A Statistical Pipeline for Identifying Physical Features that Differentiate Classes of 3D Shapes

Bruce Wang^{1,2,*}, Timothy Sudijono^{3,*}, Henry Kirveslahti^{4,*}, Tingran Gao⁵, Douglas M. Boyer⁶, Sayan Mukherjee^{4,7-9}, and Lorin Crawford^{10-12,†}

1 Data Science Initiative, Brown University, Providence, RI, USA

2 Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, NJ, USA

3 Division of Applied Mathematics, Brown University, Providence, RI, USA

4 Department of Statistical Science, Duke University, Durham, NC, USA

5 Committee on Computational and Applied Mathematics, Department of Statistics, University of Chicago, Chicago, IL, USA

6 Department of Evolutionary Anthropology, Duke University, Durham, NC, USA

7 Department of Computer Science, Duke University, Durham, NC, USA

8 Department of Mathematics, Duke University, Durham, NC, USA

9 Department of Bioinformatics & Biostatistics, Duke University, Durham, NC, USA

10 Department of Biostatistics, Brown University, Providence, RI, USA

11 Center for Statistical Sciences, Brown University, Providence, RI, USA

12 Center for Computational Molecular Biology, Brown University, Providence, RI, USA

* Authors Contributed Equally

† Corresponding E-mail: lorin_crawford@brown.edu

Abstract

The recent curation of large-scale databases with 3D surface scans of shapes has motivated the development of tools that better detect global-patterns in morphological variation. Studies which focus on identifying differences between shapes have been limited to simple pairwise comparisons and rely on pre-specified landmarks (that are often known). We present SINATRA: the first statistical pipeline for analyzing collections of shapes without requiring any correspondences. Our novel algorithm takes in two classes of shapes and highlights the physical features that best describe the variation between them. We use a rigorous simulation framework to assess our approach. Lastly, as a case study, we use SINATRA to analyze mandibular molars from four different suborders of primates and demonstrate its ability recover known morphometric variation across phylogenies.

Introduction

Sub-image analysis is an important open problem in both medical imaging studies and geometric morphometric applications. The problem asks which physical features of shapes are most important for differentiating between two classes of 3D images or shapes such as computed tomography (CT) scans of bones or magnetic resonance images (MRI) of different tissues. More generally, the sub-image analysis problem can be framed as a regression-based task: given a collection of shapes, find the properties that explain the greatest variation in some response variable (continuous or binary). One example is identifying the structures of glioblastoma tumors that best indicate signs of potential relapse and other clinical outcomes [1]. From a statistical perspective, the sub-image selection problem is directly related to the variable selection problem — given high-dimensional covariates and a univariate outcome, we want to infer which variables are most relevant in explaining or predicting variation in the observed response.

Framing sub-image analysis as a regression presents several challenges. The first challenge centers around representing a 3D object as a (square integrable) covariate or feature vector. The transformation should lose a minimum amount of geometric information and apply to a wide range of shape and imaging datasets. In this paper, we use a tool from integral geometry and differential topology called the Euler characteristic (EC) transform [1–4], which maps shapes into vectors without requiring pre-specified landmark points or pairwise correspondences. This property is central to our innovations.

After finding a vector representation of the shape, our second challenge is quantifying which topological features are most relevant in explaining variation in a continuous outcome or binary class label. We address this classic take on variable selection by using a Bayesian regression model and an information theoretic metric to measure the relevance of each topological feature. Our Bayesian method allows us to perform variable selection for nonlinear functions — we discuss the importance of this requirement in Results and Methods.

The last challenge deals with how to interpret the most informative topological features obtained by our variable selection methodology. The EC transform is invertible; thus, we can take the most informative topological features and naturally recover the most informative physical regions on the shape. In this paper, we introduce SINATRA: a unified statistical pipeline for sub-image analysis that addresses each of these challenges and is the first sub-image analysis method that does not require landmarks or correspondences.

Classically there have been three approaches to modeling random 3D images and shapes: (i) landmark-based representations [5], (ii) diffeomorphism-based representations [6], and (iii) representations that use integral geometry and excursions of random fields [7]. Landmark-based analysis uses points on shapes that are known to correspond with each other. As a result, any shape can be represented as a collection of 3D coordinates. Landmark-based approaches have two major shortcomings. First, many modern datasets are not defined by landmarks; instead, they consist of 3D CT scans [8,9]. Second, reducing these detailed mesh data to simple landmarks often results in a great deal of information loss.

Diffeomorphism-based approaches have bypassed the need for landmarks. Many tools have been developed that efficiently compare the similarity between shapes in large databases via algorithms that continuously deform one shape into another [10–14]. Unfortunately, these methods require diffeomorphisms between shapes: the map from shape A to shape B must be differentiable, as must the inverse of the map. Such functions are often called “correspondence maps” since they take two shapes and place them in correspondence. There are many applications with no such transformations because of qualitative differences. For example, in a dataset of fruit fly wings, some mutants may have extra lobes of veins [15]; or, in a dataset of brain arteries, many of the arteries cannot be continuously mapped to each other [16]. Indeed, in large databases such as the MorphoSource [9], the CT scans of skulls across many clades are not diffeomorphic. Thus, there is a real need for 3D image analysis methods that do not require correspondences.

Previous work [2] introduced two topological transformations for shapes: the persistent homology (PH) transform and the EC transform. These tools from integral geometry first allowed for pairwise comparisons between shapes or images without requiring correspondence or landmarks. Since then, mathematical foundations of the two transforms and their relationship to the theory of sheaves and fiber bundles have been established [3, 4]. Detailed mathematical analyses have also been provided [3]. A nonlinear regression framework which uses the EC transform to predict outcomes of disease free survival in glioblastoma [1] is most relevant to this paper. This work shows that the EC transform reduces the problem of regression with shape covariates into a problem in functional data analysis (FDA), and that nonlinear regression models are more accurate than linear models when predicting complex phenotypes and traits. The SINATRA pipeline further enhances the relation between FDA and topological transforms by enabling variable selection with shapes as covariates.

Beyond the pipeline, this paper includes software packaging to implement our approach and a detailed design of rigorous simulation studies which can assess the accuracy of sub-image selection methods. The

freely available software comes with several built-in capabilities that are integral to sub-image analyses in both biomedical studies and geometric morphometric applications. First, SINATRA does not require landmarks or correspondences in the data. Second, given a dataset of normalized and axis-aligned 3D images, SINATRA outputs evidence measures that highlight the physical regions on shapes that explain the greatest variation between two classes. In many applications, users may suspect *a priori* that certain landmarks have greater variation across groups of shapes (e.g. via the literature). To this end, SINATRA also provides p-values and Bayes factors that detail how likely any region is identified by chance [17].

In this paper, we describe each mathematical step of the SINATRA pipeline and demonstrate its power and utility via simulations. We also use a dataset of mandibular molars from four different genera of primates to show that our method has the ability to (i) further understanding of how landmarks vary across evolutionary scales in morphology and (ii) visually detail how known anatomical aberrations are associated to specific disease classes and/or case-control studies.

Results

Method Pipeline Overview

The SINATRA pipeline implements four key steps (Fig. 1). First, SINATRA summarizes the geometry of 3D shapes (represented as triangular meshes) by a collection of vectors (or curves) that encode changes in their topology. Second, a nonlinear Gaussian process model, with the topological summaries as input, classifies the shapes. Third, an effect size analog and corresponding association metric is computed for each topological feature used in the classification model. These quantities provide evidence that a given topological feature is associated with a particular class. Fourth, the pipeline iteratively maps the topological features back onto the original shapes (in rank order according to their association measures) via a reconstruction algorithm. This highlights the physical (spatial) locations that best explain the variation between the two groups. Details of our implementation choices are detailed in Methods, with theoretical support given in Supplementary Note.

