Gaussian Process Landmarking for Three-Dimensional Geometric Morphometrics*

Tingran Gao†, Shahar Z. Kovalsky ‡, Doug M. Boyer §, and Ingrid Daubechies ¶

Abstract. We demonstrate applications of the Gaussian process-based landmarking algorithm proposed in [T. Gao, S.Z. Kovalsky, I. Daubechies 2018] to geometric morphometrics, a branch of evolutionary biology centered at the analysis and comparisons of anatomical shapes, and compares the automatically sampled landmarks with the "ground truth" landmarks manually placed by evolutionary anthropologists; the results suggest that Gaussian process landmarks perform equally well or better, in terms of both spatial coverage and downstream statistical analysis. We provide a detailed exposition of numerical procedures and feature filtering algorithms for computing high-quality and semantically meaningful diffeomorphisms between disk-type anatomical surfaces.

Key words. Gaussian Process, Experimental Design, Active Learning, Manifold Learning, Geometric Morphometrics

AMS subject classifications. 60G15, 62K05, 65D18

1. Introduction.

Computing one-to-one correspondences between objects, or *registration*, plays an important role in a wide range of scientific disciplines, such as functional Magnetic Resonance Imaging (fMRI) in medical studies [74], shape matching in computer graphics [98], remote sensing in geophysical sciences [65], to name just a few. When the number of objects in comparison is prohibitively large, or when the information encoded in each data object is noisy or redundant, a common strategy is to reduce the comparison to *alignment* or *assignment* between sets of extracted *features*, which are typically compact and informative representations of the original data constructed with domain knowledge. Examples of this sort include the *scale-invariant feature transform* (SIFT) in computer vision [59] and *anatomical landmarks* in statistical shape analysis [25] and computational anatomy [47].

This paper stems from an attempt to apply principles of the statistics field of *optimal experimental design* to automatically detect salient feature points on anatomical surfaces. For the biological morphometrical data of interest to us, these feature points surprisingly resemble the "ground truth" landmarks seasoned evolutionary anthropologists would identify as discrete anatomical points of correspondence. The proposed algorithm, based on successively selecting points of maximum *mean squared prediction error* (MSPE) with respect to a Gaussian process modeled on a smooth manifold or from a scattered point cloud, turns out to be intimately connected with a class of greedy algorithms in *reduced basis methods* [10, 24] and

*Submitted

Funding: This work is supported by Simons Math+X Investigators Award 400837 and NSF CAREER Award BCS-1552848.

†Computational and Applied Mathematics Initiative, Department of Statistics, The University of Chicago, Chicago IL (tingrangao@galton.uchicago.edu)

‡Department of Mathematics, Duke University, Durham, NC (shaharko@math.duke.edu)

§Department of Evolutionary Anthropology, Duke University, Durham NC 27708, USA (doug.boyer@duke.edu).

¶Department of Mathematics and Department of Electrical and Computer Engineering, Duke University, Durham NC (ingrid@math.duke.edu)

enjoys fast rate of convergence. Interested readers may consult our companion paper [30] for more details on these theoretical aspects. This paper demonstrates that, not only do these Gaussian process landmarks perform equally well or better compared with the ground truth landmarks (in terms of both spatial coverage and downstream statistical analysis), they can also be leveraged to establish high-quality and semantically meaningful smooth bijections (registration maps) between biological structures; shape distances induced by these maps reach comparable explanatory power to ground truth landmarks for ordination (i.e., discriminating species groups) purposes.

The effectiveness of the proposed Gaussian process landmarking algorithm can be intuitively interpreted by the analogy between Gaussian process experimental design and the landmark selection procedure in geometric morphometrics. In a nutshell, the method we propose considers a Gaussian process on the anatomical surface, with a variance-covariance structure specified by the heat kernel of the underlying Riemannian manifold; the landmarks are then selected successively, each time picking a new landmark as the point on the surface with largest variance conditioned on all the previously selected landmarks. This procedure is highly reminiscent of the practice of *landmarking* in *geometric morphometrics*, the subfield of evolutionary biology focusing on quantifying the (dis-)similarities between pairs of two-dimensional anatomical surfaces based on their spatial configurations [105]. The main objects of study in geometric morphometrics are anatomical surfaces, such as bones and teeth, of extinct and extant animals of particular interest to evolutionary biologists. The landmarking procedure typically starts with manually identifying an equal number of geometrically or semantically meaningful feature points, or *landmarks*, on each specimen in a potentially large collection of anatomical surfaces; the landmarks are certified by domain experts to be in consistent one-to-one correspondences. The methodology of landmark selection, though difficult to articulate and is still constantly under debate (see e.g., very recent discussions [100, 101] and the references therein), emphasizes a comprehensive and balanced decision between sharp geometric features (points of application of real biomechanical forces) and points extremal to spatial configuration (points taken "farthest away" from other points under certain metrics). While geometric features could be recognized in a computational geometry manner by computing discretized curvatures, the precise meaning of "extremality" can be flexibly characterized as an "uncertainly" naturally encoded into the covariance structure of a Gaussian process with a geometry-aware kernel function. By carefully designing a curvature-reweighted heat kernel on discrete anatomical surfaces, the proposed algorithm is thus capable of selecting both sharp geometric features and extremal points in one single pass. Downstream Procrustes analysis [40, 25, 41] results presented in this paper also speak of the biological relevance of the automatically generated landmarks.

The rest of this paper is organized as follows. Section 2 provides background materials on geometric morphometrics, Gaussian processes, and the Gaussian process landmarking algorithm proposed in [30]; Section 3 explains the importance of using a reweighted kernel in the Gaussian process landmarking algorithm for our application; Section 4 details all the numerical procedures in this application; Section 5 closes with comments and discussions.

2. Background.

2.1. Geometric Morphometrics: Old and New. Statistical shape analysis, often termed *geometric morphometrics* in the context of comparative biology, is the quantitative analysis of variations and correlations among biological forms through the Cartesian coordinates of "landmarks"—biologically informative, repeatable, and in some sense corresponding anatomical loci—on surfaces representing anatomy of biological organisms [102, 1, 105]. In order for the comparisons across specimens to be meaningful, practitioners in this field often require that the landmarks be consistently annotated on each specimen in a manner reflecting the "operational homology[1]" or "biological correspondence" (as discussed in [14]) across individuated traits inherited from a common ancestry [62]. For instance, in the *generalized Procrustes analysis* (GPA) framework [40, 25, 41, 2], the *Procrustes distance* between two surfaces S_1, S_2 is computed using the following procedure:

(i) Specify two sets of operationally homologous landmarks $\left\{x_\ell^{(1)} \mid 1 \leq \ell \leq L\right\}$ and $\left\{x_\ell^{(2)} \mid 1 \leq \ell \leq L\right\}$ on S_1 and S_2, respectively;

(ii) Compute the distance between S_1 and S_2 by minimizing the energy functional

$$d_{\mathrm{cP}}\left(S_1, S_2\right) = \inf_{T \in \mathbb{E}(3)} \left(\frac{1}{L} \sum_{\ell=1}^{L} \left\| T\left(x_\ell^{(1)}\right) - x_\ell^{(2)} \right\|^2 \right)^{\frac{1}{2}} \tag{2.1}$$

where $\mathbb{E}(3)$ is the group of rigid motions in $\mathbb{R}^3$.

This idea can be generalized to analyze a collection of consistently landmarked shapes, either assuming each set of landmarks on the same shape is centered at the origin so the variational problem is defined on a product space of orthogonal groups [41, 69, 94, 67], or estimate the optimal orthogonal and translation group elements jointly without the overall centering assumption [18].

The key to successfully applying the Procrustes framework in statistical shape analysis is to obtain an equal number of consistent, operationally homologous landmarks on every shape in a potentially enormous collection of shapes. Consistently landmarking a collection of shapes relies crucially upon domain knowledge and tedious manual labor, and the skill to perform it "correctly" typically requires years of professional training; even then the "correctness" can be subject to debate among experts (see e.g., [13] for an example on *Lepilemur* teeth). In the first place, extracting a finite number of landmarks from a continuous surface inevitably loses geometric information, unless when the shapes under consideration are easily seen to be uniquely determined by the landmarks (e.g., planar polygonal shapes, as considered in [48][26]), which is rarely the case for geometric morphometricians in biology; this problem of "inadequate coverage" motivated the introduction of *semilandmarks*—additional points along curves containing critical curvature information about the morphology—to compensate for the loss of geometry in the landmarking process. Unfortunately, the essential arbitrariness of the semilandmarks along a curve induces additional uncertainty that needs to be quantified and reduced [105, §2], especially in the absence of sharp anatomical features; the constraint of picking an equal number of landmarks on each shape also turns out to be far too artificial when the anatomical forms undergo complex evolutionary and developmental processes.

[1]The term "homology" in the context of evolutionary theory bears a different meaning than in modern topology; see e.g., [82, Part IV].

To mitigate both the scalability and the subjectivity issues in the existing Procrustes analysis framework, a recent trend of research in geometric morphometrics advocates *automated* workflows to bypass the repetitive, laborious, and time-consuming process of manual landmark placement; see e.g., [2, 13, 79, 55, 12, 56, 14, 27, 99, 51, 45] and the references therein. These techniques work directly with digitized anatomical surfaces represented as discrete triangular meshes; numerical algorithms are combined with computer graphics and geometry processing to provide high-throughput, landmark-free approaches for precise phenotyping [78, 43, 44] on the discrete triangular meshes on their entirety (often consisting of thousands to millions of vertices in $\mathbb{R}^3$), without the manual landmarking stage to filter down the number of variables using potentially biased *a priori* domain knowledge. Similar to the claim made by the proponents of landmark-based morphometrics that landmark coordinates contain more information than utilized in more traditional, measurement-based morphometrics [105, §2], the precursors of automated geometric morphometrics believe that using whole surfaces as input passes even more information to the downstream analysis.

Despite the capability of generating high quality pairwise shape registrations, automated geometric morphometric methods suffer from interpretability problems: since all comparisons are performed merely pairwise, composing the obtained correspondences along a closed loop does not give rise to an identity map in general. This lack of *transitivity* (see e.g., [31, 28]) demands additional post-processing steps to translate the pairwise results into valid input for standard downstream phylogenetic analysis [75, 32], which bears a strong similarity with recent studies in *synchronization problems* [18, 29, 73]. The loop inconsistency of pairwise correspondences also challenges the interpretability of automated geometric morphometrics, since it becomes virtually impossible to identify functionally equivalent regions across distinct anatomical structures in a consistent manner. Until fully automated geometric morphometric algorithms reach the maturity with comparable explanatory power to a human practitioner of landmark-based geometric morphometrics, deeper and more systematic understanding of the landmarking process still seem of great interest and value.

