmental Biology, University of California, Santa Cruz, Santa Cruz, CA

⁴Centre for GeoGenetics, Globe Institute, University of Copenhagen, Copenhagen, Denmark

⁵Department of Biology and Living Earth Collaborative, Washington University, Saint Louis, MO

Associate editor: Harmit Malik

Raw-sequencing reads have been deposited in the NCBI Short Read Trace Archive under PRJNA341850.

*Corresponding author: E-mail: russcd@gmail.com.

Abstract

There are many compelling examples of molecular convergence at individual genes. However, the prevalence and the relative importance of adaptive genome-wide convergence remain largely unknown. Many recent works have reported striking examples of excess genome-wide convergence, but some of these studies have been called into question because of the use of inappropriate null models. Here, we sequenced and compared the genomes of 12 species of anole lizards that have independently converged on suites of adaptive behavioral and morphological traits. Despite extensive searches for a genome-wide signature of molecular convergence, we found no evidence supporting molecular convergence at specific amino acids either at individual genes or at genome-wide comparisons; we also uncovered no evidence supporting an excess of adaptive convergence in the rates of amino acid substitutions within genes. Our findings indicate that comprehensive phenotypic convergence is not mirrored at genome-wide protein-coding levels in anoles, and therefore, that adaptive phenotypic convergence is likely not constrained by the evolution of many specific protein sequences or structures.

Key words: molecular evolution, convergence, *Anolis*.

Introduction

Convergent evolution of species occupying similar selective environments has long been cited as strong evidence of adaptive evolution (Simpson 1955; Schluter 2000). Recently, research has focused on whether the evolution of phenotypic convergence is mirrored at the genetic level (Zhang 2003; Protas et al. 2006; Castoe et al. 2009; Kingsley et al. 2009; Zhen et al. 2012). Several spectacular findings have identified genetic convergence at amino acid residues (Protas et al. 2006; Zhen et al. 2012), and at different sites within the same locus (Zhang 2003; Protas et al. 2006; Kingsley et al. 2009; Zhen et al. 2012; reviewed in detail in Storz 2016). These findings have important implications for our understanding of evolutionary predictability and the extent to which constraints on genetic variation shape the direction of evolutionary change (Wake 1991; Losos 2011; Wake et al. 2011). However, in most studies, loci were selected based on a priori knowledge of functional relationships, which may bias inferences toward genes with strong individual effects and traits with relatively simple genetic bases. It is therefore unknown how commonly convergence in complex phenotypic results in molecular convergence genome-wide.

Numerous genome-wide comparisons have searched for signatures of molecular convergence correlated with phenotypic convergence (Woodard et al. 2011; Parker et al. 2013; Foote et al. 2015; Sackton et al. 2019). However, these studies have many drawbacks. First, most studies are narrowly focused, examining either identical amino acid substitutions (Parker et al. 2013; Foote et al. 2015) or accelerated protein evolution (Woodard et al. 2011; Chikina et al. 2016). Second, often only one convergent trait is examined, which might not be detectable in genome-wide comparisons (Woodard et al. 2011; Parker et al. 2013). Third, few efforts have been made to determine whether a signal of molecular convergence is greater than would be expected by chance (Parker et al. 2013). However, this is a particularly important concern because even random processes will result in some level of molecular convergence (Foote et al. 2015; Thomas and Hahn 2015; Zou and Zhang 2015; Chikina et al. 2016; Sackton et al. 2019).

The adaptive radiation of Caribbean *Anolis* lizards is one of the best-studied examples of phenotypic convergence. Anoles diversified for the most part independently on each island of the Greater Antilles, and produced the same set of habitat specialists, termed ecomorphs, on each island (Losos 2009). Phylogenetic analyses confirm that the members of the




Fig. 1. Phylogeny of *Anolis* species in this study. All nodes have a minimum of 99.8% bootstrap support using both ExaML (Kozlov et al. 2015) and Astral (Mirarab et al. 2014) species tree estimation methods. *Although *A. carolinensis* is native to the mainland United States, it is a recent colonist from Cuba and closely related to, and nearly indistinguishable from trunk-crown anoles native to that island.

same ecomorph from different islands are not closely related (fig. 1, see also, Alföldi et al. 2011). Furthermore, phenotypic convergence in anoles spans a wide array of traits (table 1). Such multivariate convergence likely involves many genetic elements and provides a good system for examining patterns of molecular convergence genome-wide.

To address the important statistical challenges described above and to determine the prevalence of genome-wide molecular convergence, we studied 12 anole species that are convergent across a broad suite of traits, applied statistical tests that explore multiple genetic levels of molecular convergence, and exhaustively compared these results against empirical null models derived from lineages that are not phenotypically convergent. Our results indicate that there is no detectable excess of genome-wide convergence at any scale that we examined including individual protein residues, rates of protein evolution within genes, and genome-wide.

Results and Discussion

Samples, Data Production, and Processing

To study the frequency of molecular convergence in anoles, we selected 12 species, consisting of four representatives of each of three ecomorphs (fig. 1 and table 1), including *Anolis carolinensis* whose genome has been sequenced previously (Alföldi et al. 2011). We sequenced each of the remaining 11 species to a mean coverage of 9× using 150-bp paired-end Illumina reads, and we aligned these data to the *A. carolinensis* genome assembly using a divergence-permissive alignment algorithm (Shrestha and Frith 2013; supplementary table S1, Supplementary Material online). Importantly, although ancestral states are often challenging to confidently identify for *Anolis* ecomorph evolution (Losos 2009), the species that we selected must contain several independent convergences for at least two of the ecomorphs included in our study.

Table 1. Ecomorphs, Species Names, and Morphological and Behavioral Characteristics of Each Ecomorph Included in the Study.