Algorithmic Overview and Implementation

To facilitate analyses, software for implementing the SINATRA pipeline is carried out in R code and is freely available at <https://github.com/lcrawlab/SINATRA>. This algorithm requires these inputs:

- axis aligned shapes represented as meshes;
- $\mathbf{y}$, a binary vector denoting shape classes;
- r , the radius of the bounding sphere for the shapes (which we usually set to 1/2 since we work with meshes normalized to the unit ball);
- c , the number of cones of directions;
- d , the number of directions within each cone;
- θ , the cap radius used to generate directions in a cone;
- l , the number of sublevel sets (i.e. filtration steps) to compute the Euler characteristic (EC) along a given direction.

In the next two sections, we discuss strategies for choosing values for the free parameters through simulation studies. A table detailing scalability for the current algorithmic implementation can be found in Supplementary Note (see Supplementary Table 1).

Simulation Study: Perturbed Spheres

We begin with a proof-of-concept simulation study to demonstrate both the power of our proposed pipeline and how different parameter value choices affect its ability to detect associated features on 3D shapes. We take 100 spheres and perturb regions, or collections of vertices, on their surfaces to create two classes with a two-step procedure:

1. We generate a fixed number of (approximately) equidistributed regions on each sphere: some number u regions to be shared across classes, and the remaining v regions to be unique to class assignment.
2. To create each region, we perturb the k closest vertices $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k\}$ by a pre-specified scale factor α and add some random normally distributed noise $\epsilon_i \sim \mathcal{N}(0, 1)$ by setting $\mathbf{x}_i^* := \mathbf{x}_i \alpha + \epsilon_i$ for $i = 1, \dots, k$.

We consider three scenarios based on the number of shared and unique regions between shape classes (Figs. 2(a)-2(c)). We choose $(u, v) = (2, 1)$ (scenario I), $(6, 3)$ (scenario II), and $(10, 5)$ (scenario III), and set all regions to be $k = 10$ vertices.

Each sequential scenario represents an increase in degree of difficulty, because class-specific regions should be harder to identify in shapes with more complex structures. We analyze fifty unique simulated datasets for each scenario. In each dataset, only the v -region vertices used to create class-specific regions are defined as true positives, and we quantify SINATRA's ability to prioritize these true vertices using receiver operating characteristic (ROC) curves plotting true positive rates (TPR) against false positive rates (FPR) (Supplementary Note Section 2). We then evaluate SINATRA's power as a function of its free parameter inputs: c number of cones, d number of directions per cone, direction generating cap radius θ , and l number of sublevel sets per filtration. We iteratively vary each parameter, while holding the others as constants $\{c = 25, d = 5, \theta = 0.15, l = 30\}$. Figures displayed in the main text are based on varying the number of cones (Figs. 2(d)-2(f)), while results for the other sensitivity analyses can be found in Supplementary Note (Supplementary Figs. 1-3).

As expected, SINATRA's performance is consistently better when shapes are defined by a few prominent regions (e.g. scenario I) versus when shape definitions are more complex (e.g. scenarios II and III), because each associated vertex makes a greater individual contribution to the overall variance between classes (i.e. $\mathbb{V}(\mathbf{y})/10 > \mathbb{V}(\mathbf{y})/30 > \mathbb{V}(\mathbf{y})/50$). Similar trends in performance have been shown during the assessment of high-dimensional variable selection methods in other application areas [18–20].

This simulation study also demonstrates the general behavior and effectiveness of the SINATRA algorithm as a function of different choices for its free input parameters. First, we assess how adjusting the number of cones of directions used to compute Euler characteristic curves changes power. Computing topological summary statistics over just a single cone of directions (i.e. $c = 1$) is ineffective at capturing enough variation to identify class-specific regions (Figs. 2(d)-2(f)), which supports the intuition that seeing more of a shape leads to an improved ability to understand its complete structure [1–3]. Our empirical results show that more power can be achieved by summarizing the shapes with filtrations taken over multiple directions. In practice, we suggest specifying multiple cones $c > 1$ and utilizing multiple directions d per cone (see monotonically increasing power in Supplementary Fig. 1).

While the other two parameters (θ and l) do not have monotonic properties, their effects on SINATRA's performance still have natural interpretations. For example, when changing the angle between directions within cones from $\theta \in [0.05, 0.5]$ radians, we observe that power steadily increases until $\theta = 0.25$ radians and then slowly decreases afterwards (Supplementary Fig. 2). This supports previous theoretical results that cones should be defined by directions in close proximity to each other [3]; but not so close that they explain the same local information with little variation.

Perhaps most importantly, we must understand how the number of sublevel sets l (i.e. the number of steps in the filtration) used to compute Euler characteristic curves affects the performance of the algorithm. As we show in the next section, this function depends on the types of shapes being analyzed.

Intuitively, for very intricate shapes, coarse filtrations with too few sublevel sets cause the algorithm to miss or “step over” very local undulations in a shape. For the spheres simulated in this section, class-defining regions are global-like features, and so finer filtration steps fail to capture broader differences between shapes (Supplementary Fig. 3); however, this failure is less important when only a few features decide how shapes are defined (e.g. scenario I). In practice, we recommend choosing the angle between directions within cones θ and the number of sublevel sets l via cross validation or some grid-based search.

As a final demonstration, we show what happens when we meet the null assumptions of the SINATRA pipeline (Supplementary Fig. 4). Under the null hypothesis, our feature selection measure assumes that all 3D regions of a shape equally contribute to explaining the variance between classes — that is, no one vertex (or corresponding topological characteristics) is more important than the others. We generate synthetic shapes under the two cases when SINATRA fails to produce significant results: (a) two classes of shapes that are effectively the same (up to some small Gaussian noise), and (b) two classes of shapes that are completely dissimilar. In the first simulation case, there are no “significantly associated” regions and thus no group of vertices stand out as important (Supplementary Fig. 4(a)). In the latter simulation case, shapes between the two classes look nothing alike; therefore, all vertices contribute to class definition, but no one feature is key to explaining the observed variation (Supplementary Fig. 4(b)).

Simulation Study: Caricatured Shapes

Our second simulation study modifies computed tomography (CT) scans of real Lemuridae teeth (one of the five families of Strepsirrhini primates commonly known as lemurs) [10] using a well-known caricaturization procedure [22]. We fix the triangular mesh of an individual tooth and specify class-specific regions centered around known biological landmarks (Fig. 3) [10]. For each triangular face contained within a class-specific region, we apply a corresponding affine transformation, positively scaled, that smoothly varies on the triangular mesh and attains its maximum at the biological landmark used to define the region (Supplementary Note Section 3). We caricature 50 different teeth with two steps (Fig. 3(a)):

1. Assign v of a given tooth’s landmarks to be specific to one class and v' to be specific to the other class.
2. Perform the caricaturization: multiply each face in the v and v' class-specific regions by a positive scalar (i.e. exaggerated or enhanced). Repeat twenty-five times (with some small noise per replicate) to create two equally-sized classes of 25 shapes.