The methodology we propose in this paper incorporates an algorithmic landmarking procedure into the automated pairwise registration algorithms. We point out that detecting morphometrically meaningful landmarks on anatomical surfaces in a completely unsupervised manner is generally a daunting task, since some of the most reliable landmarks are determined by patterns of juxtapositions of tissues—termed "Type 1 landmarks" by Bookstein [83]—which are almost always absent on the triangular meshes input to automated algorithms [11]. The selection process is further complicated by the requirement of the consistency of relative landmark positions across the data collected, as well as the specific functionality of the biological organism being studied [105]. While geometry processing algorithms (e.g., [6, 58, 17, 16, 96]) are capable of detecting sharp geometric features characterized by metric or topology (Bookstein's "Type 2 landmarks" [83]), as well as producing high quality pairwise registrations for accurate determination of operationally homologous loci, semilandmarks and a majority of "Type 3 landmarks" in Bookstein's typology of landmarks are marked simply for adequate and/or comprehensive coverage of the anatomical forms [83]. For instance, some Type 3 landmarks are included in the analysis for being "furthest away" from sharp geometric or functional features [105]. These observations motivated us to consider algorithmic analogies of geometric morphologists' daily practice beyond the scope of computational geom-

etry, shedding light upon landmark identification from the perspective of Bayesian statistics quantifying the "uncertainty" of morphometric analysis. As will be detailed in Section 4, we will first generate a set of candidate landmarks on each of the anatomical surfaces based on the "uncertainty" modeled by a Gaussian process, then apply a matching scheme that filters out non-corresponding candidate landmarks between a pair of surfaces based upon bounded conformal distortion [54].

2.2. Gaussian Processes. A *Gaussian process* (or *Gaussian random field*) on a Polish space M with mean function $m : M \to \mathbb{R}$ and covariance function $K : M \times M \to \mathbb{R}$ is defined as the stochastic process of which any finite marginal distribution on n fixed points $x_1, \cdots, x_n \in M$ is a multivariate Gaussian distribution with mean vector

$$m_n := (m(x_1), \cdots, m(x_n)) \in \mathbb{R}^n$$

and variance-covariance matrix

$$K_n := \begin{pmatrix} K(x_1, x_1) & \cdots & K(x_1, x_n) \\ \vdots & & \vdots \\ K(x_n, x_1) & \cdots & K(x_n, x_n) \end{pmatrix} \in \mathbb{R}^{n \times n}. \tag{2.2}$$

A Gaussian process with mean function $m : M \to \mathbb{R}$ and covariance function $K : M \times M \to \mathbb{R}$ will be denoted as $\mathrm{GP}(m, K)$. Under model $Y \sim \mathrm{GP}(m, K)$, given observed values $y_1, \cdots, y_n$ at locations $x_1, \cdots, x_n$, the *best linear predictor* (BLP) [95, 85] for the random field at a new point x is given by the conditional expectation

$$Y^*(x) := \mathbb{E}[Y(x) \mid Y(x_1) = y_1, \cdots, Y(x_n) = y_n] = m(x) + k_n(x)^\top K_n^{-1}(Y_n - m_n) \tag{2.3}$$

where $Y_n = (y_1, \cdots, y_n)^\top \in \mathbb{R}^n$, $k_n(x) = (K(x, x_1), \cdots, K(x, x_n))^\top \in \mathbb{R}^n$; at any $x \in M$, the expected squared error, or *mean squared prediction error* (MSPE), is defined as

$$\begin{aligned} \mathrm{MSPE}(x) :&= \mathbb{E}\left[(Y(x) - Y^*(x))^2\right] \\ &= \mathbb{E}\left[(Y(x) - \mathbb{E}[Y(x) \mid Y(x_1) = y_1, \cdots, Y(x_n) = y_n])^2\right] \\ &= K(x, x) - k_n(x)^\top K_n^{-1} k_n(x) \end{aligned} \tag{2.4}$$

which is a function over M. Here the expectation is with respect to all realizations $Y \sim \mathrm{GP}(m, K)$. Squared integral (L^2) or sup (L^∞) norms of the pointwise MSPE are often used as a criterion for evaluating the prediction performance over the experimental domain. In geospatial statistics, interpolation with (2.3) and (2.4) is known as *kriging*.

2.3. Gaussian Process Landmarking. We now provide a quick summary of the *Gaussian process landmarking* algorithm from [30]. The input to this algorithm is a triangular mesh $G = (V, E, F)$. Denote the set of vertices V as $\{x_1, \cdots, x_{|V|}\} \subset \mathbb{R}^3$, and construct the *discrete heat kernel*

$$K = (K_{ij})_{1 \le i,j \le n} := \left[\exp\left(-\frac{\|x_i - x_j\|^2}{t}\right)\right]_{1 \le i,j \le n}. \tag{2.5}$$

We calculate the mean and Gaussian curvature functions $\eta : V \to \mathbb{R}$, $\kappa : V \to \mathbb{R}$ on the triangular mesh (V, E, F) using standard algorithms in computational differential geometry [21, 3]. On a two-dimensional surface S, Gaussian curvature $\kappa(x)$ at a point $x \in S$ is the product of the two principal curvatures $k_1(x)$, $k_2(x)$ at x, i.e., $\kappa(x) = k_1(x) k_2(x)$, and the mean curvature is defined as $\eta(x) = k_1(x) + k_2(x)$. For any $\lambda \in [0, 1]$ and $\rho > 0$, define the value of the weight function $w = w_{\lambda,\rho}$ at each vertex x_i by

$$w_{\lambda,\rho}(x_i) = \frac{\lambda |\kappa(x_i)|^\rho}{\sum_{k=1}^{|V|} |\kappa(x_k)|^\rho \nu(x_k)} + \frac{(1-\lambda) |\eta(x_i)|^\rho}{\sum_{k=1}^{|V|} |\eta(x_k)|^\rho \nu(x_k)}, \qquad \forall x_i \in V \tag{2.6}$$

where $\nu(x_k)$ is the area of the Voronoi cell centered at x_i on the triangular mesh T, W is a diagonal matrix of size $n \times n$ with $w(x_k)\nu(x_k)$ at its k-th diagonal entry $(1 \le k \le |V|)$, and K is the discrete squared exponential kernel matrix (2.5). The parameter λ controls the relative proportion of mean and Gaussian curvatures, and the parameter ρ adjusts the sharpness of the energy surfaces of mean and Gaussian curvatures before they are combined to form the weight function. Denoting W for the diagonal matrix with weight $w_{\lambda,\rho}(x_k)$ at the kth diagonal element, we define the *reweighted heat kernel* $k_t^{w_{\lambda,\rho}}$ on $V \times V$ as

$$k_t^{w_{\lambda,\rho}}(x_i, x_j) = \sum_{k=1}^{|V|} k_{t/2}(x_i, x_k) k_{t/2}(x_k, x_j) w_{\lambda,\rho}(x_k) \nu(x_k) =: \left(\tilde{K}^\top W \tilde{K}\right)_{ij} \tag{2.7}$$

where the unweighted kernel $k_{t/2}$ is calculated as in (2.5) but with bandwidth parameter $t/2$ instead of t, i.e.,

$$\tilde{K} := \left(e^{-\frac{\left\|x_i - x_j\right\|^2}{t/2}} \right)_{1 \le i,j \le |V|} \in \mathbb{R}^{|V| \times |V|}. \tag{2.8}$$

Note that, even for the trivial weight function $w \equiv 1$ (i.e., when W is the identity matrix of size $|V| \times |V|$), (2.7) does not reduce to (2.5) since it is easy to check that $\tilde{K}^2 \neq K$; this is in stark contrast with the reproducing case in the continuous regime in [30, (2.8)]. Note only are heat kernels natural choices for specifying Gaussian processes on Riemannian manifolds (see e.g., a brief survey in [30, §2.1]), their infinite smoothness also ensures fast convergence of the proposed Gaussian process landmarking algorithm; see [30, Theorem 4.6].

We now detail the greedy landmarking procedures. Until a fixed total number of landmarks are collected, at step $(n+1)$ the algorithm computes the uncertainty score $\Sigma_{(k)}$ on V from the existing n landmarks $\xi_1, \cdots, \xi_n$ by

$$\Sigma_{(n+1)}(x_i) = k_t^{w_{\lambda,\rho}}(x_i, x_i) - k_t^{w_{\lambda,\rho}}\left(x_i, \xi_n^1\right)^\top k_t^{w_{\lambda,\rho}}\left(\xi_n^1, \xi_n^1\right)^{-1} k_t^{w_{\lambda,\rho}}\left(x_i, \xi_n^1\right) \tag{2.9}$$

for all $x_i \in V$, where

$$k_t^{w_{\lambda,\rho}}\left(x_i, \xi_n^1\right) := \left(k_t^{w_{\lambda,\rho}}(x_i, \xi_1), \cdots, k_t^{w_{\lambda,\rho}}(x_i, \xi_n)\right)^\top,$$

$$k_t^{w_{\lambda,\rho}}\left(\xi_n^1, \xi_n^1\right) := \begin{pmatrix} k_t^{w_{\lambda,\rho}}(\xi_1, \xi_1) & \cdots & k_t^{w_{\lambda,\rho}}(\xi_1, \xi_n) \\ \vdots & & \vdots \\ k_t^{w_{\lambda,\rho}}(\xi_n, \xi_1) & \cdots & k_t^{w_{\lambda,\rho}}(\xi_n, \xi_n) \end{pmatrix},$$

and pick the $(n+1)$-th landmark ξ_{n+1} according to the rule

$$\xi_{n+1} = \underset{x_i \in V}{\operatorname{argmax}} \Sigma_{(n+1)}(x_i).$$

If there exist more than one maximizers of $\Sigma_{(n+1)}$, we just randomly pick one; at step 1 the algorithm simply picks the vertex maximizing $x \mapsto k_t^{w_{\lambda,\rho}}(x,x)$ on V. See Algorithm 2.1 for a comprehensive description. We stick in this paper with default values $\lambda = 1/2$ and $\rho = 1$ in all examples and applications, but one may wish to alter these values to fine-tune the landscape of the weight function for a specific application. The bandwidth parameter t is set to be the average edge length of the triangular mesh, following standard practices of manifold learning (see e.g., [22]). More details can be found in [30, §3]; we only mention here that (2.6) and (2.7) are discretizations of

$$w_{\lambda,\rho}(x) = \frac{\lambda |\kappa(x)|^\rho}{\int_M |\kappa(\xi)|^\rho \operatorname{dvol}_M(\xi)} + \frac{(1-\lambda)|\eta(x)|^\rho}{\int_M |\eta(\xi)|^\rho \operatorname{dvol}_M(\xi)}, \quad \forall x \in M \tag{2.10}$$

and

$$k_t^w(x,y) = \int_M k_{t/2}(x,z)\, k_{t/2}(z,y)\, w(z) \operatorname{dvol}_M(z), \quad \forall x,y \in M \tag{2.11}$$

respectively. Since Algorithm 2.1 builds upon the heat kernel and curvature of the underlying Riemannian manifold, the landmarks are stable under isometric deformations; we expect them to be stable across near-isometric Riemannian manifolds as well, which are objects of main interest to applications in geometric morphometrics.