Ecomorph	Species Included	Morphological Characteristics	Behavioral Characteristics
Trunk-crown	<ul style="list-style-type: none">• <i>carolinensis</i>• <i>chlorocyanus</i>• <i>evermanni</i>• <i>A. grahami</i>	<ul style="list-style-type: none">• Green• Short limbs• Long tails• Large toepads• Flattened triangular heads• Sexually dimorphic in body size	<ul style="list-style-type: none">• Perches in canopy• Sit-and-wait and active forager• Consumes fruit/nectaras well as insects
Trunk-ground	<ul style="list-style-type: none">• <i>cratatellus</i>• <i>cybotes</i>• <i>lineatopus</i>• <i>A. sagrei</i>	<ul style="list-style-type: none">• Brown• Long hindlimbs• Long tails• Stocky bodies• Large heads• Small toepads• Sexually dimorphic in body size	<ul style="list-style-type: none">• Perches low on trunks• Sit-and-wait predator• Males actively defend small territories
Twig	<ul style="list-style-type: none">• <i>angusticeps</i>• <i>occultus</i>• <i>insolitus</i>• <i>A. valencienni</i>	<ul style="list-style-type: none">• Gray• Very short limbs• Short tails• Slender bodies• Long, narrow head• Not very sexually dimorphic in body size	<ul style="list-style-type: none">• Perches on the smallest branches• Active foraging predator• Eats small arboreal insects• Relies on crypsis and slow stealthy movements to avoid predators

Because here we are aligning short read data to a moderately diverged reference genome, it is possible that this approach will perform poorly, particularly for the most highly divergent samples. We therefore confirmed that our reference alignment-based approach is sufficient for analyses of these data through extensive simulations. To do this, we introduced increasingly large numbers of divergent sites relative to the reference genome, simulated reads from these divergent genomes, and aligned and analyzed the output as we do for our real data. For the levels of divergence that we observe among the species we sampled, we found that our approach produced unbiased, reliable genotypes (supplementary figs. S1 and S2 and table S2, Supplementary Material online). After applying stringent quality control filters (see Materials and Methods), we extracted and retained 11,747 protein-coding genes for further analysis. Nonetheless, it is possible that some rapidly evolving genes or gene regions are excluded by this approach because short read data were unalignable.

Phylogenetic Analyses

Our first goal was to estimate a robust phylogeny, an essential prerequisite for tests of molecular convergence. We applied both maximum likelihood (Kozlov et al. 2015) and coalescent-based (Mirarab et al. 2014) specie-tree estimation methods, and obtained identical topologies in all analyses (fig. 1). All nodes have >99.8% bootstrap support in both analyses, allowing us to confidently resolve the deepest internal nodes of the *Anolis* phylogeny that have proved challenging for previous studies (Alföldi et al. 2011). Furthermore, we found relatively little evidence of incomplete lineage sorting using Bayesian concordance analyses (Ane et al. 2006; supplementary figs. S3 and S4, Supplementary Material online), consistent with the relatively low reported rates of polymorphism within *Anolis* populations (Tollis et al. 2012). Importantly, given the time scales considered and the modest impact of phylogenetic discordance, most molecular convergence in these species would be driven by recurrent de novo mutation

rather than selection on ancestral standing variation. Additionally, we found that each ecomorph set of species has a similar mean sequencing depth, mean distance to the reference genome, and total branch lengths. Furthermore, none is an outlier relative to any set of four *Anole* lineages selected at random (see below, supplementary fig. S5, Supplementary Material online). We therefore conclude that there is little evidence for lineage-specific bioinformatics biases indicating this phylogeny and data set is appropriate for tests of molecular convergence.

Molecular Convergence at Individual Protein Residues (MC_{site})

To search for molecular convergence, we developed and applied two statistical tests. In the first, MC_{site}, we looked for an excess of identical amino acid substitutions in phenotypically convergent species (some authors distinguish between “parallel” and “convergent” substitution; Zhang 2003; Zhen et al. 2012; Storz 2016; where we refer to these in aggregate as convergent). Specifically, with MC_{site}, we used a simulation procedure that rejects the null hypothesis of no molecular convergence for genes that harbor an excess of convergent mutations within an ecomorph relative to the genome-wide rates of homoplasious substitution (see Materials and Methods). If natural selection is driving molecular convergence in excess to the background rate of neutral homoplasious substitution, then we expect to observe a significantly larger proportion of convergent substitutions within species selected from the same ecomorph class than for comparisons of nonphenotypically convergent lineages (i.e., between species in different ecomorph classes). In comparisons to permuted sets of ecomorphs (fig. 2 and supplementary table S3, Supplementary Material online), we found that there is no evidence for an excess of MC_{site} within any of the three ecomorph classes that we examined (fig. 2 and supplementary fig. S6, Supplementary Material online). Additionally, when we performed traditional branch-site tests for positive