We explore two scenarios by varying the number of class-specific landmarks v and v' that determine the caricaturization in each class. First we set both $v, v' = 3$; next, we fix $v, v' = 5$. Like the simulations with perturbed spheres, the difficulty of the scenarios increases with the number of caricatured regions. We evaluate SINATRA’s ability to identify the vertices involved in the caricaturization using ROC curves (Supplementary Note Section 2), and we assess this estimate of power as a function of the algorithm’s free parameter inputs. While varying each parameter, we hold the others as constants $\{c = 15, d = 5, \theta = 0.15, l = 50\}$. Figures in the main text are based on varying the number of cones c (Figs. 3(b) and 3(c)); results for the other sensitivity analyses can be found in Supplementary Note (Supplementary Figs. 5-7).

Overall, using fewer caricatured regions results in better (or at least comparable) performance. Like the simulations with perturbed spheres, SINATRA’s power increases monotonically with an increasing number of cones and directions used to compute the topological summary statistics (Figs. 3(b), 3(c), and Supplementary 5). For example, at a 10% FPR with $c = 5$ cones, we achieve 30% TPR in scenario I experiments and 35% TPR in scenario II. Increasing the number of cones to $c = 35$ improves power to 52% and 40% TPR for scenarios I and II, respectively. Trends from the previous section continue when choosing the angle between directions within cones (Supplementary Fig. 6) and the number of sublevel sets (Supplementary Fig. 7). Results for the perturbed spheres suggest that there is an optimal cap

radius for generating directions in a cone. Since we are analyzing shapes with more intricate features, finer filtrations lead to more power.

Recovering Known Morphological Variation Across Genera of Primates

As an application of our pipeline, with “ground truth” or known morphological variation, we consider a dataset of CT scans of $n = 59$ mandibular molars from two suborders of primates: Haplorhini (which include tarsiers and anthropoids) and Strepsirrhini (which include lemurs, galagos, and lorises). From the haplorhine suborder, 33 molars came from the genus *Tarsius* [10, 23, 24] and 9 molars from the genus *Saimiri* [25]. From the strepsirrhine suborder, 11 molars came from the genus *Microcebus* and 5 molars from the genus *Mirza* [10, 23, 24]; both are lemurs. The meshes of all teeth were aligned, centered at the origin, and normalized within a unit sphere (Methods and Supplementary Fig. 8).

We chose this specific collection of molars because morphologists and evolutionary anthropologists understand variations of the paraconid, the cusp of a primitive lower molar. The paraconids are retained only by *Tarsius* and do not appear in the other genera (Fig. 4(a)) [25, 26]. Using phylogenetic analyses of mitochondrial genomes across primates, Pozzie et al. estimate divergence dates of the subtree composed of *Microcebus* and *Mirza* from *Tarsius* at 5 million years before the branching of *Tarsius* from *Saimiri* [27]. We want to see if SINATRA recovers the information that the paraconids are specific to the *Tarsius* genus. We also investigate if variation across the molar is associated to the divergence time of the genera.

Since *Tarsius* is the only genus with the paraconid in this sample, we use SINATRA to perform three pairwise classification comparisons (*Tarsius* against *Saimiri*, *Mirza*, and *Microcebus*, respectively), and assess SINATRA’s ability to prioritize/detect the location of the paraconid as the region of interest (ROI). Based on our simulation studies, we run SINATRA with $c = 35$ cones, $d = 5$ directions per cone, a cap radius of $\theta = 0.25$ to generate each direction, and $l = 75$ sublevel sets to compute topological summary statistics. In each comparison, we evaluate the evidence for each vertex based on the first time that it appears in the reconstruction: this is the evidence potential for a vertex (Methods). A heatmap for each tooth (Fig. 4(b)) provides visualization of the physical regions that are most distinctive between the genera.

To assess SINATRA’s ability to find *Tarsius*-specific paraconids, we use a null-based scoring method. We place a paraconid landmark on each *Tarsius* tooth, and consider the $K = \{10, 50, 100, 150, 200\}$ nearest vertices surrounding the landmark’s centermost vertex. This collection of $K + 1$ vertices defines our ROI. Within each ROI, we weight the SINATRA-computed evidence potentials by the surface area (or area of the Voronoi cell) encompassed by their corresponding vertices, and then sum the scaled potentials together across the ROI vertices. This aggregated value, which we denote as τ^* , represents a score of association for the ROI. To construct a “null” distribution and assess the strength of any score τ^* , we randomly select $N = 500$ other “seed” vertices across the mesh of each *Tarsius* tooth and uniformly generate N -“null” regions that are K -vertices wide. We then compute similar (null) scores $\tau_1, \dots, \tau_N$ for each randomly generated region. A “p-value”-like quantity (for the i -th molar) is then generated by:

$$P_i = \frac{1}{N + 1} \sum_{t=1}^N \mathbb{I}(\tau_i^* \leq \tau_t), \quad (1)$$

where $\mathbb{I}(\cdot)$ is an indicator function, and a smaller P_i means more confidence in SINATRA’s ability to find the desired paraconid landmark. To ensure the robustness of this analysis, we generate the N -random null regions in two ways: (i) using a K -nearest neighbors (KNN) algorithm on each of the N -random seed vertices [28], or (ii) manually constructing K -vertex wide null regions with surface areas equal to that of the paraconid ROI (Supplementary Note Section 4). In both settings, we take the median of the P_i values in Equation (1) across all teeth, and report them for each genus and choice of K combination (see the first half of Table 1).

Using p-values as a direct metric of evidence can cause problems. For example, moving from $P = 0.03$ to $P = 0.01$ does not increase evidence for the alternative hypothesis (or against the null hypothesis) by a factor of 3. To this end, we use a calibration formula that transforms a p-value to a bound/approximation of a Bayes factor (BF) [17], the ratio of the marginal likelihood under the alternative hypothesis H_1 versus the null hypothesis H_0 :

$$BF(P_i)_{10} = [-e P_i \log(P_i)]^{-1}, \quad (2)$$

for $P_i < 1/e$ and $BF(P_i)_{10}$ is an estimate of $\Pr(H_1 | \mathcal{M})/\Pr(H_0 | \mathcal{M})$, where $\mathcal{M}$ are the molars as meshes and H_0 and H_1 are the null and alternative hypotheses, respectively. Table 1 reports the calibrated Bayes factor estimates.

Overall, the paraconid ROI is more strongly enriched in the comparisons between the *Tarsius* and either of the strepsirrhine primates, rather than for the *Tarsius-Saimiri* comparison. We suspect this difference is partly explained by the divergence times between these genera: *Tarsius* is more recently diverged from *Saimiri* than from the strepsirrhines. This conjecture is consistent with the intuition from our simulation studies, where classes of shapes with sufficiently different morphology result in more accurate identification of unique ROI. On the other hand, the *Tarsius-Saimiri* comparison is analogous to the simulations under to the null model: with too-similar molars, no region appears key to explaining the variance between the two classes of primates.

Discussion

In this paper, we introduce SINATRA: the first statistical pipeline for sub-image analysis that does not require landmarks or correspondence points between images. We use simulations to demonstrate properties of SINATRA and we illustrate the practical utility of SINATRA on real data. SINATRA's current formulation and software is limited to binary classification, but we believe that extensions to multi-class problems and regression with continuous responses are trivial.