Guided by uncertainty specified through such curvature-reweighted covariance structure, the Gaussian process landmarking often identifies landmarks of abundant biological information—for instance, the first few Gaussian process landmarks are often highly biologically informative, and demonstrate comparable level of coverage with the observer landmarks manually picked by human experts. See Figure 1 for a visual comparison between the automatically generated landmarks with the observer landmarks manually placed by evolutionary anthropologists on a different digitized fossil molar.

Algorithm 2.1 greedily picks as the next landmark the point on the manifold M that maximizes the pointwise prediction error. This sequential optimization approach is reminiscent of a popular approximation scheme in entropy-based experimental design [53, 71, 85]. The key observation here is the equivalence between minimizing the *conditional entropy* (or maximizing the *information gain*) and maximizing the determinant of the marginal covariance matrix at the design points; see e.g., [85, §6.2.1], [53, §3.1], [71, §2.2]. More concretely, if we denote the maximum entropy of a Gaussian process $\mathrm{GP}(m,K)$ on a manifold M with respect to any n observations as

$$\mathrm{OPT}(n) := \max_{\{x_1,\cdots,x_n\} \subset M} \det K_n$$

where K_n is the variance-covariance matrix defined in (2.2), and write $\mathrm{GPL}(n)$ for the value

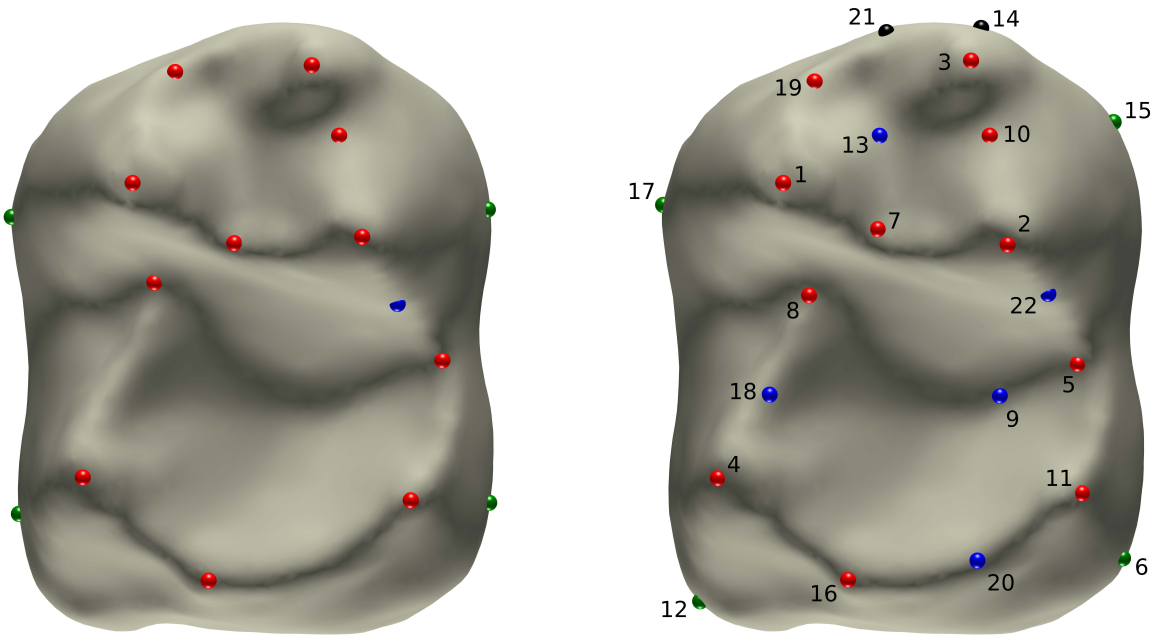

(a) Observer Landmarks **(b) Gaussian Process Landmarks**

Figure 1: **Left:** Observer landmarks on a digitized fossil molar of a *Teilhardina* (one of the oldest known fossil primates closely related with living tarsiers and anthropoids [7]) identified manually by an evolutionary anthropologist as ground truth, first published in [13]. **Right:** Landmarks picked by Gaussian Process Landmarking (Algorithm 2.1). The numbers next to each landmark indicate the order of appearance. These automatically computed landmarks strikingly resemble the observer landmarks; the 22 first Gaussian process landmarks already capture the same features selected by an expert anthropologist: the red landmarks (Number 1-5, 7, 8, 10, 11, 16, 19) signal sharp geometric features (cusps or saddle points corresponding to local maximum/minimum Gaussian curvature); the blue landmarks sit either along the curvy cusp ridges and grooves (Number 13, 18, 20, 22) or at the basin (Number 9), serving the role often played by semilandmarks; the four green landmarks (Number 6, 12, 15, 17) approximately delimit the "outline" of the tooth in occlusal view. The ordering of the Gaussian process landmarks does not affect the bounded distortion surface registration algorithm in Subsection 4.2.

of $\det K(X_1, \cdots, X_n)$ where $X_1, \cdots, X_n$ are generated using Algorithm 2.1, then by the *submodularity* of the entropy function [50] we conclude, using classic results [68], that

$$\mathrm{OPT}(n) \geq \mathrm{GPL}(n) \geq (1 - 1/e)\,\mathrm{OPT}(n).$$

In other words, the entropy of the greedy algorithm is equivalent to the maximum entropy up to a multiplicative constant. While this general submodularity-based framework justifies the information-theoretic asymptotic near-optimality of Gaussian Process Landmarking (Algorithm 2.1), in [30] we established a stronger optimality result in the context of Gaussian processes: up to a multiplicative constant, the MSPE with respect to n Gaussian process

Algorithm 2.1 Gaussian Process Landmarking with Reweighted Heat Kernel

procedure GPL(T, L, $\lambda \in [0,1]$, $\rho > 0$, $\epsilon > 0$)▷ Triangular Mesh $T = (V, E)$, number of landmarks L
  $\kappa, \eta \leftarrow$ DiscreteCurvatures(T) ▷ calculate discrete Gaussian curvature κ and mean curvature η on T
  $\nu \leftarrow$ VoronoiAreas(T) ▷ calculate the area of Voronoi cells around each vertex x_i
  $w_{\lambda,\rho} \leftarrow$ CalculateWeight($\kappa, \eta, \lambda, \rho, \nu$) ▷ calculate weight function $w_{\lambda,\rho}$ according to (2.10)
  $W \leftarrow \left[\exp\left(-\left\|x_i - x_j\right\|^2 \Big/ \epsilon\right)\right]_{1\le i,j\le |V|} \in \mathbb{R}^{|V|\times|V|}$
  $\Lambda \leftarrow \operatorname{diag}\left(w_{\lambda,\rho}(x_1)\,\nu(x_1), \cdots, w_{\lambda,\rho}\left(x_{|V|}\right)\nu\left(x_{|V|}\right)\right) \in \mathbb{R}^{|V|\times|V|}$
  $\xi_1, \cdots, \xi_L \leftarrow \emptyset$ ▷ initialize landmark list
  $\Psi \leftarrow 0$
  $\ell \leftarrow 1$
  $K_{\text{full}} \leftarrow W^\top \Lambda W \in \mathbb{R}^{|V|\times|V|}$
  $K_{\text{trace}} \leftarrow \operatorname{diag}(K_{\text{full}}) \in \mathbb{R}^{|V|}$
  while $\ell < L + 1$ **do**
    if $\ell = 1$ **then**
      $\Sigma \leftarrow K_{\text{trace}}$
    else
      $\Sigma \leftarrow K_{\text{trace}} - \operatorname{diag}\left(\Psi^\top\left(\Psi\left[[\xi_1, \cdots, \xi_\ell], :\right] \setminus \Psi\right)\right) \in \mathbb{R}^{|V|}$ ▷ calculate uncertainty scores by (2.9)
    end if
    $\xi_\ell \leftarrow \operatorname{argmax}\Sigma$
    $\Psi \leftarrow K_{\text{full}}[:, [\xi_1, \cdots, \xi_\ell]]$
    $\ell \leftarrow \ell + 1$
  end while
  return $\xi_1, \cdots, \xi_L$
end procedure

landmarks coincides with the optimal MSPE attainable over all sets of n points on M, for any positive integer $n \in \mathbb{Z}_{\ge 0}$. More details about this can be found in [30, §4].

3. The Role of Reweighting in the Kernel Construction. An essential construction in the application of Algorithm 2.1 to geometric morphometrics is the reweighted kernel (2.11). Intuitively, the reweighting step modifies the Euclidean heat kernel (2.5) by amplifying the influence of locations with relatively high weights. To investigate the role played by the reweighted kernel in greater detail, we study the behavior of the reweighted kernel in the asymptotic regime when the number of samples increases to infinity. As will be demonstrated in Theorem 3.1, the reweighted kernel defines a diffusion process on the manifold with a backward Kolmogorov operator conjugate to the Witten Laplacian, of which the eigenfunctions has desired localization behavior near critical points of the weight function.

To motivate Theorem 3.1, let us consider i.i.d. (with respect to the standard volume

measure) samples $\{x_i\}_{i=1}^N$ on a closed Riemannian manifold (M, g), as well as an arbitrary function $f \in C^2(M)$. To simplify the discussion, assume for the moment that the samples are uniformly distributed on M with respect to the normalized volume form on M. It is well known in manifold learning (see e.g., [89, 22]) that, for any $x \in M$,

$$\frac{1}{N}\sum_{j=1}^{N} e^{-\frac{\left\|x-x_j\right\|^2}{2\epsilon}} f(x_j) \longrightarrow \frac{1}{\mathrm{Vol}(M)}\int_M e^{-\frac{\|x-y\|^2}{2\epsilon}} f(y)\,\mathrm{dVol}(y) \tag{3.1}$$

as $n \to \infty$. Denote the positive weight function (2.10) used in (2.11) by

$$w(x) = e^{-V(x)}, \quad V(x) \geq 0 \qquad \forall x \in M. \tag{3.2}$$

Note that the normalization in (2.10) ensures that the value of the weight function is bounded between 0 and 1, which is implied in (3.2), for numerical stability. Repeatedly using the "law of large numbers" argument (3.1), we have the following convergence for the reweighted kernel:

$$\frac{1}{N^2}\sum_{k=1}^{N} e^{-\frac{\|x-x_k\|^2}{2\epsilon}} e^{-V(x_k)} \sum_{j=1}^{N} e^{-\frac{\left\|x_k-x_j\right\|^2}{2\epsilon}} f(x_j) \longrightarrow$$
$$\frac{1}{[\mathrm{Vol}(M)]^2}\int_M\int_M e^{-\frac{\|x-z\|^2}{2\epsilon}} e^{-V(z)} e^{-\frac{\|z-y\|^2}{2\epsilon}} f(y)\,\mathrm{dVol}(z)\,\mathrm{dVol}(y) \quad \text{as } N \to \infty.$$

This motivates us to consider the asymptotic behavior of the integral operator on the right hand side in the asymptotic regime $\epsilon \to \infty$, a recurring theme in manifold learning (see e.g., [8, 22, 90, 28]). For simplicity of notation, we will denote the integral operator with respect to kernel $K : M \times M \to \mathbb{R}$ as

$$T_K f(x) := \frac{1}{|\mathrm{Vol}(M)|}\int_M K(x,y) f(y)\,\mathrm{dVol}(y).$$

Theorem 3.1. *Let M be a d-dimensional closed manifold. For any smooth functions $f, V \in C^\infty(M)$ with $V \geq 0$, for kernel function*

$$K_V^\epsilon(x,y) := \frac{1}{|\mathrm{Vol}(M)|}\int_M e^{-\frac{\|x-z\|^2}{2\epsilon}} e^{-V(z)} e^{-\frac{\|z-y\|^2}{2\epsilon}}\,\mathrm{dVol}(z) \tag{3.3}$$

we have

$$\frac{T_{K_V^\epsilon} f(x)}{T_{K_V^\epsilon} 1} \to f(x) + \epsilon\Big(\Delta f(x) - \nabla f(x)\cdot\nabla V(x)\Big) + O\left(\epsilon^{\frac{3}{2}}\right) \quad \text{as } \epsilon \to 0 \tag{3.4}$$

where 1 stands for the constant function on M taking value 1.