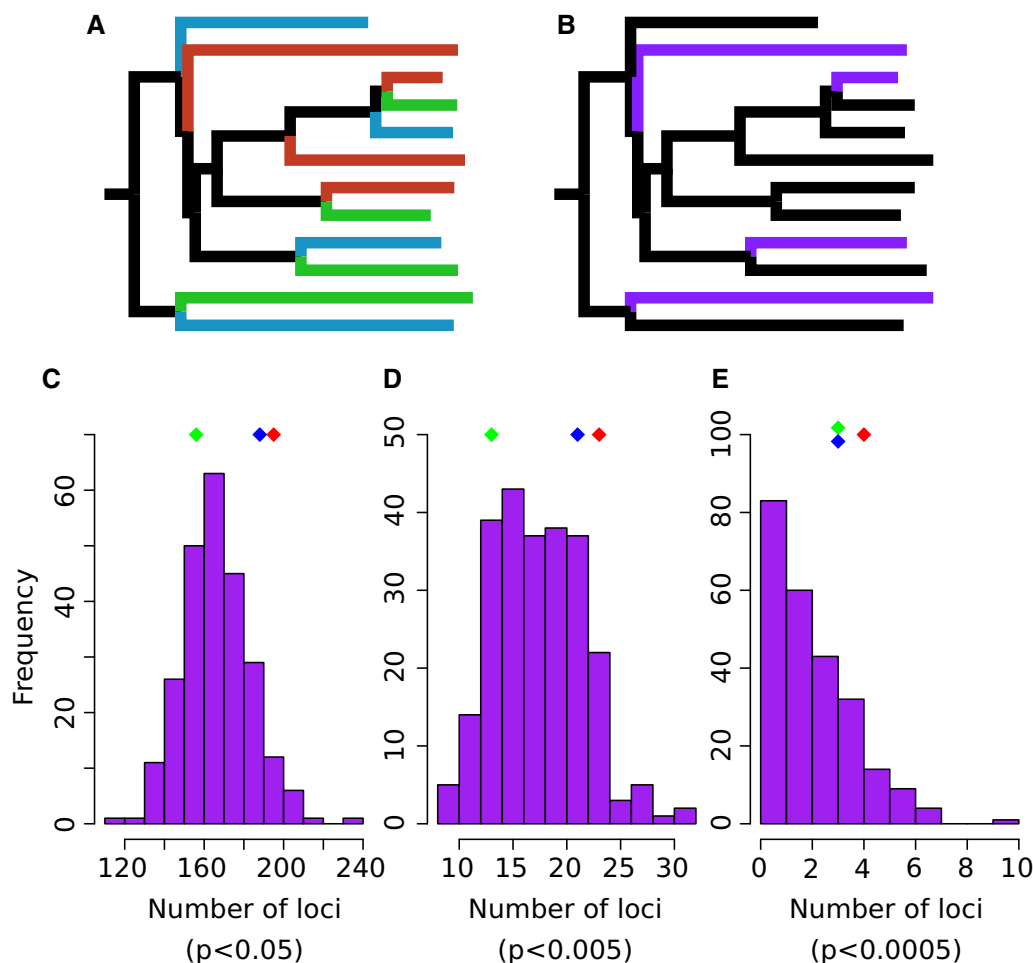


FIG. 2. The molecular convergence hypothesis testing framework developed in this work and its application for detecting amino acid-based convergence MC_{site}. (A) The true sets of Anole lineages used for tests of molecular convergence, and (B) an example permuted set of lineages used to establish the null expectation for the rates of molecular convergence among nonconvergent lineages. The number of genes identified using the MC_{site} simulation test at the uncorrected $P < 0.05$ (C), $P < 0.005$ (D), and $P < 0.0005$ (E). Comparisons are color-coded as follows: trunk-ground (grey), trunk-crown (light grey), twig (dark grey), and permuted sets of lineages are shown in the histogram (black). In all cases shown, the number of genes at a given P value threshold for each ecomorph is fewer than the top 5% tail of the permuted sets of lineages.

selection on terminal lineages, we found no excess of molecular adaptation within ecomorphs relative to permuted sets (supplementary table S4, Supplementary Material online). Nonetheless, we believe that these results may be useful in guiding future functional work in these species.

We next searched for evidence of MC_{site} on a genome-wide scale. Briefly, in this test, we contrasted the numbers of sites where two lineages evolved different amino acid substitutions with the number of sites where two lineages independently evolved identical amino acids (Castoe et al. 2009; Thomas and Hahn 2015). Here again, we find no evidence for excess MC_{site} within an ecomorph ($P = 0.21$, $P = 0.16$, $P = 0.85$ for trunk-crown, trunk-ground, and twig, fig. 3). In combination with the single gene-based tests of MC_{site}, our results suggest that MC_{site} due to selection for phenotypically convergent evolution is infrequent across the genomes of *Anolis* species at genome-wide levels.

Molecular Convergence in Rates of Protein Evolution

To investigate the second potentially fundamental type of molecular convergence, convergence in rates of protein

evolution or MC_{locus}, we next sought for convergence in the rates of protein evolution associated with the ecomorph lineages. Although many authors have compared the rates of protein-coding evolution (i.e., dN/dS ratios; Woodard et al. 2011), this approach does not correct for lineages with an elevated genome-wide relative rate of nonsynonymous substitutions. In particular, we noted that two twig anole lineages display relatively high dN/dS ratios resulting in a signal of MC_{locus} using this uncorrected test (supplementary fig. S7, Supplementary Material online).

We therefore sought to control for the effects of individual lineage branch lengths in our search for MC_{locus}. To do this, we first estimated the rate of amino acid substitution on each lineage for each protein, and then we normalized branch lengths of each lineage by the mean rate of amino acid substitution genome-wide (following Chikina et al. 2016). After applying this correction, we find no evidence for excess MC_{locus} within any of the ecomorph classes considered here (fig. 3 and supplementary fig. S8, Supplementary Material online). This indicates that tests that examine uncorrected branch lengths may yield false positives for MC_{locus}.

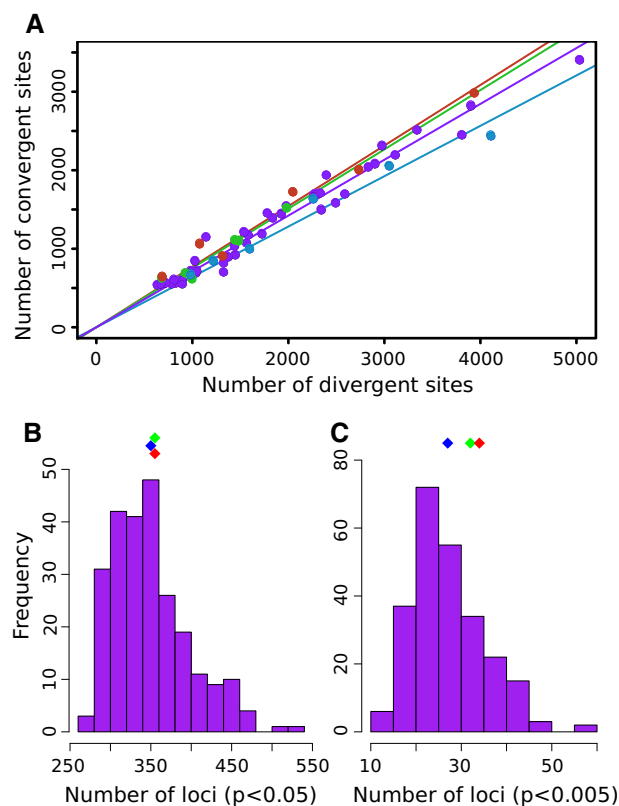


FIG. 3. Genome-wide tests of convergence at amino acids (MC_{site}) and tests of convergent protein evolutionary rates (MC_{locus}). (A) To evaluate evidence for MC_{site} genome-wide, we followed (Castoe et al. 2009; Thomas and Hahn 2015), and we contrasted the numbers of fixations of divergent and convergent amino-acid substitutions. Pairs of different ecomorphs are shown in purple. We removed all sister lineage comparisons and specific cases of excess phylogenetic discordance because ancestral states cannot be reliably inferred. (B and C) The number of genes identified using the MC_{locus} simulation test at the uncorrected $P < 0.05$ (B) and $P < 0.005$ (C). Comparisons are color-coded as follows: trunk-ground (grey), trunk-crown (light grey), twig (dark grey), and permuted sets of lineages are shown in the histogram (black).

However, some permuted sets of lineages show an excess of MC_{locus} even after applying this correction and despite little apparent phenotypic convergence. We note that these groups tend to include lineages for which the nearest node in the phylogeny shows evidence for elevated rates of phylogenetic discordance (supplementary figs. S3 and S4, Supplementary Material online). Therefore, in addition to correcting for mean branch lengths, it may also be necessary to correct for differences among genealogies—which may arise due to, for example, incomplete lineage sorting or introgression after speciation—among the loci examined to generate an accurate null model for molecular convergence (Mendes and Hahn 2016).

Molecular Convergence at Additional Molecular Levels

Our genome-wide analysis of molecular convergence does not support the hypothesis that Anole ecomorph convergence is driven by convergent protein evolution.