To analyze continuous traits and phenotypes in many evolutionary applications, one must first disentangle adaptation and heredity [29–32]. The standard approach for this disentanglement is to explicitly account for the hierarchy of descent by adding genetic covariance or kinship across species to the likelihood either via phylogenetic regression [33] or linear mixed models (e.g. the animal model) [34]. Modeling covariance structures also arises in statistical and quantitative genetics applications where individuals are related [35–37]. The SINATRA framework uses a Bayesian hierarchical model that is straightforward to adapt to analyze complex covariance structures in future work.

Acknowledgements

The authors would like to thank Ani Eloyan, Anthea Monod, Jenny Tung, Katharine Turner, and Christine Wall for helpful conversations and suggestions. This research was partly supported by grants P20GM109035 (COBRE Center for Computational Biology of Human Disease; PI Rand) and P20GM103645 (COBRE Center for Central Nervous; PI Sanes) from the NIH NIGMS, 2U10CA180794-06 from the NIH NCI and the Dana Farber Cancer Institute (PIs Gray and Gatsonis), as well as by an Alfred P. Sloan Research Fellowship (No. FG-2019-11622) awarded to LC. A majority of this research was conducted using computational resources and services at the Center for Computation and Visualization (CCV), Brown University. SM would like to acknowledge partial funding from HFSP RGP005, NSF DMS 17-13012, NSF BCS 1552848, NSF DBI 1661386, NSF IIS 15-46331, NSF DMS 16-13261, as well as high-performance computing partially supported by grant 2016-IDG-1013 from the North Carolina Biotechnology Center. Lastly, TG was supported by NSF Grant No. DMS-1439786 while in residence at the Institute for Computational and Experimental Research in Mathematics (ICERM) in Providence, RI, during the Computer

310 Vision Semester Program. Any opinions, findings, and conclusions or recommendations expressed in this
311 material are those of the author(s) and do not necessarily reflect the views of any of the funders.

312 **Author Contributions Statement**

313 LC conceived the study. SM and LC developed the methods. BW, TS, and HK developed the algorithms
314 and implemented the software. DB designed sampling strategy for the molar analysis. All authors
315 performed the analyses, interpreted the results, and wrote and revised the manuscript.

316 **Competing Financial Interests**

317 The authors have declared that no competing interests exist.

Methods

Topological Summary Statistics for 3D Shapes

In the first step of the SINATRA pipeline, we use a tool from integral geometry and differential topology called the Euler characteristic (EC) transform [1–4]. For a mesh $\mathcal{M}$, the Euler characteristic is an accessible topological invariants derived from:

$$\chi = \#V(\mathcal{M}) - \#E(\mathcal{M}) + \#F(\mathcal{M}), \quad (3)$$

where $\{\#V(\mathcal{M}), \#E(\mathcal{M}), \#F(\mathcal{M})\}$ denote the number of vertices (corners), edges, and faces of the mesh, respectively. An EC curve $\chi_\nu(\mathcal{M})$ tracks the change in the Euler characteristic with respect to a given filtration of length l in direction ν (Figs. 1(a) and (b)). First we specify a height function $h_\nu(\mathbf{x}) = \mathbf{x}^\top \nu$ for vertex $\mathbf{x} \in \mathcal{M}$ in direction ν . We then use this height function to define sublevel sets (or subparts) of the mesh $\mathcal{M}_\nu^a$ in direction ν , where $h_\nu(\mathbf{x}) \leq a$. The EC curve is $\chi(\mathcal{M}_\nu^a)$ over a range of l filtration steps over a (Fig. 1(b)).

The EC transform is the collection of EC curves across a set of directions $\nu = 1, \dots, m$, and maps a 3D shape into a concatenated $p = (l \times m)$ -dimensional feature vector. For a study with n -shapes, an $n \times p$ design matrix $\mathbf{X}$ is statistically analyzed, where the columns denote the Euler characteristic computed at a given filtration step and direction. Each sublevel set value, direction, and set of shape vertices used to compute an EC curve are stored for the association mapping and projection phases of the pipeline. Previously, Curry et al. proved sufficiency, stating that m number of directions and the l range of sublevel set values required for the EC transform preserve all information for a family of shapes [3]. In this paper, we used simulations to outline empirical procedures and developed intuition behind these quantities.

Statistical Model for Shape Classification

In the second step of the SINATRA pipeline, we use (weight-space) Gaussian process probit regression to classify shapes based on their topological summaries generated by the EC transformation. Namely, we specify the following (Bayesian) hierarchical model [38–42]

$$\mathbf{y} \sim \mathcal{B}(\boldsymbol{\pi}), \quad g(\boldsymbol{\pi}) = \Phi^{-1}(\boldsymbol{\pi}) = \mathbf{f}, \quad \mathbf{f} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}), \quad (4)$$

where $\mathbf{y}$ is an n -dimensional vector of Bernoulli distributed class labels, $\boldsymbol{\pi}$ is an n -dimensional vector representing the underlying probability that a shape is classified as a “case” (i.e. $y = 1$), $g(\cdot)$ is a probit link function with $\Phi(\cdot)$ the cumulative distribution function (CDF) of the standard normal distribution, and $\mathbf{f}$ is an n -dimensional vector estimated from the data.

The key objective of SINATRA is to use the topological features in $\mathbf{X}$ to find the physical 3D properties that best explain the variation across shape classes. To do so, we use kernel regression, where the utility of generalized nonparametric statistical models is well-established due their ability to account for various complex data structures [43–48]. Generally, kernel methods posit that $\mathbf{f}$ lives within a reproducing kernel Hilbert space (RKHS) defined by some (nonlinear) covariance function, which implicitly account for higher-order interactions between features, leading to more complete classifications of data [49–51]. To this end, we assume $\mathbf{f}$ is normally distributed with mean vector $\mathbf{0}$, and covariance matrix $\mathbf{K}$ defined by the radial basis function $\mathbf{K}_{ij} = \exp\{-\theta\|\mathbf{x}_i - \mathbf{x}_j\|^2\}$ with bandwidth θ set using the median heuristic [52]. The full model specified in Equation (4) is commonly referred to as “Gaussian process classification” or GPC.

Interpretable Feature (Variable) Selection

To estimate the model in Equation (4), we use an elliptical slice sampling Markov chain Monte Carlo (MCMC) algorithm (Supplementary Note Section 1.1). This allows samples from the approximate posterior distribution of $\mathbf{f}$ (given the data), and also allows for the computation of an effect size analog for each topological summary statistic [53–55]

$$\boldsymbol{\beta} = (\mathbf{X}^\top \mathbf{X})^\dagger \mathbf{X}^\top \mathbf{f}, \quad (5)$$

where $(\mathbf{X}^\top \mathbf{X})^\dagger$ is the generalized inverse of $(\mathbf{X}^\top \mathbf{X})$.

These effect sizes represent the nonparametric equivalent to coefficients in linear regression using generalized ordinary least squares. SINATRA uses these weights and assigns a measure of relative centrality to each summary statistic (first panel Fig. 1(c)) [55]. This criterion evaluates how much information in classifying each shape is lost when a particular topological feature is removed from the model. This loss is determined by computing the Kullback-Leibler divergence (KLD) between (i) the conditional posterior distribution $p(\boldsymbol{\beta}_{-j} | \beta_j = 0)$ with the effect of the j -th topological feature set to zero, and (ii) the marginal posterior distribution $p(\boldsymbol{\beta}_{-j})$ with the effects of the j -th feature integrated out:

$$\text{KLD}(\beta_j) = \int_{\boldsymbol{\beta}_{-j}} \log \left(\frac{p(\boldsymbol{\beta}_{-j})}{p(\boldsymbol{\beta}_{-j} | \beta_j = 0)} \right) p(\boldsymbol{\beta}_{-j}) d\boldsymbol{\beta}_{-j} \quad j = 1, \dots, p, \quad (6)$$

which has a closed form solution when the posterior distribution of the effect sizes is assumed to be (approximately) Gaussian (Supplementary Note 1.2). Finally, we normalize to obtain an association metric for each topological feature, $\gamma_j = \text{KLD}(\beta_j) / \sum \text{KLD}(\beta_i)$.