The proof of Theorem 3.1 can be found in Appendix A. Theorem 3.1 indicates that a proper normalization of the reweighted kernel (2.11) gives rise to an approximation to the heat kernel of the backward Kolmogorov operator

$$L = -\Delta + \nabla V \cdot \nabla.$$

Note that this operator is also the infinitesimal generator of the diffusion process determined by the stochastic differential equation

$$\mathrm{d}X_t = -\nabla V\left(X_t\right)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}W_t$$

where W_t is the standard Wiener process defined on M. In particular, Theorem 3.1 suggests that our construction of the reweighted kernel encodes information from the weight function into the dynamics of the stochastic process on M.

Remark 3.2. It is interesting to compare the result of Theorem 3.1 with the related but different kernel construction in diffusion maps [22, 66]. The integral operator in (3.4) is a properly normalized version of integrating a smooth function against the kernel

$$\widetilde{K}_V\left(x,y\right) := \int_M e^{-\frac{\|x-z\|^2}{2\epsilon}} e^{-V(z)} e^{-\frac{\|z-y\|^2}{2\epsilon}}\,\mathrm{dVol}\left(z\right)$$

which is obtained from sandwiching $\exp\left[-V\left(\cdot\right)\right]$ with the squared exponential kernel. If we pick a different order of sandwiching, namely, construct the kernel as

$$K_V\left(x,y\right) := e^{-V(x)} e^{-\frac{\|x-y\|^2}{2\epsilon}} e^{-V(y)},$$

then a quick computation in the same spirit as [9, Theorem 1] leads to

$$\frac{T_{K_V} f\left(x\right)}{T_{K_V} 1} \to f\left(x\right) + \epsilon\left(\frac{1}{2}\Delta f\left(x\right) - \nabla f\left(x\right)\cdot\nabla V\left(x\right)\right) + O\left(\epsilon^{\frac{3}{2}}\right) \quad \text{as } \epsilon \to 0.$$

In other words, the infinitesimal generator differs from the one calculated in (3.4) only by a multiplicative factor in front of the Laplace-Beltrami operator. We actually observe very similar numerical results when replacing $\widetilde{K}_V$ with K_V in Algorithm 2.1, when appropriate bandwidth parameters are chosen. Empirically, we find the bandwidth parameter easier to tune for kernel $\widetilde{K}_V$ than for K_V.

With these preparation, we now illustrate the effect of kernel reweighting by replacing (3.2) with

$$w_\epsilon\left(x\right) = e^{-V(x)/\epsilon}, \quad V\left(x\right) \geq 0 \qquad \forall x \in M. \tag{3.5}$$

Note that we can always redefine the potential function V so that (3.2) is written in the form of (3.5), for any fixed $\epsilon > 0$; in particular, replacing $V\left(x\right)$ by $V\left(x\right)/\epsilon$ does not change its critical points. An almost identical calculation as in the proof of Theorem 3.1 leads to

Corollary 3.3. *Under the same assumptions as in Theorem 3.1,*

$$\frac{T_{K^\epsilon_{V/\epsilon}} f\left(x\right)}{T_{K^\epsilon_{V/\epsilon}} 1} \to f\left(x\right) + \epsilon\Big(\Delta f\left(x\right) - \nabla f\left(x\right)\cdot\frac{1}{\epsilon}\nabla V\left(x\right)\Big) + O\left(\epsilon^{\frac{3}{2}}\right) \quad \textit{as } \epsilon \to 0 \tag{3.6}$$

where K^ϵ_V/ϵ is obtained from (3.3) by replacing the potential V with $\frac{1}{\epsilon}V$.

Similar to Theorem 3.1, Corollary 3.3 shows that a proper normalization of the kernel reweighted by (3.5) approximates the heat kernel of the operator

$$L_\epsilon := -\Delta + \frac{1}{\epsilon}\nabla V \cdot \nabla.$$

The dependence on ϵ is of particular interest. It is well known in the literature of semi-classical analysis that ϵL_ϵ is conjugate to the *semiclassical Witten Laplacian* [103] on 0-forms:

$$\Delta_{V,\epsilon} = \epsilon e^{-V/2\epsilon} \left(\epsilon L_\epsilon\right) e^{V/2\epsilon}$$

where

$$\Delta_{V,\epsilon} = -\epsilon^2 \Delta + \frac{1}{4}\left|\nabla V\right|^2 - \frac{\epsilon}{2}\Delta V.$$

Since the eigenfunctions of $\Delta_{V,\epsilon}$ corresponding to the leading small eigenvalues concentrate near the critical points of the potential function V for sufficiently small $\epsilon > 0$ (see e.g., [46] or [77, Theorem 3.9] for mathematically more precise statements), the eigenfunctions corresponding to the leading small eigenvalues of L_ϵ also concentrate near the critical points of V after being multiplied by $e^{-V/2\epsilon}$, which, by the proof of Theorem 3.1, can be approximated by the square root of the denominator of (3.6). Note that the matrix K_{full} in Algorithm 2.1 corresponds to the integral kernel in the numerator of (3.6) if we choose Λ to be the diagonal matrix with $\exp\left[-V\left(v_i\right)/\epsilon\right]$ at its i-th diagonal entry, where v_i is the i-th vertex on the triangular mesh; setting D to be the diagonal matrix with the i-th row sum of K_{full} at its i-th diagonal entry, the Witten Laplacian $\Delta_{V,\epsilon}$ can then be approximated by $D^{-1/2}K_{\text{full}}D^{-1/2}$ up to a scalar multiplication by the bandwidth ϵ. We would thus expect the first few eigenfunctions corresponding to the smallest eigenvalues of $D^{-1/2}K_{\text{full}}D^{-1/2}$ to concentrate near the critical points of V. This can be easily verified with numerical experiments; see Figure 2 for an example.

Remark 3.4. The localization effects on eigenfunctions can also be achieved by adding a diagonal matrix, which represents a potential function on the domain of interest, to the kernel matrix; this idea has found applications in biomedical data analysis [23] and computer graphics [61]. The insight we gained from relating the reweighted kernel to Witten Laplacian reveals that the localization can be equivalently realized through a simple nonlinear procedure. This phenomenon can be explained as follows: while adding a diagonal potential to the kernel provides a direct discretization of the Schödinger operator $L_V = -\Delta + V$, the reweighted kernel (2.11) can be thought of as a discretization of the heat kernel of L_V by taking only "one-hop" paths in the Feynman-Kac formula for e^{-tL_V}. An extensive discussion along this direction is beyond the scope of the current paper; interested readers may find useful the works [92, 93, 42] and the references therein.

The connection between the reweighted kernel (2.11) and the semi-classical analysis of the Witten Laplacian provides insights for the behavior of the Gaussian process landmarks generated from Algorithm 2.1. Given a Gaussian process $\mathrm{GP}\left(m, K\right)$ defined on a manifold M, the eigenfunctions of K—properly reweighted by their corresponding eigenvalues— gives

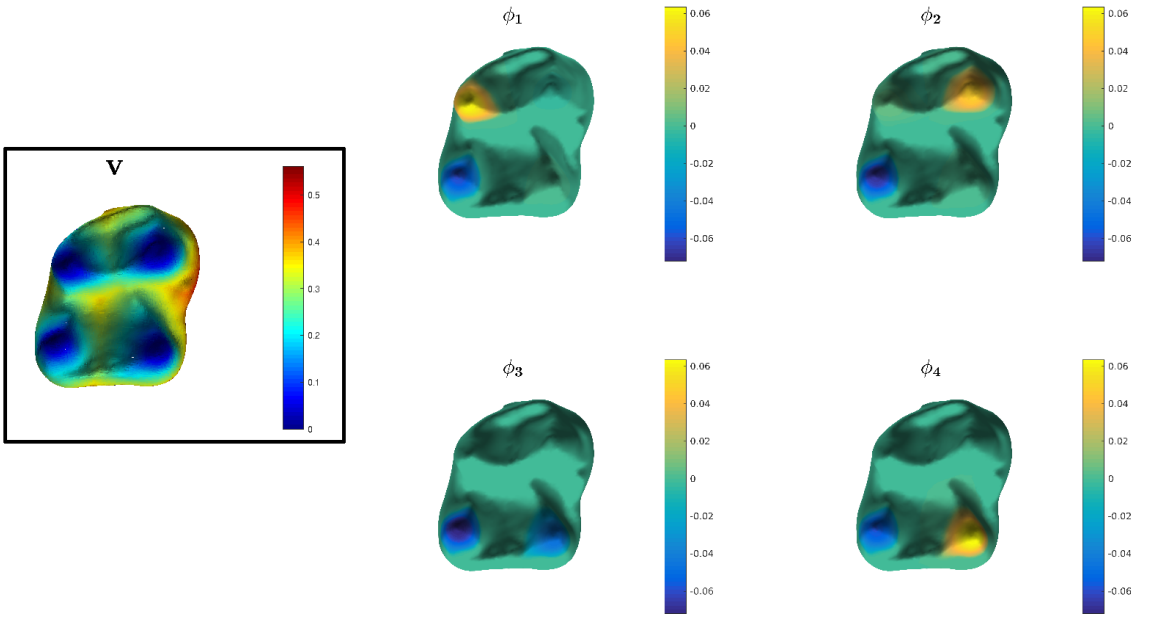

Figure 2: Concentration of the eigenfunctions of the Witten Laplacian with respect to a potential function V. The left boxed subplot illustrates the potential function V using a heat map; we constructed this potential to have 4 wells centered around 4 cusps manually picked on a *Lorisidae* mandibular molar. The 4 eigenfunctions corresponding to the 4 smallest eigenvalues of the Witten Laplacian with potential V are depicted on the 2×2 panel on the right; it can be read off from the colorbars that the support of each of these 4 eigenfunctions are concentrated around critical points of V.

rise to the Karhunen-Loève basis for $\mathrm{GP}(m, K)$, with respect to which the sample paths of $\mathrm{GP}(m, K)$ adopt expansions with i.i.d. standard normal coefficients. If the low-frequency components of these expansions tend to concentrate at certain regions on M, when fitting an unknown function in $\mathrm{GP}(m, K)$ using an active learning procedure, it could be much more efficient if one begins with the inquiry for the function value at those regions of concentration. After information at regions of concentration are collected to some extent, it becomes more beneficial to also incorporate ambient information that resides in the complement of these regions in M. Therefore, spreading landmarks on M using a Gaussian process with reweighted kernel balances out the information prioritized by the weight function and the space-filling experimental design strategies. The main reason of not selecting landmarks solely based on the weight function is that the weight function can be too spurious to produce reliable, semantically meaningful landmarks, at least for our application in geometric morphometrics; see Figure 3 for an example for such a comparison.