Nonetheless, we cannot rule out a contribution of convergent evolution at additional molecular levels. For example, studies of molecular convergence in the evolution of gene expression, and gene family evolution, could further elucidate the degree to which genetic constraints influence the course of adaptive phenotypic convergence (Sackton and Clark 2019; Sackton et al. 2019). Distinguishing between these hypotheses is fundamental to developing a comprehensive understanding of the causes of molecular convergence. This will help to identify the processes of molecular convergence and constraint in shaping recurrent patterns of evolution and should therefore be a primary goal of future analyses. Because they display such a wide array of convergent adaptive phenotypes (table 1), studies of convergence on additional molecular levels in anoles could figure prominently into efforts aimed at disentangling these sophisticated and nuanced hypotheses.

Candidate Convergent Genes

Although we found no evidence consistent with a genome-wide signature of molecular convergence, a small subset of convergent genes might still contribute to convergent phenotypic evolution. We therefore applied gene ontology analyses to identify functionally related genes that are enriched for signatures of MC_{site} and MC_{locus} . Of particular interest, we discovered enrichment for MC_{site} in twig anoles for genes related to neuron and head development (supplementary table S5, Supplementary Material online). This suggests that molecular convergence at a few loci is involved in their derived foraging behaviors or distinct head shape associated with this ecomorph (table 1). This finding is potentially consistent with previous work showing that genes associated with forebrain evolution are accelerated across *Anolis* lineages (Tollis et al. 2018). Additionally, we found enrichment for MC_{locus} in trunk-ground anoles for genes associated with regulation of tissue regeneration (supplementary table S6, Supplementary Material online). This could be consistent with the modified demands associated with a sit-and-wait foraging strategy. We caution that although there are many potentially exciting avenues to follow resulting from these analyses, substantial functional work will be necessary to determine which, if any, molecularly convergent genes are responsible for causing adaptive phenotypic outcomes.

Conclusion

Our central result is that phenotypic convergence in anoles does not predict protein molecular convergence at the levels of individual mutations or rates of evolution, in genome-wide comparisons. Given that strikingly similar ecomorph communities have evolved independently and quickly on each island of the Greater Antilles—essentially whenever opportunity for ecological diversification arises (Losos 2009)—a plausible explanation for the lack of a genome-wide excess of molecular convergence is that there is a vast array of molecular mechanisms through which a lineage can rapidly evolve the characteristic phenotypes associated with a given ecomorph. As many of the characteristic ecomorph traits are quantitative, this explanation is consistent with recent efforts to characterize the genotype–phenotype map in humans, which is

often complex and highly polygenic (Boyle et al. 2017). Seemingly paradoxically, this implies that the exceptional levels of natural replication that make the anole system such an appealing model for studying phenotypic convergence also indicate that phenotypic change may be unlikely to be constrained to evolve genetically in a limited number of ways. Thus, convergence protein-coding sequences at genome-wide scales are unlikely to be required for the evolution of phenotypically convergent anole ecomorphs.

Our data do not support a role of constraints of specific protein-coding mechanisms on genome-wide levels to resolving optimal phenotypic outcomes across the anole radiation. Whether similar results can be found in other lineages that show similarly striking phenotypic convergence is a fundamental and open question (Storz 2016; Sackton and Clark 2019). Using the comparative statistical frameworks that we and others have described, it is now possible to comprehensively search, quantify, and compare the rates of molecular convergence across a broad array of taxonomic groups—an essential goal that is vital in unraveling the relative importance and genomic underpinnings of molecular constraints on phenotypic convergence.

Materials and Methods

Sample Preparation

We selected 11 species that in combination with *A. carolinensis* represent each independent evolution of the three ecomorphs in this study. We extracted DNA from each sample using DNeasy Blood and Tissue kits and following the manufacturer's instructions. DNA was eluted in 200 μ l of pH 8.0 Tris, and stored until library preparation at -80°C . We then prepared the DNA for library preparation by shearing gDNA to a target size of 350 bp using a Covaris S220 sonicator. Libraries were produced using reagents from the NEXTflex Bioo scientific genomic DNA library preparation kit, and using an Apollo 324 automated library preparation protocol. We then performed five PCR cycles on each library.

Sequencing

We multiplexed the 11 samples and sequenced them using a NextSeq 500. In total, we ran four lanes, and sequenced to a target depth of $9\times$ genomic coverage per library using 150-bp paired-end reads. Mean depths and insert size statistics obtained for each library are listed in [supplementary table S1, Supplementary Material](#) online.

Short Read Alignment

We aligned all short read data to the *A. carolinensis* genome reference sequence (Alföldi et al. 2011) using a short read alignment method that is designed for highly divergent sequences (Shrestha and Frith 2013). We used otherwise default parameters, but included the command line options “-Q1 -r5 -q5 -a35 -b5 -i1,” which is the set of parameters that the authors recommend for Illumina short read data in order to make Illumina quality scores. We opted for this short read reference-based alignment approach, rather than a de novo genome assembly for each species for two reasons. First, the

modest sequencing depths in our study, and lack of long read data or mate-pair or linking libraries, preclude our assembling highly contiguous genome assemblies for these complex vertebrate genomes. Second, in our analyses, we aimed only to study protein-coding regions, which tend to be highly conserved and therefore enable a reference-alignment strategy to be effective. After producing our initial alignments, we realigned short reads, removed potential PCR duplicates, and performed genotyping using the Genome-Analysis Toolkit (McKenna et al. 2010) using default parameters except that we required haploid genotypes, that is, using the command line option “-ploidy 1,” when producing final sequences for analysis.

We then extracted consensus sequences for the longest annotated transcript for each gene requiring a minimum quality of 25 and a minimum depth of two reads in annotated coding regions of the *A. carolinensis* genome. Post hoc examination of potentially heterozygous sites (i.e., sites identified as heterozygotes in using standard GATK), indicated that the variant with the highest number of supporting reads is almost always selected if a confident genotype is emitted. Such sites are rare relative to homozygous divergences and it is unlikely that a systematic bias with respect to heterozygous sites has significantly altered our results. However, future work that aims to understand the impacts of molecular convergence in highly polymorphic species should consider the effect of polymorphic sites on downstream analyses. Such positions may have functionally different properties than fixed differences with respect to the evolution of convergent phenotypes. To assist with multiple sequence alignment (below), we included sites that did not pass filters or for which there were no reads aligned in the *A. carolinensis* reference allele in the sequence, and masked those sites after multiple alignment.

Alignment Filtration

Prior to multiple alignment, we applied a number of quality control filters to reduce false positives due to sequence misalignment. Because recent gene duplications may cause one copy to evolve under weaker evolutionary constraint, including reads derived from these structural mutations in our analyses could affect our results. Duplicated genes, or paralogs for which one copy is missing from the reference sequence, would tend to cause an increase in read depth, hence we identified all genes for which the mean sequencing depth exceeded the single-tailed Poisson 95% confidence interval for any species in our study, and we discarded those transcripts. However, it would also be possible for sequences with weak overall homology (e.g., a single highly conserved protein domain in an otherwise highly divergent paralog), to affect our inferences without causing a significant increase in the average depth across the entire alignment.