There are two main takeaways from this formulation. First, the KLD is non-negative, and equals zero if and only if the posterior distribution of $\boldsymbol{\beta}_{-j}$ is independent of the effect β_j . Intuitively, this says that removing an unimportant shape feature has no impact on explaining the variance between shape classes. Second, γ is bounded on the unit interval $[0, 1]$ with the natural interpretation of providing relative evidence of association for shape features; higher values suggest greater importance. For this metric, the null hypothesis assumes that every feature equally contributes to the total variance between shape classes, while the alternative proposes that some features are more central than others [55]. As we show in the Supplementary Note, when the null assumption is met, SINATRA displays association results that appear uniformly distributed and effectively indistinguishable.

Shape Reconstruction

After obtaining association measures for each topological feature, we map this information back onto the physical shape (second panel Fig. 1(c) and 1(d)). We refer to this process as *reconstruction*, as this procedure recovers regions that explain the most variation between shape classes (Supplementary Note Section 1.3). Intuitively, we want to identify vertices on the shape that correspond to the topological features with the greatest association measures.

Begin by considering d directions within a cone of cap radius or angle θ , which we denote as $\mathcal{C}(\theta) = \{\nu_1, \dots, \nu_d | \theta\}$. Next, let $\mathcal{Z}$ be the set of vertices whose projections onto the directions in $\mathcal{C}(\theta)$ are contained within the collection of “significant” topological features — for every $z \in \mathcal{Z}$, the product $z \cdot \nu$ is contained within a sublevel set (taken in the direction $\nu \in \mathcal{C}(\theta)$) that shows high evidence of association in the feature selection step.

A reconstructed region is then defined as the union of all mapped vertices from each cone, or $\mathcal{R} := \bigcup_i \mathcal{Z}_i$. We use cones because vectors of Euler characteristics taken along directions close together express comparable information. That similarity lets us leverage findings between them to increase our power of detecting truly associated shape vertices and regions — as opposed to antipodal directions where the

lack of shared information may do harm when determining reconstructed manifolds (Supplementary Note Section 1.4) [3, 56, 57].

Visualization of Enriched Shape Regions

Once shapes have been reconstructed, we can visualize the relative importance or “evidence potential” for each vertex on the mesh with a simple procedure. First, we sort the topological features from largest to smallest according to their association measures $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_p$. Next, we iterate through the sorted measures $T_k = \gamma_k$ (starting with $k = 1$), and reconstruct the vertices corresponding to the topological features in the set $\{j : \gamma_j \geq T_k\}$.

The evidence potential for each vertex is defined as the largest threshold T_k at which it is reconstructed for the first time, because vertices with earlier “birth times” in the reconstruction are more important relative to vertices that appear later. We illustrate these values via heatmaps over the reconstructed meshes (Fig. 1(d)). For consistency across different applications and case studies, we set the coloring of these heatmaps on a scale from $[0 - 100]$. A maximum value of 100 represents the threshold value at which the first vertex is born, while 0 denotes the threshold when the last vertex on the shape is reconstructed. Under the null hypothesis, where there are no meaningful regions differentiating between two classes of shapes, (mostly) all vertices appear to be born relatively early and at the same time (Supplementary Fig. 4). This is not the case under the alternative.

Code Availability

Code for implementing the SINATRA pipeline is freely available at <https://github.com/lcrawlab/SINATRA>, and is written in R (version 3.5.3). As part of this procedure: (i) inference for the Gaussian process classification (GPC) model using elliptical slice sampling was carried out using the R package **FastGP** (version 1.2) [58] and (ii) the computation of effect sizes and association measures for the Euler characteristic curves was done with the “RelAtive cEntRality (RATE)” source code in R (version 1.0.0; <https://github.com/lorinanthony/RATE>) [55].

Visualizing the reconstructed regions outputted by SINATRA was done using the package **rgl** (version 0.100.19) [59], and general utility functions for triangular meshes from the package **Rvcg** (version 0.18) [60]. Furthermore, preprocessing steps for the meshes examined in the study were performed using **Morpho** (version 2.60) [60, 61] and **auto3dgm** (Version 1.00) [62].

Data Availability

The current study makes use of two real shape datasets. The first consists of Lemuridae teeth, a specific genera of Cercopithecidae (Old World monkeys; <http://www.wisdom.weizmann.ac.il/~yflipman/CPsurfcomp/>) [10]. The second is comprised of mandibular molars from two different suborders of the primate: Haplorhini (“dry-nosed” primates; <https://gaotingran.com/codes/codes.html>) and Strepsirrhini (“moist-nosed” primates; http://morphosource.org/Detail/ProjectDetail/Show/project_id/89). From the first suborder, we have 33 molars from the *Tarsius* [10, 23, 24] and 9 molars from the *Saimiri* [25] genera. In the second suborder, we have 11 molars from the *Microcebus* and 5 molars from the *Mirza* genera [10, 23, 24]. Prior to any analysis, the meshes of all teeth were aligned using **auto3dgm** [62]. This algorithm establishes correspondences between uniformly placed landmarks on each tooth such that each mesh has the same orientation (e.g. Fig 8). After alignment, the molars were translated to be centered at the origin and normalized to be enclosed within a unit ball. These quality controlled meshes were then used to demonstrate the utility of the SINATRA pipeline.

Figures and Tables

Figure 1. Schematic overview of SINATRA: a novel statistical framework for feature selection and association mapping with 3D shapes. (a) The SINATRA algorithm requires the following inputs: (i) aligned shapes represented as meshes; (ii) $\mathbf{y}$, a binary vector denoting shape classes; (iii) r , the radius of the bounding sphere for the shapes; (iv) c , the number of cones of directions; (v) d , the number of directions within each cone; (vi) θ , the cap radius used to generate directions in a cone; and (vii) l , the number of sublevel sets (i.e. filtration steps) to compute the Euler characteristic (EC) along a given direction. (b) We select initial positions uniformly on a unit sphere. Then for each position, we generate a cone of d directions within angle θ using Rodrigues' rotation formula [63], resulting in a total of $m = c \times d$ directions. For each direction, we compute EC curves with l sublevel sets. We concatenate the EC curves along all the directions for each shape to form vectors of topological features of length $p = l \times m$. Thus, for a study with n -shapes, an $n \times p$ design matrix is statistically analyzed using a Gaussian process classification model. (c) Evidence of association for each topological feature vector are determined using relative centrality measures. We reconstruct corresponding shape regions by identifying the vertices (or locations) on the shape that correspond to “statistically associated” topological features. (d) This enables us to visualize the enrichment of physical features that best explain the variance between the two classes. The heatmaps display vertex evidence potential on a scale from $[0 - 100]$. A maximum of 100 represents the threshold at which the first shape vertex is reconstructed, while 0 denotes the threshold when the last vertex is reconstructed.