4. Gaussian Process Landmarking for Automated Geometric Morphometrics. This section is divided into three parts. Subsection 4.1 compares Gaussian process landmarks generated from Algorithm 2.1 with "ground truth" landmarks manually picked by comparative biologists on a real dataset of anatomical surfaces, demonstrating comparable levels of coverage

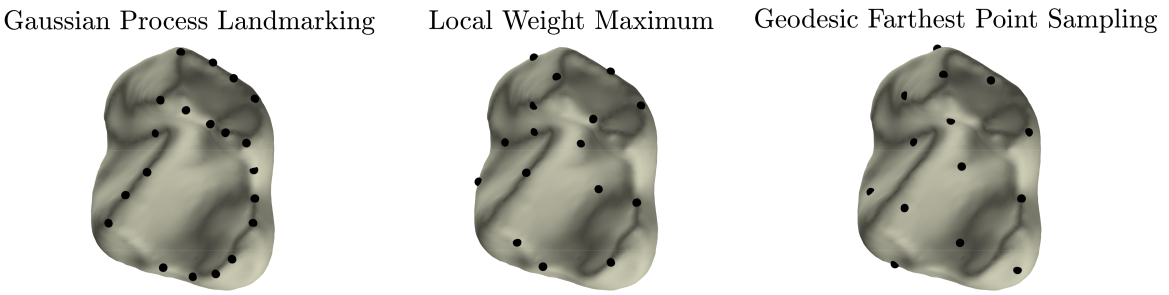

Figure 3: An equal number (20) of landmark points on the same surface generated using three different strategies. **Left:** Gaussian process landmarks generated using Algorithm 2.1 with a exponential squared kernel function reweighted by weight function w; **Middle:** The local maxima of the weight function used in the reweighed kernel, which appears semantically less meaningful from the perspective of geometric morphometrics; **Right:** Points produced by *geodesic farthest point sampling* (see e.g., [38, 64]), a greedy algorithm commonly used for generating uniformly sampled (or approximately space-filling) points on a triangular mesh. Comparing these three sampling approaches, Gaussian process landmarks has the advantage of tending to fill up the manifold while prioritizing the choice of semantically meaningful features for geometric morphometrics.

of biologically significant features. The results presented in this section provide quantitative evidence that Gaussian process landmarks are capable of capturing geometric features encoding important information for comparative biologists on individual anatomical surfaces. Subsection 4.2 adapts the image feature matching algorithm of [57] for the registration of pairs of anatomical surfaces via matching Gaussian process landmarks computed on each individual surface. We compare the resulting shape correspondence maps with a baseline obtained using previously developed *continuous Procrustes analysis* in [2, 13]; the results suggest that, though Gaussian process landmarks are generated on each individual shape separately, they implicitly encode operationally homologous features that can be compared and contrasted across shapes, and such pairwise comparison results are comparable with those obtained from standard Procrustes shape analysis based on "ground truth" observer landmarks. The pairwise shape distances induced from matching landmarks are turned into ordination plots (two-dimensional embeddings of distance matrices, commonly employed for visualizing the "morphospace" characterizing shape variances) in Appendix C; it turns out that Gaussian process landmarks lead to a favorable ordination, which is in better agreement with observations made by comparative biologists and paleontologists in existing literature.

We remark that the idea of using landmarks to guide the computation of shape correspondence maps is very natural in the geometry processing community. For instance, [72] proposes to choose landmarks based on a conditional number quantifying the stability of the matching problem for a pair of shapes, and demonstrated significantly enhanced map quality with very few landmarks specified in this manner; [49] generates small numbers of feature points with large fractions of semantic correspondences between shapes, and use these points to guide the search for a highly complicated map composed from blending multiple simpler maps. These methods, as well as many other state-of-the-art techniques in geometry processing, are not

suitable for general geometric morphometrics applications as they are often fail to capture all anatomical correspondences indispensable for repeatable and reliable analysis.

4.1. Unsupervised Landmarking on Individual Anatomical Surfaces. Gaussian process landmarks are generated on each anatomical surface individually, regardless of the total size of the shape collection; the landmarks manually selected by human experts (observers), however, may well depend on the information gradually exposed to the human expert as he/she moves through a collection of surfaces. It is thus surprising that the individually generated, "local" Gaussian process landmarks bear striking similarity with the observer landmarks selected with certain extent of collection-wise or "global" knowledge, as illustrated in [30, Figure 2] through an example fossil molar. This subsection is devoted to a more thorough and quantitative comparison between Gaussian process landmarks and "ground truth" observer landmarks placed by human experts.

We begin by collecting in Figure 4 results obtained by applying the Gaussian Process Landmarking algorithm to several different types of anatomical surfaces, including a subset of the second mandibular molars of primates (first published in [13]) and a subset of the astragalus and calcaneus bones of tarsiers first published in [14]. It can be recognized from Figure 4 that the algorithm is capable of consistently capturing both geometric features ("Type 3 landmarks" [83]) and semilandmarks (delineating ridges and grooves).

To quantitatively validate the biological informativeness of Gaussian process landmarks, we use a dataset of second mandibular molars on which the "observer landmarks," or the landmarks manually selected by experienced comparative biologists, are readily available (see [13]). We calculate, on each anatomical surface, the median geodesic distance from an observer landmark to its closest Gaussian process landmark. This median geodesic distance can obviously be calculated for any other type of landmarks in place of the Gaussian process landmarks; we shall refer to it as the *median observer-to-automatic landmark distance*, for the sake of simplicity. We compute the median observer-to-automatic landmark distances for a varying number of automatic landmarks, obtaining a curve that encodes the rate with which the median observer-to-automatic landmark distances decay to zero as the number of automatic landmarks increases. Comparing curves obtained from different automatic landmark generation methods then provides us with a way to evaluate how closely each type of automation "mimics" the observer landmarks. When we perform such curve comparisons for a large collection of surfaces in a dataset, one can in principle construct statistics (e.g., "mean" or "standard deviation" in the "space of curves") along the lines of *Functional Data Analysis* (FDA, see e.g., [81, 80] and the references therein), but we will have to defer such a statistically systematic treatment to future work.

Using this strategy, we compare Gaussian process landmarks with landmarks obtained from *Geodesic Farthest Point Sampling* (GFPS) [38, 64], a popular downsampling technique in automated geometric morphometrics, on a dataset of 116 second mandibular molars of prosimian primates and closely related non-primates first published in [13]. GFPS is known to produce approximately uniformly distributed points on surfaces, with respect to the canonical surface volume measure. In our numerical experiments in this subsection, we choose the first point in GFPS to be the same as the first point obtained by Gaussian process landmarking, to eliminate the effects of random initialization. We further compare with two additional

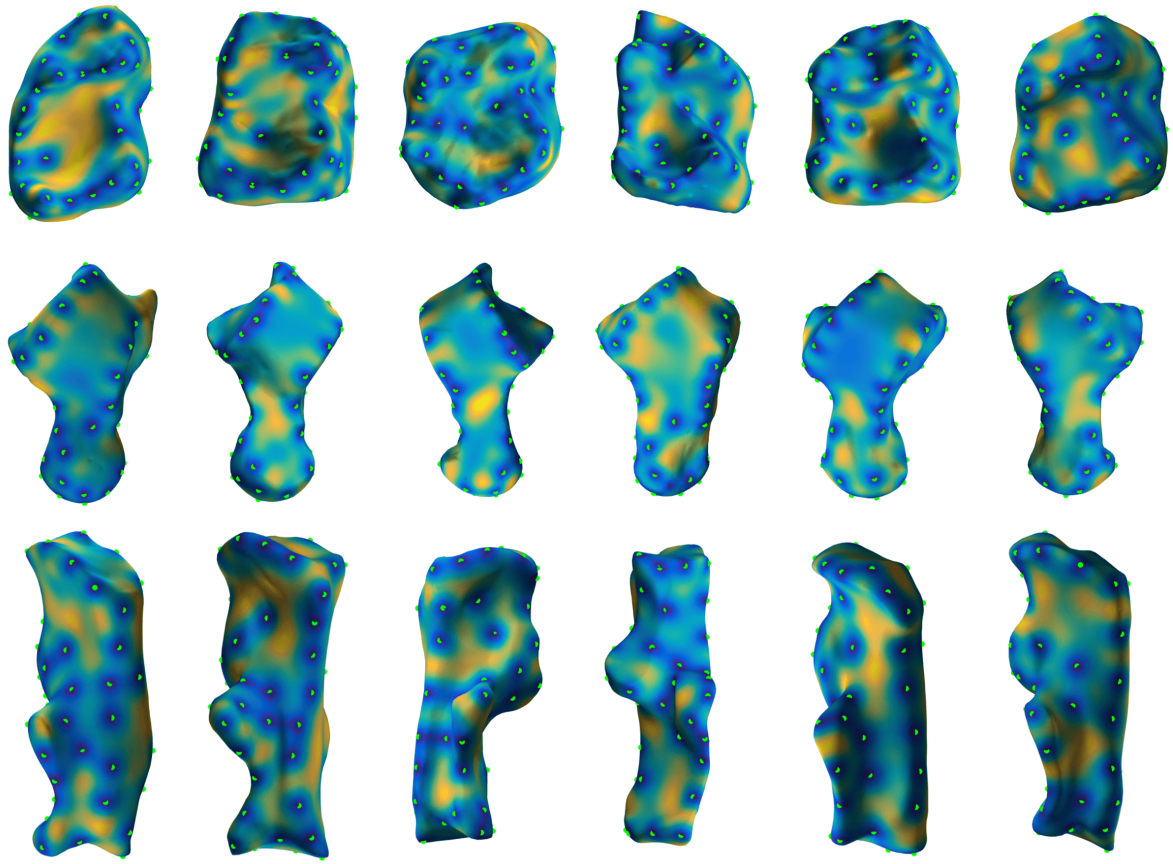

Figure 4: Gaussian process landmarks on several different types of anatomical surfaces. All triangular meshes are acquired from μCT scans. **Top Row:** Six second mandibular molars of prosimian primates and non-primate close relatives, from a dataset of 116 molars first published in [13]. **Middle Row:** Six astragalus bones of tarsiers from a dataset of 40 astragali first published in [14]. **Bottom Row:** Six calcaneus bones of tarsiers from a dataset of 40 calcanei first published in [14]. On all three types of anatomical surfaces, Gaussian process landmarks tend to play the roles of both landmarks and semilandmarks (c.f. [105, §2]): the curvature-reweighted kernel promotes automatically selecting sharp peaks or saddle points on the anatomical surface; after most of the prominent geometric features—normally recognized as Type 2 landmarks [105]—are captured, the uncertainty-based criterion encourages the identification of semilandmark-type points along ridges and grooves.