Multiple Alignment

We produced multiple alignments for each transcript using MAFFT v7.130 (Katoh et al. 2002; Katoh and Standley 2013). We ran this software using default parameters. We then masked all low-quality sites (i.e., those with a phred-scaled

quality score of <20 , corresponding to a nominal error rate of 1 in 100), and removed all codons for which any single base was low quality, or for which no sequence data were available from the sample. We then excluded multiple alignments for which $<75\%$ of codons contained a high-quality alignment for at least 10 (of 12) *Anolis* species. Varying the percent of high-quality codons aligned required to consider an alignment had little effect on the genes considered in our analyses because the distribution of sites covered in ten or more species is bimodal with a peak just $<100\%$ coverage and an additional smaller peak near to 0% coverage.

Short Read Error Rate Estimation

To estimate error rates in short read-sequencing data, we relied on the relatively high depth on the mitochondrion in each of our samples. Due to exceptionally high-sequencing depth and homozygosity of the mitochondrion, it is possible to confidently assign consensus genotypes to the majority of positions in this genome. It is then straightforward to estimate the per site error rate by counting the number of discrepant bases relative to the consensus sequence and dividing this by the total number of bases aligned to the mitochondrion. To do this, we first genotyped the mitochondrion using the GATK under otherwise default parameters, except that we selected the “-ploidy 1” option. We then summed the number of discrepant base pairs at sites with a minimum coverage of 50 or greater and a minimum genotype quality score of 50 or greater, and we then divided this by the total of high-quality aligned base pairs at those sites. Estimated error rates are reported in [supplementary table S1, Supplementary Material online](#), and range from 0.011 to 0.017.

Alignment Simulation

Shrestha and Firth (2013) showed that the alignments produced by their algorithm are highly accurate, even for applications wherein the reference genome is derived from a relatively distantly related species (i.e., Xenomapping). Nonetheless, to confirm that this alignment algorithm will produce reliable results for our specific applications, and to determine if increasing divergence relative to the reference genome was likely to bias the results we obtained, we performed simulations using this alignment algorithm to determine what levels of divergence and error rates would bias the alignment and genotyping outcomes.

To assess this possibility, we simulated data for divergence from the reference sequence between 0.01 and 0.1 per base pair within coding sequences, and at a rate five times greater (i.e., between 0.05 and 0.5) per base pair within the surrounding noncoding regions (e.g., introns, 3′-, and 5′-UTRs). We further included indels into each sequence type at a rate of 0.1 per single nucleotide mutation. Indel lengths were geometrically distributed, with a mean of 15 in noncoding regions, and a mean of 3 in coding regions. The goal was to approximate patterns of nucleotide divergence, and to increase the rates in noncoding regions, which is conservative for the analyses presented in this work. For each level of divergence considered in this analysis, we simulated data using *dwgmsim* v0.1.11 (<https://github.com/nh13/DWGSIM>, last

accessed April 2015), at two uniform per base-pair error rates of 0.01 and 0.02. These divergence levels and error rates were selected to span the ranges we observed in empirical data (0.011–0.017 error/bp, see Short Read Error Rate Estimation). Read lengths and insert size distributions were set to values that match empirical distributions derived from our short read alignments ([supplementary table S1, Supplementary Material online](#)). All other parameters were set to the program defaults.

Alignment of Short Read Data Sets

We aligned short reads to the *A. carolinensis* reference genome (Alföldi et al. 2011) using the same program and parameters as we described above. We then computed the coverage, divergence rates and error rates for each pairwise alignment to the reference genome for each combination of divergence and short read error rates. The results of this analysis are summarized in [supplementary table S2, Supplementary Material online](#). It should be acknowledged that although we have endeavored to be conservative with the assumptions we have made in simulating these data for alignment to the reference, it is infeasible to simulate all of the pertinent aspects of real data.

In simulated data sets, true sequence divergence and estimated sequence divergence are approximately linearly correlated until reaching an asymptote near to a divergence rate of 0.06–0.07 divergent sites/aligned coding bp ([supplementary fig. S1 and table S1, Supplementary Material online](#)). Because all samples considered have a mean divergence from *A. carolinensis* that is less than the point where divergence estimates appear to asymptote (the maximum observed pairwise divergence to *A. carolinensis* is 0.0536 for *A. occultus*, [supplementary table S1, Supplementary Material online](#)), and because our simulation error estimation method is very conservative, we believe that reference bias is unlikely to present a challenge to the results and interpretations we present in this work. However, it is possible that the most rapidly evolving genes cannot be aligned using this approach and full de novo assembly may be preferable in future work.

We note that the raw number of sequencing reads produced is very strongly linearly correlated with the aligned sequencing depth (Pearson's $r = 0.963$, $P = 1.901\text{E-}6$; [supplementary fig. S2 and table S1, Supplementary Material online](#)), which lends additional support to the conclusion that reference bias has a minimal effect. Furthermore, when we performed multiple regressions including both divergence to *A. carolinensis* and the raw number of sequencing reads as predictors of the mean aligned sequence depth, we found that a model including only the raw number of sequencing reads as a predictor for sequencing depth provided an equivalent fit to the data, and is favored because fewer parameters are required for this model ($P = 0.6713$, likelihood ratio test).

Phylogeny and Species Tree Estimation

To produce a species tree, we employed a concatenation-based phylogeny estimation approach, as well as a coalescent method that can account for some aspects of discordance in genealogies between partially linked and unlinked loci. In the

first approach, we used EXaML v2.0.4 (Kozlov et al. 2015), partitioned our loci by gene, and produced 1,000 bootstrap replicates of this analysis. For the coalescent analysis, we first produced individual gene genealogies for each coding sequence separately using RAXML v8.1.5 (Stamatakis 2014) using the GTRGAMMA substitution model and a randomly generated starting tree, but with otherwise default parameters. We then provided the resulting output trees to the program ASTRAL v4.7.5 (Mirarab et al. 2014), and used the program under default conditions, except that we used the exact algorithm as our data set contains a sufficiently small number of taxa to make this calculation feasible. We further quantified uncertainty in the topology produced by ASTRAL by bootstrapping genealogies 1,000 times, and rerunning ASTRAL on these. Both analyses produced identical species tree topologies, and all but one of the nodes has 100% bootstrap support in both analyses. The final node is supported by 99.8% of bootstrap replicates in EXaML, and 100% of bootstrap replicates in the ASTRAL analysis.