Figure 2. Power analysis for detecting associated vertices across different classes of perturbed spheres. We generate 100 shapes by partitioning unit spheres into regions 10 vertices-wide, centered at 50 equidistributed points. Two classes (50 shapes per class) are defined by shared (blue protrusions) and class-specific (red indentations) characteristics. The shared or “non-associated” features are chosen by randomly selecting u regions and pushing the sphere outward at each of these positions. To generate class-specific or “associated” features, v distinct regions are chosen for a given class and perturbed inward. We vary these parameters and analyze three increasingly more difficult simulation scenarios: **(a)** $u = 2$ shared and $v = 1$ associated; **(b)** $u = 6$ shared and $v = 3$ associated; and **(c)** $u = 10$ shared and $v = 5$ associated. In panels **(d)**–**(f)**, ROC curves depict the ability of SINATRA to identify vertices located within associated regions, as a function of increasing the number of cones of directions used in the algorithm. These results give empirical evidence that seeing more of a shape (i.e. using more unique directions) generally leads to an improved ability to map back onto associated regions. Other SINATRA parameters were fixed: $d = 5$ directions per cone, $\theta = 0.15$ cap radius used to generate directions in a cone, and $l = 30$ sublevel sets per filtration. Results are based on fifty replicates in each scenario.

Figure 3. Power analysis for detecting associated vertices across different classes of caricatured shapes. (a) We modify real Lemuridae molars using the following caricaturization procedure: (i) Fix the triangular mesh of an individual tooth. (ii) Assign v of the known landmarks for the tooth [10] to be specific to one class and v' to be specific to the other. The caricaturization is performed by positively scaling each face within these regions so that class-specific features are exaggerated. We repeat twenty-five times (with some small added noise) to create two classes of 25 shapes. (iii) The synthetic shapes are analyzed by SINATRA to identify the associated regions. We consider two scenarios by varying the number of class-specific landmarks that determine the caricaturization in each class. In scenario I, we set $v, v' = 3$; and in scenario II, $v, v' = 5$. In panels (b) and (c), ROC curves depict the ability of SINATRA to identify vertices located within associated regions, as a function of the number of cones of directions used in the algorithm. Other SINATRA parameters were fixed: $d = 5$ directions per cone, $\theta = 0.15$ cap radius used to generate directions in a cone, and $l = 50$ sublevel sets per filtration. Results are based on fifty replicates in each scenario.

Figure 4. Real data analysis aimed at detecting unique paraconids in molars belonging to primates in *Tarsius* genus. We carry out three pairwise comparisons to analyze the physical differences between *Tarsius* molars and teeth from (i) *Saimiri*, (ii) *Mirza*, and (iii) *Microcebus* genus. In panel (a), we depict the phylogenetic relationship between these groups. Morphologically, we know that tarsier teeth have an additional high-cusp (highlighted in red), which allows this genus of primate to eat a wider range of foods [64]. We want to assess SINATRA's ability to find this region of interest (ROI). In panel (b), we show an example of the reconstruction from each comparison. Results are consistent with the phylogeny of the primates, as well as with our previous simulation studies. Genetically, *Tarsius* differ more from the *Mirza* and *Microcebus* genera, rather than from *Saimiri*. As a result, SINATRA finds the unique paraconid in the former two comparisons because of the appropriate genetic distance, rather than in the latter case where molar structures are more similar. The heatmaps display vertex evidence potential on a scale from [0 – 100]. A maximum of 100 represents the threshold at which the first shape vertex is reconstructed, while 0 denotes the threshold when the last vertex is reconstructed.

	Test	Region Size	<i>Tarsius</i> vs. <i>Saimiri</i>	<i>Tarsius</i> vs. <i>Mirza</i>	<i>Tarsius</i> vs. <i>Microcebus</i>
<i>P-Values (P)</i>	KNN	10	4.75×10^{-1}	3.39×10^{-1}	2.14×10^{-1}
		50	2.89×10^{-1}	2.10×10^{-1}	1.56×10^{-1}
		100	2.14×10^{-1}	2.20×10^{-2}	6.19×10^{-2}
		150	1.99×10^{-1}	1.80×10^{-2}	6.59×10^{-2}
		200	2.22×10^{-1}	2.99×10^{-2}	9.18×10^{-2}
	Equal-Area	10	3.21×10^{-1}	2.10×10^{-1}	1.84×10^{-1}
		50	2.81×10^{-1}	1.72×10^{-1}	1.26×10^{-1}
		100	2.40×10^{-1}	4.39×10^{-2}	8.78×10^{-2}
		150	2.59×10^{-1}	3.79×10^{-2}	8.18×10^{-2}
		200	2.55×10^{-1}	4.39×10^{-2}	9.98×10^{-2}
<i>Bayes Factors (BF)</i>	KNN	10	—	1.003	1.115
		50	1.025	1.122	1.269
		100	1.115	4.381	2.136
		150	1.145	5.087	2.053
		200	1.101	3.505	1.678
	Equal-Area	10	1.009	1.122	1.181
		50	1.031	1.215	1.409
		100	1.074	2.681	1.722
		150	1.051	3.016	1.796
		200	1.055	2.681	1.599

Table 1. Null region experiment to evaluate SINATRA’s ability to find paraconids in *Tarsius* molars. We assess how likely it is that SINATRA finds the region of interest (ROI) by chance. We generate 500 “null” regions on each *Tarsius* tooth using (i) a KNN algorithm and (ii) an equal-area approach (Supplementary Note Section 4). Next, for each region, we sum the evidence potential or “birth times” of all its vertices. We compare how many times the aggregate scores for the ROI is less than those for the null regions. The median of these “p-values”, and their corresponding calibrated Bayes factors (BF) when median $P < 1/e$, across all teeth are provided above for the three primate comparisons. Results with values p-values less than 0.1 and BF’s greater than 1.598 are in bold.