GFPS variants that use a predetermined set of feature points for initialization: Curv-FPS first selects points of critical curvature and then applies geodesic farthest point sampling; we used the setup of [13] for the detection of isolated points of locally critical Gaussian- and mean-curvatures. GP-FPS takes the $k = 5$ first Gaussian process landmark points and then proceeds with geodesic farthest point sampling. As a baseline, we also calculate the median observer-to-automatic landmark distances for completely randomly picked vertices on the triangular meshes in this dataset. The results are presented in Figure 5, in which each curve is obtained

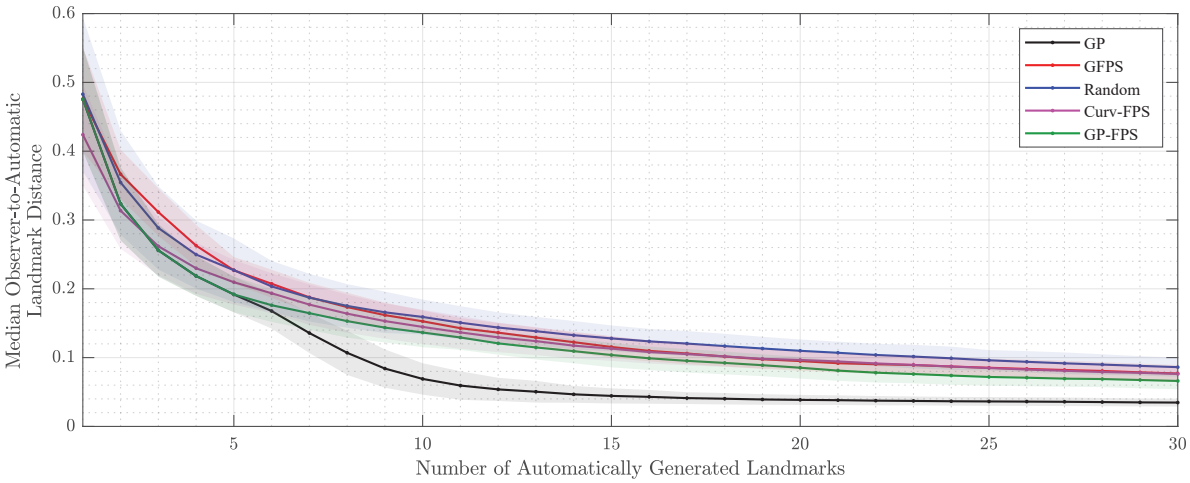

Figure 5: Median geodesic distance from an observer landmark to its nearest automatically generated landmark, with respect to different numbers of automatic landmarks, on a collection of 116 second mandibular molars of prosimian primates and closely related non-primates (see [13] for a more detailed description of this dataset and the observer landmark acquisition). Each point on any of the three solid curves is obtained by averaging the 116 median observer-to-automatic landmark distances over the entire dataset, and the transparent bands represent confidence intervals of one standard deviation. The types of automatic landmarks are Gaussian process landmarks (GP), geodesic farthest point sampling landmarks (GFPS), random landmarks uniformly selected from the vertices of each triangular mesh, and geodesic farthest point sampling landmarks initialized by either critical points of curvature (Curv-FPS) or the first few Gaussian process landmarks (GP-FPS). The random landmarks are only used as a baseline in this experiment.

by averaging individual curves over the entire shape collection; confidence intervals of one standard deviation are also plotted in shades of transparency. Figure 5 suggests that Gaussian process landmarks consistently outperforms GFPS landmarks and the random baseline in terms of coverage of observer landmarks. Figure 14, provided in the supplementary materials, includes additional quantiles of the geodesic distances between observer landmarks and their closest landmarks obtained with each of the automatic methods. The results for the first metatarsals and radii datasets are similar; see Figure 14 and Figure 13 in the supplementary materials.

The comparisons in Figure 5, Figure 12, Figure 13, and Figure 14 are all conducted between the Gaussian process landmarks and the expert observer landmarks manually prepared in [13]. These landmarks were treated as "golden standards" accepted by domain experts, as show in [13] and subsequent domain applications in geometric morphometrics [2, 31]. It is interesting to investigate into the reproducibility and stability of these observer landmarks — will Gaussian process landmarks be close to landmarks manually selected by an independent evolutionary anthropologist following the same landmarking protocol (e.g., pre-fixed ordering and/or number of landmarks of each type in the typology summarized in Subsection 2.1, for ensuring the reproducibility of landmark-based geometric morphometrics; see e.g., [11, 105])? To this end, we collected six additional sets of observer landmarks on the molar teeth in

Figure 1, picked by independent human experts following the landmarking protocol detailed in the supplementary materials of [13], and compared them with the observer landmarks in [13] and the Gaussian process landmarks. The result is shown in Figure 6. Additional views of this example are provided in the supplementary materials; see Figure 15. It is clear from these figures that observer landmarks following the same protocol can be robustly reproduced, especially for geometric features and semilandmarks. A thorough investigation of the stability and reproducibility of observer landmarks is beyond the scope of this paper, but the experimental results in Figure 6 and Figure 15 suggest that Gaussian process landmarks resemble independent observer landmarks under the same protocol as in [13].

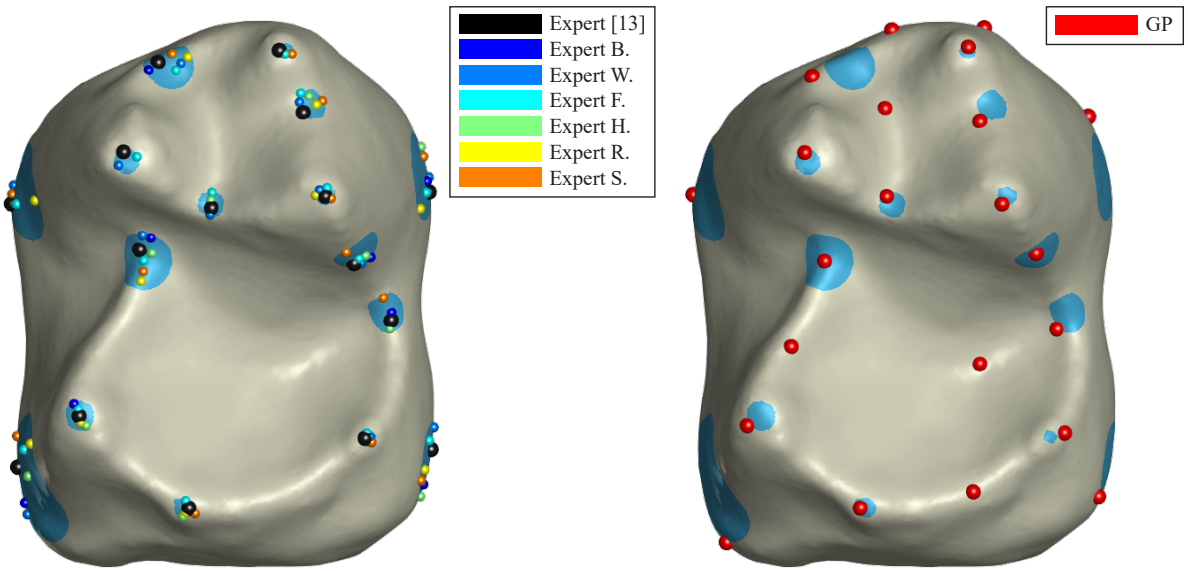

Figure 6: (a) Observer landmarks placed by multiple human experts on the fossil molar of *Teilhardina* in Figure 1. Landmarks obtained in [13], shown in black color, were identified manually by an evolutionary anthropologist and used throughout this paper as ground truth. Six additional sets of landmarks are shown in multiple colors; these landmarks were provided by experts guided to manually place landmarks following the same landmarking scheme, so as to capture the same morphological features; see Figure 16 for a visualization of the landmarks provided by each individual expert. Cyan-shaded areas depict the smallest geodesic disk enclosing each group of observer landmarks. (b) Gaussian process landmarks computed with Algorithm 2.1 are shown in red along the cyan disks representing the distribution of human expert landmarks. This further exemplifies the potential of the first few Gaussian process landmarks in identifying biologically-informative feature points.

4.2. Surface Registration via Matching Gaussian Process Landmarks. We demonstrate in this subsection the benefit of using Gaussian process landmarks for establishing correspondences between pairs of surfaces. In particular, although Gaussian process landmarks are obtained independently on each surface, they turn out to encode geometrically significant features — shared more often among similar or related shapes — that are recognized as "operational homologous" [14] loci by comparative biologists. In order to obtain biologi-

cally meaningful correspondences, we use reduced bandwidth parameter in the discrete heat kernel when applying Algorithm 2.1 that is between 3/5 of the bandwidth parameter used to generate the illustrative Figure 1. Statistical analysis also suggest that these correspondence maps reach comparable explanatory power to observer landmarks placed by human experts, in terms of induced shape distances. The comparison between the morphospaces characterized by these shape distances is deferred to Appendix C.

4.2.1. Experimental Setup and Methodology.

Bounded Distortion Gaussian Process Landmark Matching. Let S_1, S_2 be two-dimensional disk-type surfaces (conformally equivalent to planar disks by the Uniformization Theorem; see e.g., [2, 55, 56]) and let $\xi_1^{(1)}, \ldots, \xi_{L_1}^{(1)} \in S_1$ and $\xi_1^{(2)}, \ldots, \xi_{L_2}^{(2)} \in S_2$ be two sets of Gaussian process landmarks computed using Algorithm 2.1 on S_1 and S_2, respectively. Note that the algorithm we present in this subsection works equally well for $L_1 \neq L_2$, though we choose $L_1 = L_2 = 40$ throughout this paper to simplify the discussion. Adopting the approach suggested in [57] for feature-based image matching, we devise the following two-step approach for establishing geometrically-consistent matchings between the two sets of Gaussian process landmarks:

(1) *Parametrization*: For each surface S_j $(j = 1, 2)$ we follow [91, 52] to compute an as-isometric-as-possible (AIAP) two-dimensional parametrization, which is a diffeomorphism $\Phi_j : S_j \to \Omega_j \subset \mathbb{R}^2$ from S_j to a connected planar domain Ω_j minimizing the (discretization of the) *isometric distortion energy*

$$F(\phi) := \int_{S_j} \left(|\nabla \phi(x)|^2 + \left|\nabla \phi^{-1}(x)\right|^2 \right) \mathrm{dvol}_M(x). \tag{4.1}$$

Each landmark $\xi_\ell^{(j)}$ is mapped to a unique corresponding point $\zeta_\ell^{(j)} = \Phi_j\left(\xi_\ell^{(j)}\right) \in \Omega_j \subset \mathbb{R}^2$, where

$$\Phi_j := \operatorname*{argmin}_{\phi: S_j \to \mathbb{R}^2} F(\phi).$$

Since both "left-handed" and "right-handed" shapes exist in our dataset (as illustrated in Figure 4), we attempt two parametrizations — with or without reflection — for each surface, and only keep the orientation that leads to a smaller Procrustes score (see (4.3) below). A byproduct of this implementation detail is a consistent re-arrangement of the surface orientations.