Gene Tree Discordance

It is possible that a subset of genealogies is discordant perhaps due to incomplete lineage sorting or hybridization between lineages. To investigate this, we conducted a Bayesian gene tree concordance analysis using Bucky V1.4.4 (Ane et al. 2006; Larget et al. 2010). We included all genealogies for loci with total length >3,000 bp in order to enrich for loci where there would be sufficient phylogenetic signal within the locus to resolve the gene genealogy. However, we obtained similar results using all genealogies. Genealogies were estimated as described above for each locus using RAXML. The concordance tree is identical to that found in the above analyses, and most nodes are supported by the majority of genealogies. However, three nodes, those adjacent to the shortest internal branches, show evidence for significant gene tree discordance (supplementary figs. S3 and S4, Supplementary Material online). Given that most nodes are strongly supported and no two lineages of the same ecomorph are related through poorly supported nodes, it is unlikely the incomplete lineage sorting of functional genetic elements is related to convergent phenotypic evolution of ecomorphs. Furthermore, given that we found no significant excess of convergent loci in the MC_{site} test for ecomorphs, this lends further support to the idea that incomplete lineage sorting is unlikely to contribute to the recurrent phenotypic evolution of *Anolis* ecomorphs.

Ecomorph Permutation

Although it is possible to identify statistically significant evidence of molecular convergence within ecomorphs (see below), this does not necessarily indicate that the molecular convergence is a result of convergent phenotypic evolution. To provide support for that assertion, one must show that there is an excess of molecular convergence in lineages that are phenotypically convergent relative to lineages with no such phenotypic convergence (Thomas and Hahn 2015; Zou and Zhang 2015; Chikina et al. 2016). We therefore created permuted sets of ecomorphs, so that we could establish the null expectation for molecular convergence in the

absence of phenotypic convergence. To do this, we identified all combinations of four lineages, and we excluded those sets that contain three or more lineages of the same ecomorph, and those that contained sister lineages.

We then analyzed all 246 possible permuted sets of ecomorphs, against which we can test for significance of molecular convergence. We note that leaf nodes are not exchangeable because of differences in branch lengths between branches in the tree and the underlying correlation structure imposed by the tree. If the branches leading to a specific ecomorph have different properties, for example, different lengths, elevated dN/dS ratios, or an excess of phylogenetic discordance at the preceding node, than random branches in the tree, this would violate the exchangeability assumption underlying permutation tests, hence some caution is warranted in interpreting the results of analyses based on comparisons to permuted ecomorphs.

Molecular Convergence at Individual Sites (MC_{site})

We developed a simulation test, which tests the null hypothesis of no enrichment of convergent identical amino acid mutations within a single protein-coding transcript. We used PAML (specifics for these analyses are described below) to provide maximum likelihood estimates of the ancestral codon states at each site in each node. From these estimates, we identified pairs of lineages with identical amino acid changes at the same codon position within a protein. We define “a convergence” as a position in which two or more lineages have experienced a mutation that generated identical amino acids at the same position. A “k-lineage convergence” is a convergence in which *k* lineages leading to leaf nodes have experienced the exact same changes. All 2-lineage convergences can be divided into intra- or intermorph convergences depending on whether the two leaf nodes associated with the convergences are from the same or from different ecomorphs (e.g., twig anoles), respectively. A *k*-lineage convergence may involve both intra- and intermorph convergences if *k* > 2. The test statistic we use is defined, for gene *g*, as follows:

$$t_g = \sum_{j=2}^n \binom{j}{2} w_{jg},$$

where w_{jg} is the number of intramorph *j*-lineage convergences in gene *g*. From the genome-wide data, we estimate the conditional probability that a *k*-lineage convergence with *j* intramorph convergences as follows:

$$\hat{p}_{kj} = \frac{\sum_{i=1}^S w_{kji}}{\sum_{i=1}^S c_{ki}},$$

where c_{ki} is the total number of *k*-lineage convergences in gene *i*, w_{kji} is the number of *k*-lineage convergences with *j* intramorph convergences in gene *i*, and *S* is the total number of genes in the genome. We then simulate new data for each gene, keeping the total number of *k*-lineage convergences in

the gene fixed, but assigning each convergence to include j intramorph with probability \hat{p}_{kj} . This procedure is repeated 10^6 times, and the resulting simulated data are compared with the observed data to produce a P value in a one-sided test for each gene.

This test rejects the null model if there are more intramorph convergences than expected given the average ratio of intramorph to intermorph convergences, and given the total number of convergences in the gene. This test is robust to multiple factors that otherwise have the potential to create artifactual evidence of selection for molecular convergence. In particular, there are likely many positions in the genome that only allow a few specific amino acids. Such positions will tend neutrally to switch back and forth between a small number of states. The patterns observed in such sites would mimic the pattern of adaptive convergent evolution inferred by naïve methods that do not use apparent convergences between pairs of different ecomorphs as a built-in control. For the purposes of this analysis, we limited our inferences to codons wherein all lineages have a high-quality aligned codon, and proteins for which at least five sites are convergent between two or more lineages.

To determine whether the observed convergences may be attributed to phenotypic convergence, we then applied the same test to the 246 permuted sets of ecomorphs (see above) for each locus. We compared the P value distributions for each ecomorph with those obtained from permuted ecomorph sets, which should be more robust than assuming a uniform P value distribution. There was no excess of significant tests in any ecomorph relative to permuted sets (supplementary fig. S6, Supplementary Material online).

MC_{site} with Amino Acid Biochemical Properties

We further extended this test to include a broader range of convergent amino acid mutations. That is, for each amino acid mutation, we assigned it to a group based on the amino acid classification system employed in Zhang (2000), and we recoded all amino acid positions with their corresponding amino acid group numbers. The idea was to broaden the above test to include mutations to biochemically similar amino acids without requiring strictly identical amino acid substitutions. We then looked for instances wherein two lineages independently converged on the same amino acid group at a site, and apply the same simulation procedure as described above, where again, we are substituting amino acid groups for specific amino acid residues. Here, we do not observe a significant enrichment for convergent radical amino acid class substitutions within morphologically convergent lineages both at individual loci, and the genome-wide MC_{site} analysis. We therefore conclude that neither a genome-wide excess of narrow-sense nor a broader-sense MC_{site} is supported by the data we have obtained.