References

1. Crawford L, Monod A, Chen AX, Mukherjee S, Rabadán A. Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis. J Am Stat Assoc. In Press; Available from: <https://arxiv.org/abs/1611.06818>.
2. Turner K, Mukherjee S, Boyer DM. Persistent homology transform for modeling shapes and surfaces. Inf Inference. 2014;3(4):310–344. Available from: <https://academic.oup.com/imaia/article-abstract/3/4/310/724811?redirectedFrom=fulltext>.
3. Curry J, Mukherjee S, Turner K. How many directions determine a shape and other sufficiency results for two topological transforms. arXiv. 2018;p. 1805.09782. Available from: <https://arxiv.org/abs/1805.09782>.
4. Ghrist R, Levanger R, Mai H. Persistent homology and Euler integral transforms. J Appl and Comput Topology. 2018;2(1-2):55–60. Available from: <https://link.springer.com/article/10.1007/s41468-018-0017-1>.
5. Kendall DG. A Survey of the Statistical Theory of Shape. Statist Sci. 1989;4(2):87–99. Available from: <https://doi.org/10.1214/ss/1177012582>.
6. Dupuis P, Grenander U. Variational problems on flows of diffeomorphisms for image matching. Q Appl Math. 1998;LVI(3):587–600. Available from: <http://dl.acm.org/citation.cfm?id=298828.298844>.
7. Worsley KJ. Estimating the number of peaks in a random field using the Hadwiger characteristic of excursion Sets, with applications to medical images. Ann Stat. 1995;23(2):640–669. Available from: <https://doi.org/10.1214/aos/1176324540>.
8. Goswami A. Phenome10K: a free online repository for 3-D scans of biological and palaeontological specimens; 2015. Available from: www.phenome10k.org.
9. Boyer DM, Gunnell GF, Kaufman S, McGeary TM. Morphosource: archiving and sharing 3-D digital specimen data. The Paleontological Society Papers. 2016;22:157–181.
10. Boyer DM, Lipman Y, St Clair E, Puente J, Patel BA, Funkhouser T, et al. Algorithms to automatically quantify the geometric similarity of anatomical surfaces. Proc Natl Acad Sci U S A. 2011;108(45):18221. Available from: <http://www.pnas.org/content/108/45/18221.abstract>.
11. Ovsjanikov M, Ben-Chen M, Solomon J, Butscher A, Guibas L. Functional maps: a flexible representation of maps between shapes. ACM Trans Graph. 2012;31(4):30:1–30:11. Available from: <http://doi.acm.org/10.1145/2185520.2185526>.
12. Boyer DM, Puente J, Gladman JT, Glynn C, Mukherjee S, Yapuncich GS, et al. A new fully automated approach for aligning and comparing shapes. Anat Rec (Hoboken). 2015;298(1):249–276. Available from: <https://onlinelibrary.wiley.com/doi/abs/10.1002/ar.23084>.
13. Gao T, Kovalsky SZ, Daubechies I. Gaussian process landmarking on manifolds. SIAM J Math Data Sci. 2019;1(1):208–236. Available from: <https://epubs.siam.org/doi/abs/10.1137/18M1184035>.
14. Gao T, Kovalsky S, Boyer D, Daubechies I. Gaussian process landmarking for three-dimensional geometric morphometrics. SIAM J Math Data Sci. 2019;1(1):237–267. Available from: <https://epubs.siam.org/doi/abs/10.1137/18M1203481>.

- 481 15. Miller E. Fruit flies and moduli: interactions between biology and mathematics. Notices of the
482 AMS. 2015;62(10):1178–1184.
- 483 16. Bendich P, Marron JS, Miller E, Pieloch A, Skwerer S. Persistent homology analysis of brain
484 artery trees. Ann Appl Stat. 2016 03;10(1):198–218. Available from: [https://doi.org/10.1214/](https://doi.org/10.1214/15-AOAS886)
485 15-AOAS886.
- 486 17. Sellke T, Bayarri MJ, Berger JO. Calibration of p-values for testing precise null hypotheses. Am
487 Stat. 2001;55(1):62–71.
- 488 18. Li F, Zhang T, Wang Q, Gonzalez MZ, Maresh EL, Coan JA. Spatial Bayesian variable selection
489 and grouping for high-dimensional scalar-on-image regression. Ann Appl Stat. 2015;9(2):687–713.
490 Available from: <https://projecteuclid.org:443/euclid.aoas/1437397107>.
- 491 19. Crawford L, Zeng P, Mukherjee S, Zhou X. Detecting epistasis with the marginal epistasis test in
492 genetic mapping studies of quantitative traits. PLoS Genet. 2017;13(7):e1006869–. Available from:
493 <https://doi.org/10.1371/journal.pgen.1006869>.
- 494 20. Zhu X, Stephens M. Large-scale genome-wide enrichment analyses identify new trait-associated
495 genes and pathways across 31 human phenotypes. Nat Comm. 2018;9(1):4361.
- 496 21. Lipman Y. Collection of Lemuridae teeth; 2011. Available from: [http://www.wisdom.weizmann.](http://www.wisdom.weizmann.ac.il/~ylipman/CPsurfcomp/)
497 [ac.il/~ylipman/CPsurfcomp/](http://www.wisdom.weizmann.ac.il/~ylipman/CPsurfcomp/).
- 498 22. Sela M, Aflalo Y, Kimmel R. Computational caricaturization of surfaces. Comput Vis Image
499 Underst. 2015;141:1–17. Available from: <https://doi.org/10.1016/j.cviu.2015.05.013>.
- 500 23. Gao T. Hypoelliptic diffusion maps and their applications in automated geometric morphometrics.
501 Duke University; 2015.
- 502 24. Gao T. The diffusion geometry of fibre bundles: horizontal diffusion maps. arXiv. 2016;p.
503 1602.02330.
- 504 25. St Clair EM, Boyer DM. Lower molar shape and size in prosimian and platyrrhine primates. Am
505 J Phys Anthropol. 2016;161(2):237–258.
- 506 26. Guatelli-Steinberg D. Primate dentition: an introduction to the teeth of non-human primates. Am
507 J Phys Anthropol. 2003;121(2):189–189. Available from: <https://doi.org/10.1002/ajpa.10194>.
- 508 27. Pozzi L, Hodgson JA, Burella AS, Raaumb RL, Disotell TR. Primate phylogenetic relationships and
509 divergence dates inferred from complete mitochondrial genomes. Mol Phylogenet Evol. 2014;75:165–
510 183.
- 511 28. Cover T, Hart P. Nearest neighbor pattern classification. IEEE Trans Inf Theor. 2006;13(1):21–27.
512 Available from: <https://doi.org/10.1109/TIT.1967.1053964>.
- 513 29. Gienapp P, Teplitsky C, Alho JS, Mills JA, Merila J. Climate change and evolution: disentangling
514 environmental and genetic responses. Mol Ecol. 2008;17(1):167–178.
- 515 30. Anderson JT, Willis JH, Mitchell-Olds T. Evolutionary genetics of plant adaptation. Trends Genet.
516 2011;27(7):258–266. Available from: <https://www.ncbi.nlm.nih.gov/pubmed/21550682>.
- 517 31. Chen J, Källman T, Ma X, Gyllenstrand N, Zaina G, Morgante M, et al. Disentangling the roles
518 of history and local selection in shaping clinal variation of allele frequencies and gene expression
519 in Norway spruce (*Picea abies*). Genetics. 2012;191(3):865–881. Available from:
520 <http://www.genetics.org/content/191/3/865.abstract>.