(2) *Bounded Distortion Matching*: Following [57], we search within the set of planar diffeomorphisms between Ω_1 and Ω_2 with *conformal distortion* [54] bounded by a pre-fixed constant $K \geq 1$. This algorithm strives to find a maximal subset of geometrically-consistent correspondences within an initial set of candidate matches. In the extreme case of $K = 1$, the search is constrained within the set of strictly angle-preserving (conformal) maps for the continuous isometric distortion energy and the set of planar rigid transformations for the discretized isometric distortion energy; we select $K = 1.5$ in this experiment to slightly enlarge the search space of candidate maps.

As input to this matching algorithm, for each Gaussian process landmark $\zeta_\ell^{(1)}$ on S_1, we choose $T \geq 2$ Gaussian process landmarks $\zeta_{\ell\to 1}^{(2)}, \ldots, \zeta_{\ell\to T}^{(2)}$ as initial putative candidate matches from the Gaussian process landmarks $\left\{\zeta_j^{(2)} \mid 1 \leq j \leq L_2\right\}$ on S_2; the algorithm then searches for the bounded distortion map $\Psi : \Omega_1 \to \Omega_2$ that approximately minimizes the mismatch count

$$\sum_{\ell=1}^{L_1} \sum_{k=1}^{T} \left\| \Psi\left(\zeta_\ell^{(1)}\right) - \zeta_{\ell\to k}^{(2)} \right\|^0, \tag{4.2}$$

where, following the notations of [57], $\|\cdot\|^0$ denotes the mixed $(2,0)$-norm

$$\left\| \Psi\left(\zeta_\ell^{(1)}\right) - \zeta_{\ell\to k}^{(2)} \right\|^0 = \begin{cases} 1 & \text{if } \Psi\left(\zeta_\ell^{(1)}\right) \neq \zeta_{\ell\to k}^{(2)}, \\ 0 & \text{otherwise.} \end{cases}$$

In practice, we follow [57] and approximate the mixed $(2,0)$-norm with a mixed $(2,p)$-norm with p decreasing during the optimization. The initial candidate matches in our experiments are generated by comparing the *Wave Kernel Signature* (WKS) [5] of the Gaussian process landmarks — the T Gaussian process landmarks on S_2 with most similar WKS's (measured in Euclidean distances) with that of each $\zeta_\ell^{(1)} \in S_1$ are selected as $\left\{\zeta_{\ell\to t}^{(2)} \mid 1 \leq t \leq T\right\}$. Ideally, "incorrect" initial matches, which potentially lead to large conformal distortions, will be filtered out by minimizing (4.2) under the bounded conformal distortion constraint. Our implementation uses $T = 2$; see Figure 7(b) for an illustrative example.

The pairwise registration algorithm outputs a subset of $1 \leq L \leq \min\{L_1, L_2\}$ one-to-one correspondences $\tilde{\zeta}_\ell^{(1)} \leftrightarrow \tilde{\zeta}_\ell^{(2)}$, $\ell = 1, \ldots, L$, from the initial $L_1 T$ candidate Gaussian process landmark matches; the number of matched landmarks, L, is automatically determined by the final bounded distortion map and could be vary between different pairs of surfaces. See Figure 7(c)(d) for an illustrative example. A final step in this algorithm pipeline "interpolates" the L pairs of matched Gaussian process landmarks to obtain a diffeomorphism between the surfaces S_1 and S_2; see Figure 7(d) for an illustrative example. For this purpose, we use the technique developed in [91, 52] which computes a map $\tilde{\Psi} : \Omega_1 \to \Omega_2$ that minimizes the isometric distortion energy (4.1) subject to the L linear equality constraints $\tilde{\Psi}\left(\tilde{\zeta}_\ell^{(1)}\right) = \tilde{\zeta}_\ell^{(2)}$ $(1 \leq \ell \leq L)$ representing the "sparse" correspondences between Gaussian process landmarks. The composition $f_{12} = {\Phi_2}^{-1} \circ \tilde{\Psi} \circ \Phi_1$ finally produces the desired map $f : S_1 \to S_2$ that approximately minimizes the mismatch count (4.2). This last interpolation step is indispensable for the purpose of interpretability (producing full-surface registrations for visual comparisons) as well as evaluation (inducing shape Procrustes distances for ordination, see Appendix C). Figure 7 outlines the complete workflow of landmark matching and surface registration for a pair of molars surfaces from [13]; a few examples are shown in Figure 8.

Evaluation Metrics and the Baseline. We demonstrate the rich taxonomic information captured by the maps output from the proposed bounded distortion Gaussian process landmark matching (GP-BD) algorithm, by comparing the shape *Procrustes distances* induced by maps

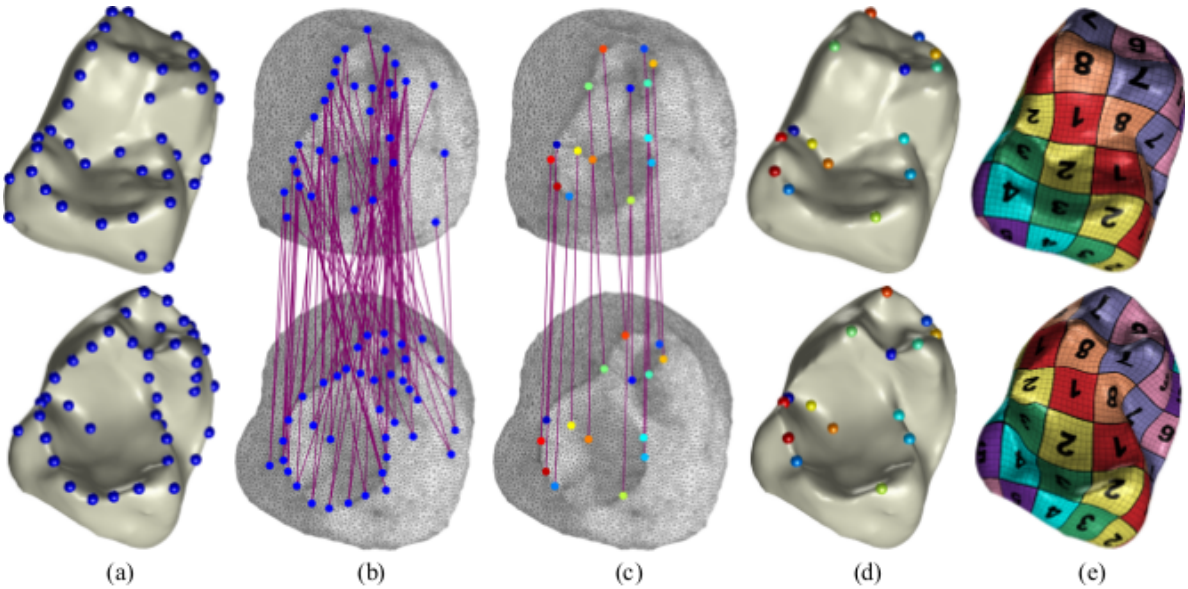

Figure 7: The outline of our approach for landmark matching and surface registration. (a) Gaussian process landmarks computed using Algorithm 2.1 on two surfaces; (b) planar parameterizations of the surfaces that minimize (4.1), overlaid with putative landmark matches (indicated by purple lines); (c) Bounded-Distortion correspondences: a subset of geometrically-consistent matches computed by approximately minimizing (4.2); (d) Pairs of corresponding Gaussian process landmarks found by bounded distortion matching, illustrated by spheres with matching colors; (e) Texture-map visualization of a correspondence map interpolating the landmark correspondences depicted in (d).

computed from GP-BD with those obtained by the same feature matching methodology but alternative choices of landmarking schemes. For any correspondence map $f : S_1 \to S_2$, we define the *Procrustes distance induced by* f as

$$d_{\mathrm{P}}(f) = \left(\min_{R \in \mathbb{E}_3} \int_{S_1} \|f(x) - R(x)\|^2 \, \mathrm{dvol}_{S_1}(x) \right)^{\frac{1}{2}}, \tag{4.3}$$

where $\mathbb{E}_3$ stands for the rigid motion group in $\mathbb{R}^3$. The distance $d_{\mathrm{P}}(f)$ measures the spatial registration error induced by the map f between the two surfaces. To ensure that the Procrustes distances induced by the various methods listed in Table 1 is comparable, we normalize each surface to have unit surface area. When the map f is the continuous Procrustes map produced from the algorithm presented in [2], d_{P} gives exactly the *continuous Procrustes distance* between S_1 and S_2.

The alternatives considered in this section include two other different types of Gaussian process landmarks ($\mathrm{GP}_{\mathrm{nW}}$-BD and $\mathrm{GP}_{\mathrm{Euc}}$-BD) with alternative kernel functions (different from (2.7)), two different strategies of utilizing the "ground truth" user-placed observer landmarks (GT-BD and GT^2-BD) as "oracles" representing the "ideal" landmarks for feature matching, as well as the baseline *continuous Procrustes maps* (CPM) reported in [2, 13]; Table 1 provides a summary of these variants. Specifically, $\mathrm{GP}_{\mathrm{Euc}}$-BD and $\mathrm{GP}_{\mathrm{nW}}$-BD are based

Abbreviation	Description of the pairwise surface registration method
GP-BD	BD-filtering for GP landmarks computed with Algorithm 2.1
GP_{Euc}-BD	BD-filtering for GP landmarks with the standard Euclidean heat kernel (2.5)
GP_{nW}-BD	BD-filtering for GP landmarks with non-weighted kernel ($w \equiv 1$ in (2.7))
GT-BD	BD-filtering for GT landmarks
GT^2-BD	BD-filtering for GT landmarks, also using the ground truth correspondences as candidate initial matches
CPM	Continuous Procrustes Maps (CPM) computed using the method of [13, 2]

Table 1: The pairwise surface registration methods compared in Subsection 4.2 and Appendix C. BD-filtering stands for *bounded distortion maps based feature matching filtering*; GP stands for "Gaussian process"; GT stands for "Ground Truth," i.e., those landmarks placed by experienced comparative biologists.

on Gaussian processes with the standard squared exponential kernel (2.5) and the trivially weighted kernel (setting $w \equiv 1$ in (2.11)), respectively. GT-BD simply replaces the GP landmarks with "ground truth" (GT) landmarks, which then follows through with candidate selection via WKS and bounded distortion maps based feature matching filtering (BD-filtering); GT^2-BD also takes ground truth landmarks as inputs, but skips the candidate selection by setting the ground truth correspondences between observer landmarks as candidate matches, though BD-filtering still applies and prunes out potential "geometrically incompatible" correspondences leading to large conformal distortions. A detailed description of observer landmark acquisition can be found in [13].