MC_{site} Genome-Wide

In the absence of selection increasing molecular convergence, the number of convergent substitutions is expected to increase as the number of divergent substitutions increases between two lineages (Castoe et al. 2009; Thomas and

Hahn 2015). To assess the impact of selection in creating MC_{site} genome-wide, we contrasted the number of divergent mutations (i.e., sites where two lineages each fixed an amino acid changing mutation) with the number of convergent mutations between all pairwise combinations of lineages genome-wide. Because the mutation rates for these two classes of sites are approximately equal, they are expected to show a linear, positive correlation. We therefore fit linear regression models to all pairwise observations within each ecomorph class separately, and we then fit those same models for sets of six pairs of lineages drawn randomly without replacement and we repeated this procedure 100 times and again discarded comparisons that included more than two comparisons between taxa of the same ecomorph. We again excluded all pairs of sister lineages and pairwise comparisons between lineages separated by nodes displaying an excess of phylogenetic discordance (i.e., concordance <0.9, supplementary fig. S3, Supplementary Material online) and that deviate substantially from the linear relationship obtained for the other taxa. Specifically, we removed *A. chlorocyanus*–*A. insolitus*, *A. lineatopus*–*A. valencienni*, and *A. grahami*–*A. valencienni*, which are the data points with the largest relative residuals in our linear regression model (below, 0.49, 0.82, and 0.75, respectively). However, errors in reconstructed ancestral states might impact other pairwise comparisons as well.

We then fit linear regression models to all pairwise observations within each ecomorph class separately forcing the intercept through the origin. We fit those same models for sets of six pairs of lineages drawn randomly without replacement and we repeated this procedure 100 times. The resulting P value for each ecomorph reflects the proportion of permutations whose slope exceeds what we obtained for the set of lineages of each ecomorph.

Molecular Convergence in Rates of Evolution at a Locus (MC_{locus})

To broaden our search for molecular convergence, we further implemented an analysis to search for transcripts in which the rate of amino acid substitution is greater in association with phenotypic convergence than expected by chance. This analysis differs from MC_{site} tests (above), in that it aims at identifying correlations between the rate of amino acid substitution and phenotypic convergence, but does not require that mutations produce convergent amino acids. Here, we employ a branch test framework in PAML v4.8 (Yang 2007). In the null model, dN/dS between all terminal branches is equal. In the second model, the alternative hypothesis, we fit a model wherein all branches corresponding to one ecomorph evolve a different dN/dS ratio from all of the other ecomorph lineages (see Woodard et al. 2011 for a related approach). Because the ancestral ecomorph states are unknown and because understanding their evolution is not a specific goal of this analysis, all internal branches are treated as a third type in both analyses. Because these models are nested, we compare them using a likelihood ratio test with one degree of freedom.

Although we find some evidence supporting an excess of MC_{locus} within the twig ecomorph class (supplementary fig. S7, Supplementary Material online), we were concerned that the elevated dN/dS ratios of the two longest twig lineages might impact the test statistic. We therefore applied an additional test of MC_{locus} (Chikina et al. 2016). Briefly, this was done by first computing the mean rate of amino acid substitution genome-wide for each branch. Then, for each locus, we factored out the genome-wide mean rate from the locus-specific rates of evolution, the resulting rate vector correspond to the “relative rates” of amino acid substitution at a given locus. To test for evidence of MC_{locus} at each gene, we used Wilcoxon rank-sum test to ask if a given locus showed evidence for increased or decreased relative rates in association with each ecomorph (Chikina et al. 2016). We further applied this test to all permuted sets to ensure we completely captured the null distribution for this statistic in the absence of phenotypic convergence.

For each ecomorph and each test, we assessed significance based on whether the number of significant genes at a given P value threshold was in the top 5% of the empirical distribution obtained from the permuted sets. Although we find no evidence for MC_{locus} or MC_{site} in either individual loci or genome-wide, it should be acknowledged that a low rate of molecular convergence might not be detected by this analysis, as infrequent adaptive convergences could still be evolutionarily significant, but difficult to distinguish from the background rate of homoplasious substitution.

Branch-Site Tests

Although not a test of MC_{site} per se, we further performed traditional branch-site tests of positive selection within ecomorph lineages using PAML. Here, the null model is one in which a portion of sites within an alignment evolve neutrally (i.e., dN/dS = 1), and the remaining sites experience negative selection (dN/dS < 1). In the alternate model, rather than neutrally, a subset of sites evolves at increased rates (i.e., dN/dS > 1). We therefore fit two models for each ecomorph and compare the models again using a likelihood ratio test. After FDR correction, we identified 80, 160, and 127 genes that show significant evidence for positive selection at the $q = 0.05$ level for trunk-crown, trunk-ground, and twig ecomorphs, respectively. However, we obtained similar results in the first 100 permuted sets of ecomorphs (range 68–201), suggesting that there is little enrichment for significant branch-site tests associated with ecomorphs specifically.

Exploring Bioinformatic Artifacts

We computed a number of quality control metrics to ask if any or all of the lineages of each ecomorph are outliers. Briefly, we computed the mean branch lengths, the mean depth of permuted sets, and the mean genetic distance to *A. carolinensis* (supplementary fig. S5, Supplementary Material online). In all cases, the true ecomorphs are within the 90% empirical confidence interval obtained for the permuted sets suggesting that these factors have little impact on our analyses and that if there are inadvertent impacts, the effect is captured by our permuted sets of lineages as well.

Gene Ontology Analyses

To explore more subtle signatures of molecular convergence that might drive phenotypic convergence but fail to exceed genome-wide expectations, we performed exploratory gene ontology (GO) analyses. To do this, we used the Panther GO analysis framework (Mi et al. 2010), where we included the top 5% of genes, ranked by P value, for each test and assessed enrichment relative to the background set of genes which included all loci included in a given analysis.

Supplementary Material

Supplementary data are available at *Molecular Biology and Evolution* online.

Acknowledgments

The authors thank Tom Sanger, Adam Freedman, Rich Glor, and Michelle Johnson for providing samples, and members of the Corbett-Detig, Nielsen, Hartl, and Losos labs for providing insights. This work was supported in part by an Alfred P. Sloan award to R.B.C.-D.

Author Contributions

R.B.C.-D. and S.L.R. produced sequencing libraries. R.B.C.-D. analyzed the data. R.B.C.-D., J.B.L., and R.N. wrote the article. R.B.C.-D. and J.B.L. conceived and designed the study.