- 521 32. Lai YT, Yeung CKL, Omland KE, Pang EL, Hao Y, Liao BY, et al. Standing genetic variation as
522 the predominant source for adaptation of a songbird. *Proc Natl Acad Sci U S A*. 2019;116(6):2152–
523 2157. Available from: <http://www.pnas.org/content/116/6/2152.abstract>.
- 524 33. Graven A. The phylogenetic regression. *Philos Trans R Soc Lond B Biol Sci*. 1989;326(1233):87–99.
- 525 34. Henderson CRCR, of Guelph U. Applications of linear models in animal breeding. Guelph, Ont. :
526 University of Guelph; 1984. Includes index.
- 527 35. Kang HM, Zaitlen NA, Wade CM, Kirby A, Heckerman D, Daly MJ, et al. Efficient control of
528 population structure in model organism association mapping. *Genetics*. 2008;178(3):1709–1723.
529 Available from: <http://www.genetics.org/content/178/3/1709.abstract>.
- 530 36. Kang HM, Sul JH, Service SK, Zaitlen NA, Kong Sy, Freimer NB, et al. Variance component model
531 to account for sample structure in genome-wide association studies. *Nat Genet*. 2010;42(4):348–354.
532 Available from: <http://dx.doi.org/10.1038/ng.548>.
- 533 37. Zhou X, Stephens M. Genome-wide efficient mixed-model analysis for association studies. *Nat*
534 *Genet*. 2012;44(7):821–825.
- 535 38. Neal RM. Monte Carlo implementation of Gaussian process models for Bayesian regression and-
536 Monte Carlo implementation of Gaussian process models for Bayesian regression and classification.
537 Dept. of Statistics, University of Toronto; 1997. 9702.
- 538 39. Neal RM. Regression and classification using Gaussian process priors. *Bayesian Anal*. 1998;6:475.
- 539 40. Williams CKI, Barber D. Bayesian classification with Gaussian processes. *IEEE Trans Pattern Anal*
540 *Mach Intell*. 1998;20(12):1342–1351. Available from: [https://ieeexplore.ieee.org/document/](https://ieeexplore.ieee.org/document/735807/)
541 [735807/](https://ieeexplore.ieee.org/document/735807/).
- 542 41. Rasmussen CE, Williams CKI. Gaussian processes for machine learning. Cambridge, MA: MIT
543 Press; 2006.
- 544 42. Nickisch H, Rasmussen CE. Approximations for binary Gaussian process classification. *J Mach*
545 *Learn Res*. 2008;9(10):2035–2078.
- 546 43. Zhang Z, Dai G, Jordan MI. Bayesian generalized kernel mixed models. *J Mach Learn Res*.
547 2011;12:111–139.
- 548 44. Heckerman D, Gurdasani D, Kadie C, Pomilla C, Carstensen T, Martin H, et al. Linear mixed
549 model for heritability estimation that explicitly addresses environmental variation. *Proc Natl*
550 *Acad Sci U S A*. 2016;113(27):7377–7382. Available from: [http://www.pnas.org/content/113/](http://www.pnas.org/content/113/27/7377.abstract)
551 [27/7377.abstract](http://www.pnas.org/content/113/27/7377.abstract).
- 552 45. Swain PS, Stevenson K, Leary A, Montano-Gutierrez LF, Clark IBN, Vogel J, et al. Inferring time
553 derivatives including cell growth rates using Gaussian processes. *Nat Comm*. 2016;7(12):13766.
554 Available from: <https://doi.org/10.1038/ncomms13766>.
- 555 46. Cheng L, Ramchandran S, Vatanen T, Lietzén N, Lahesmaa R, Vehtari A, et al. An additive
556 Gaussian process regression model for interpretable non-parametric analysis of longitudinal data.
557 *Nat Comm*. 2019;10(1):1798. Available from: <https://doi.org/10.1038/s41467-019-09785-8>.
- 558 47. McDonald KR, Broderick WF, Huettel SA, Pearson JM. Bayesian nonparametric models char-
559 acterize instantaneous strategies in a competitive dynamic game. *Nat Comm*. 2019;10(1):1808.
560 Available from: <https://doi.org/10.1038/s41467-019-09789-4>.

- 561 48. Rodriguez-Nieva JF, Scheurer MS. Identifying topological order through unsupervised machine
562 learning. *Nat Phys*. 2019; Available from: <https://doi.org/10.1038/s41567-019-0512-x>.
- 563 49. Schölkopf B, Herbrich R, Smola AJ. A generalized representer theorem. In: *Proceedings of the*
564 *14th Annual Conference on Computational Learning Theory and and 5th European Conference on*
565 *Computational Learning Theory*. London, UK, UK: Springer-Verlag; 2001. p. 416–426. Available
566 from: <http://dl.acm.org/citation.cfm?id=648300.755324>.
- 567 50. Pillai NS, Wu Q, Liang F, Mukherjee S, Wolpert R. Characterizing the function space for Bayesian
568 kernel models. *J Mach Learn Res*. 2007;8:1769–1797.
- 569 51. Jiang Y, Reif JC. Modeling epistasis in genomic selection. *Genetics*. 2015;201:759–768.
- 570 52. Chaudhuri A, Kakde D, Sadek C, Gonzalez L, Kong S. The mean and median criteria for kernel
571 bandwidth selection for support vector data description. *Data Mining Workshops (ICDMW), 2017*
572 *IEEE International Conference on*. 2017;p. 842–849. Available from: [https://ieeexplore.ieee.](https://ieeexplore.ieee.org/abstract/document/8215749/)
573 [org/abstract/document/8215749/](https://ieeexplore.ieee.org/abstract/document/8215749/).
- 574 53. Singleton KR, Crawford L, Tsui E, Manchester HE, Maertens O, Liu X, et al. Melanoma thera-
575 peutic strategies that select against resistance by exploiting MYC-driven evolutionary convergence.
576 *Cell Rep*. 2017;21(10):2796–2812.
- 577 54. Crawford L, Wood KC, Zhou X, Mukherjee S. Bayesian approximate kernel regression with variable
578 selection. *J Am Stat Assoc*. 2018;113(524):1710–1721. Available from: [https://doi.org/10.](https://doi.org/10.1080/01621459.2017.1361830)
579 [1080/01621459.2017.1361830](https://doi.org/10.1080/01621459.2017.1361830).
- 580 55. Crawford L, Flaxman SR, Runcie DE, West M. Predictor variable prioritization in nonlinear
581 models: a genetic association case study. *Ann Appl Stat*. 2019;13(2):958–989. Available from:
582 <https://projecteuclid.org/euclid.aoas/1560758434>.
- 583 56. Fasy BT, Micka S, Millman DL, Schenfisch A, Williams L. Challenges in reconstructing shapes
584 from Euler characteristic curves. *arXiv*. 2018;p. 1811.11337.
- 585 57. Oudot S, Solomon E. Inverse problems in topological persistence. *arXiv*. 2018;p. 1810.10813.
586 Available from: <https://arxiv.org/abs/1810.10813>.
- 587 58. Gopalan G, Bornn L. FastGP: An R package for Gaussian processes. *arXiv*. 2015;p. 1507.06055.
588 Available from: <https://arxiv.org/abs/1507.06055>.
- 589 59. Adler D, Nenadic O, Zucchini W. RGL: an R-library for 3D visualization with OpenGL. In:
590 *Proceedings of the 35th Symposium of the Interface: Computing Science and Statistics, Salt Lake*
591 *City*. vol. 35; 2003. p. 1–11.
- 592 60. Schlager S, Zheng G, Li S, Székely G. 9. In: *Morpho and Rvcg – Shape analysis in R:*
593 *R-packages for geometric morphometrics, shape analysis and surface manipulations*. Academic
594 Press; 2017. p. 217–256. Available from: [http://www.sciencedirect.com/science/article/](http://www.sciencedirect.com/science/article/pii/B9780128104934000110)
595 [pii/B9780128104934000110](http://www.sciencedirect.com/science/article/pii/B9780128104934000110).
- 596 61. Schlager S, Profico A, Di Vincenzo F, Manzi G. Retrodeformation of fossil specimens based
597 on 3D bilateral semi-landmarks: implementation in the R package “Morpho”. *PloS one*.
598 2018;13(3):e0194073.
- 599 62. Puente J. Distances and algorithms to compare sets of shapes for automated biological morpho-
600 metrics. Princeton University. Princeton, NJ; 2013.

- 601 63. Belongie S. Rodrigues' rotation formula. From MathWorld—A Wolfram Web Resource, created by
602 Eric W Weisstein <http://mathworld.wolfram.com/RodriguesRotationFormula.html>. 1999;.
- 603 64. Crompton RH, Savage R, Spears IR. The mechanics of food reduction in *Tarsius bancanus*. Hard-
604 object feeder, soft-object feeder or both? *Folia Primatol* (Basel). 1998;69(Suppl 1):41–59.