References

- Alföldi J, Di Palma F, Grabherr M, Williams C, Kong L, Mauceli E, Russell P, Lowe CB, Glor RE, Jaffe JD, et al. 2011. The genome of the green anole lizard and a comparative analysis with birds and mammals. *Nature* 477(7366):587–591.
- Ane C, Larget B, Baum DA, Smith SD, Rokas A. 2006. Bayesian estimation of concordance among gene trees. *Mol Biol Evol.* 24(2):412–426.
- Boyle EA, Li YI, Pritchard JK. 2017. An expanded view of complex traits: from polygenic to omnigenic. *Cell* 169(7):1177–1186.
- Castoe TA, de Koning APJ, Kim H-M, Gu W, Noonan BP, Naylor G, Jiang ZJ, Parkinson CL, Pollock DD. 2009. Evidence for an ancient adaptive episode of convergent molecular evolution. *Proc Natl Acad Sci U S A*. 106(22):8986–8991.
- Chikina M, Robinson JD, Clark NL. 2016. Hundreds of genes experienced convergent shifts in selective pressure in marine mammals. *Mol Biol Evol.* 33(9):2182–2192.
- Foot AD, Liu Y, Thomas GWC, Vinař T, Alföldi J, Deng J, Dugan S, van Elk CE, Hunter ME, Joshi V, et al. 2015. Convergent evolution of the genomes of marine mammals. *Nat Genet.* 47(3):272–275.
- Katoh K, Misawa K, Kuma K, Miyata T. 2002. MAFFT: a novel method for rapid multiple sequence alignment based on fast Fourier transform. *Nucleic Acids Res.* 30(14):3059–3066.
- Katoh K, Standley DM. 2013. MAFFT: multiple sequence alignment software version 7: improvements in performance and usability. *Mol Biol Evol.* 30(4):772–780.
- Kingsley EP, Manceau M, Wiley CD, Hoekstra HE. 2009. Melanism in *Peromyscus* is caused by independent mutations in *Agouti*. *PLoS One* 4(7):e6435–e6439.
- Kozlov AM, Aberer AJ, Stamatakis A. 2015. ExaML version 3: a tool for phylogenomic analyses on supercomputers. *Bioinformatics* 31:2577–2579.
- Larget BR, Kotha SK, Dewey CN, Ane C. 2010. BUCKy: gene tree/species tree reconciliation with Bayesian concordance analysis. *Bioinformatics* 26(22):2910–2911.
- Losos JB. 2009. Lizards in an evolutionary tree: ecology and adaptive radiation of anoles. Berkeley (CA): University of California Press.

- Losos JB. 2011. Convergence, adaptation, and constraint. *Evolution* 65(7):1827–1840.
- McKenna A, Hanna M, Banks E, Sivachenko A, Cibulskis K, Kernytzky A, Garimella K, Altshuler D, Gabriel S, Daly M, et al. 2010. The Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. *Genome Res.* 20(9):1297–1303.
- Mendes FK, Hahn MW. 2016. Gene tree discordance causes apparent substitution rate variation. *Syst Biol.* 65(4):711–721.
- Mi H, Dong Q, Muruganujan A, Gaudet P, Lewis S, Thomas PD. 2010. PANTHER version 7: improved phylogenetic trees, orthologs and collaboration with the Gene Ontology Consortium. *Nucleic Acids Res.* 38(Database issue):D204–D210.
- Mirarab S, Reaz R, Bayzid MS, Zimmermann T, Warnow T. 2014. ASTRAL: genome-scale coalescent-based species tree estimation. *Science* 30:541–548.
- Parker J, Tsagkogeorga G, Cotton JA, Liu Y, Provero P, Stupka E, Rossiter SJ. 2013. Genome-wide signatures of convergent evolution in echolocating mammals. *Nature* 502(7470):228–231.
- Protas ME, Hersey C, Kochanek D, Zhou Y, Wilkens H, Jeffery WR, Zon LI, Borowsky R, Tabin CJ. 2006. Genetic analysis of cavefish reveals molecular convergence in the evolution of albinism. *Nat Genet.* 38(1):107–111.
- Sackton TB, Clark N. 2019. Convergent evolution in the genomics era: new insights and directions. *Philos Trans R Soc B.* 374(1777):20190102.
- Sackton TB, Grayson P, Cloutier A, Hu Z, Liu JS, Wheeler NE, Gardner PP, Clarke JA, Baker AJ, Clamp M, et al. 2019. Convergent regulatory evolution and loss of flight in paleognathous birds. *Science* 364(6435):74–78.
- Schluter D. 2000. The ecology of adaptive radiation. Oxford: Oxford University Press.
- Shrestha A, Frith MC. 2013. An approximate Bayesian approach for mapping paired-end DNA reads to a reference genome. *Bioinformatics* 29(8):965–972.
- Simpson GG. 1955. Major features of evolution. New York: Columbia University Press.
- Stamatakis A. 2014. RAxML version 8: a tool for phylogenetic analysis and post-analysis of large phylogenies. *Bioinformatics* 30(9):1312–1313.
- Storz JF. 2016. Causes of molecular convergence and parallelism in protein evolution. *Nat Rev Genet.* 17(4):239–250.
- Thomas GWC, Hahn MW. 2015. Determining the null model for detecting adaptive convergence from genomic data: a case study using echolocating mammals. *Mol Biol Evol.* 32(5):1232–1236.
- Tollis M, Ausubel G, Ghimire D, Boissinot S. 2012. Multi-locus phylogeographic and population genetic analysis of *Anolis carolinensis*: historical demography of a genomic model species. *PLoS One* 7(6):e38474.
- Tollis M, Hutchins ED, Stapley J, Rupp SM, Eckalbar WL, Maayan I, Lasku E, Infante CR, Dennis SR, Robertson JA, et al. 2018. Comparative genomics reveals accelerated evolution in conserved pathways during the diversification of anole lizards. *Genome Biol Evol.* 10(2):489–506.
- Wake DB, Wake MH, Specht CD. 2011. Homoplasy: from detecting pattern to determining process and mechanism of evolution. *Science* 331:1032–1035.
- Wake DB. 1991. Homoplasy: the result of natural selection, or evidence of design limitations? *Am Nat.* 138(3):543–567.
- Woodard SH, Fischman BJ, Venkat A, Hudson ME, Varala K, Cameron SA, Clark AG, Robinson GE. 2011. Genes involved in convergent evolution of eusociality in bees. *Proc Natl Acad Sci U S A.* 108(18):7472–7477.
- Yang Z. 2007. PAML 4: phylogenetic analysis by maximum likelihood. *Mol Biol Evol.* 24(8):1586–1591.
- Zhang J. 2000. Rates of conservative and radical nonsynonymous nucleotide substitutions in mammalian nuclear genes. *J Mol Evol.* 50(1):56–68.
- Zhang J. 2003. Parallel functional changes in the digestive RNases of ruminants and colobines by divergent amino acid substitutions. *Mol Biol Evol.* 20(8):1310–1317.
- Zhen Y, Aardema ML, Medina EM, Schumer M, Andolfatto P. 2012. Parallel molecular evolution in an herbivore community. *Science* 337:1634–1637.
- Zou Z, Zhang J. 2015. No genome-wide protein sequence convergence for echolocation. *Mol Biol Evol.* 32(5):1237–1241.