The shape distances will be compared in several different ways. We will first compare cumulative distributions of the pairwise shape distance values, followed by two statistical tests addressing (i) the correlation between each automatic shape distance and the observer-determined landmarks Procrustes distance (ODLP), and (ii) the capability of each shape distance at distinguishing taxonomic groups. A qualitative but more intuitive comparison of the *morphospaces* [63] characterized by these shape distances will be deferred to supplementary materials Appendix C, in the form of ordination plots (two dimensional embedding of the shape distance matrices as visual representations for shape variances across species groups). The discussion in Appendix C is however oriented slightly more towards readers with some background in comparative biology.

4.2.2. Comparison Results. We compared the pairwise surface registration methods listed in Table 1 on three different dataset of anatomical surfaces: (i) 116 second mandibular molars of prosimian primates and closely related non-primates; (ii) 57 proximal first metatarsals of prosimian primates, New and Old World monkeys; (iii) 45 distal radii of apes and humans. Detailed descriptions of all 3 datasets can be found in [13]. We computed correspondences between each pair of surfaces within this dataset, totaling in over $13,000$ correspondence maps for dataset (i), over $1,500$ for dataset (ii), and nearly $1,000$ for dataset (iii). Figure 8 shows example pairs of surfaces, visualizing correspondence maps induced by GP-BD correspondences

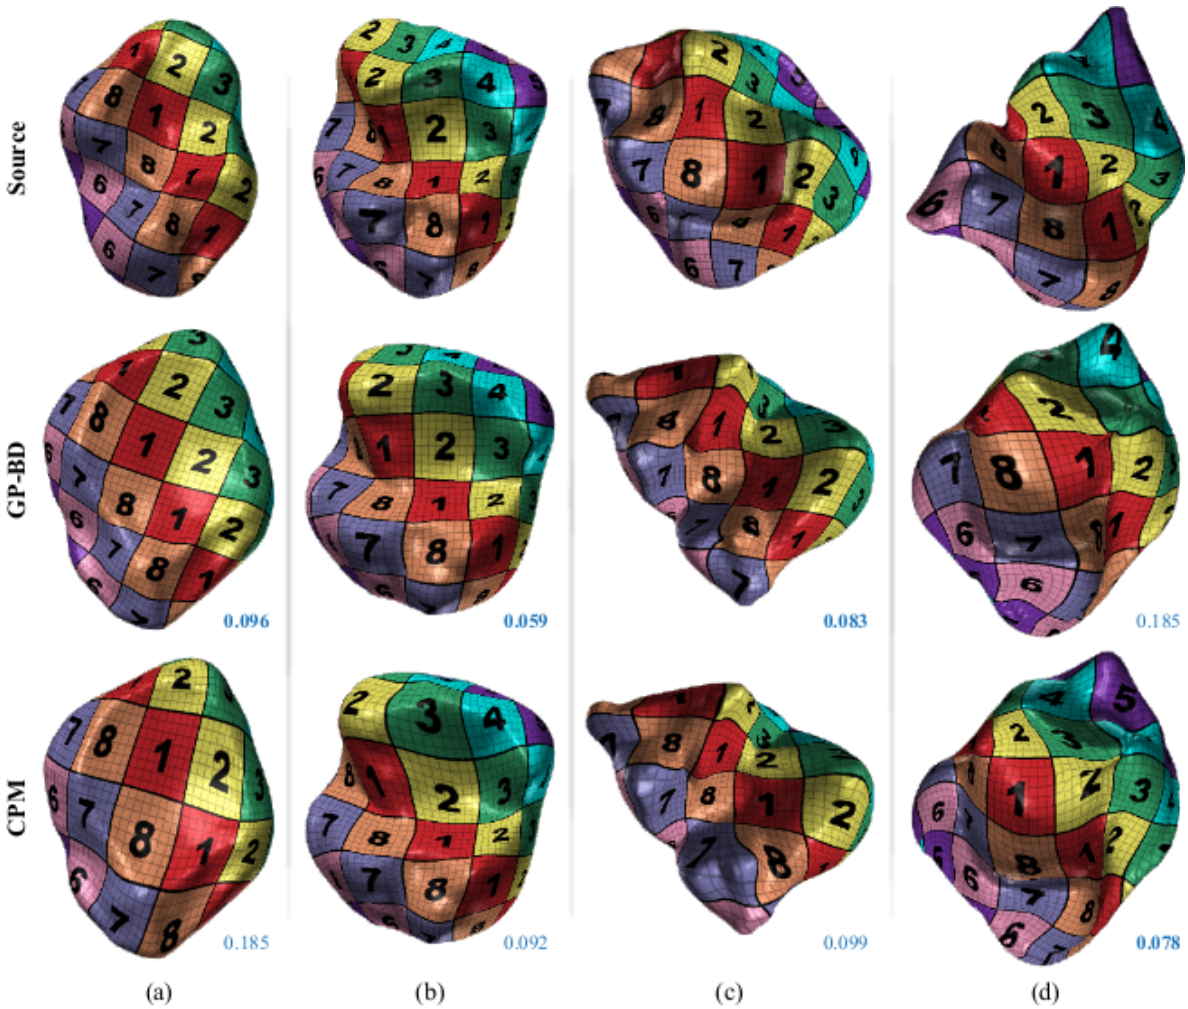

Figure 8: Correspondence maps and induced Procrustes distances: texture overlaid on the Source surface (top row) is mapped by $f : S_1 \to S_2$ so as to visualize correspondence. The second row (GP-BD) shows maps induced by correspondences established between Gaussian process landmarks computed with Algorithm 2.1. The bottom row (CPM) compares the baseline continuous procrustes maps. GP-BD outperforms CPM in examples (a)-(c), wherein the map it induces establishes a better correspondence between morphological features of the surfaces; example (d) shows a case in which GP-BD is outperformed by CPM. The inset values are the the Procrustes distance d_{P} associated with each map.

as well as the baseline CPM maps. These examples illustrate typical differences between maps computed from GP-BD and CPM along with their Procrustes distances; GP-BD maps often offer an improvement over CPM in terms of both visual quality and their ability to relate biologically meaningful and operationally homologous regions.

Comparing Cumulative Distributions of Distance Values. We first provide a crude comparison across the various methods listed in Table 1. Figure 9 plots the cumulative distributions of

Procrustes distances induced by each type of pairwise registration method in Table 1, which serves as a direct comparison of the proportions of pairwise correspondences for which the Procrustes distance d_{P} (4.3) is below a given threshold. The figure further includes a curve for the Observer-Determined Landmarks Procrustes (ODLP) distances, computed using the standard Procrustes analysis between sets of corresponding landmarks manually placed by human experts [13]. Recall from Subsection 4.2.1 that the surface areas are normalized so as to ensure that Procrustes distances induced by different shape correspondences are comparable.

Noticeably, the cumulative distribution of GP-BD distances most closely resemble that of GT-BD, obtained by replacing the Gaussian process landmarks with the ground-truth observer landmarks (but otherwise using exactly the same algorithm involving WKS and BD-filtering). Also, GP-BD outperforms $\mathrm{GP}_{\mathrm{nW}}$-BD (Gaussian process with a trivially weighted kernel) and $\mathrm{GP}_{\mathrm{Euc}}$-BD (Gaussian process with the standard Euclidean heat kernel). This comparison suggests that the Gaussian process landmarks computed with Algorithm 2.1 provide a good proxy for geometrically significant features needed to determine meaningful correspondences between surfaces.

Expectedly, GP-BD falls short compared to distances that rely on both the ground-truth observer landmarks and their true correspondences (i.e., GT^2-BD and ODLP). Comparing GP-BD to the baseline continuous Procrustes maps (CPM) in terms of distance distributions is equivocal as the two curves in Figure 9 cross each other when the Procrustes distance threshold is about 0.15. Nonetheless, comparing GP-BD to CPM indicates that the former produces more pairwise correspondences with shape distances less than 0.1 (at which the vertical gap between the two curves reaches its maximum). This, along with the correlation reported in [31] between smaller continuous Procrustes (cP) distances and better morphometric interpretability of the associated maps, implies that Gaussian process landmarks potentially lead to more stable and interpretable comparative biological analysis if combined with other globally transitive geometric morphometric methods.

Statistical Tests Comparing the Distance Matrices. We now provide a more thorough comparison, using nonparametric statistical tests, for the shape distance matrices induced by the various methods listed in Table 1. Treating the OPLD (Observer-Determined Landmarks Procrustes distance) matrix as a proxy for the "ground truth" accepted among geometric morphometricians (c.f. [13]), we first compare the correlation between each automatically computed shape distance matrix with ODLP using the *Mantel correlation test* [60]; the explanatory power of the shape distances (in terms of the ability to distinguish taxonomic groups) are then compared using a variant of the multivariate analysis of variance (MANOVA) for distance matrices.

Due to the dependence of entries in a distance matrix (e.g., constrained by the triangular inequality), assessments of correlation between distance matrices, or between a distance matrix and a (continuous or categorical) variable, often involves repeatedly permuting the rows and columns of the distance matrix (see e.g., [76, 70, 39] and the references therein). Table 2 demonstrates the results of Mantel correlation analysis [60] between the distance matrices computed in this paper against the ODLP distance matrix. In all mantel tests in this paper, we set the number of permutations to 9999 since it is recommended in [4] that at least 5000 permutations should be done for tests with an α-level of 0.01. GT-BD and GT^2-BD correlate best with ODLP, which is as expected since they both rely directly on the same set of observer-

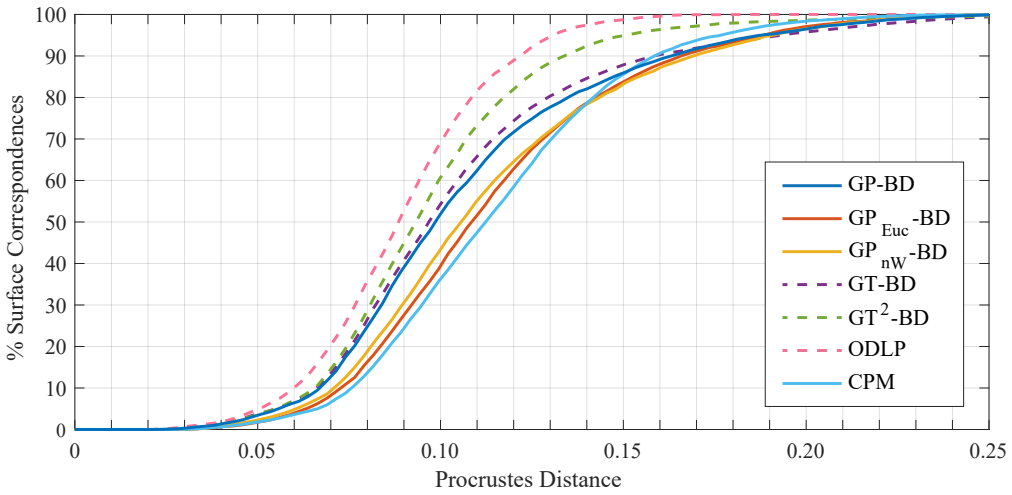

Figure 9: Comparing GP-BD to the alternative methods listed in Table 1. Each curve represents the proportion of correspondences, among a collection of 116 surfaces, for which the Procrustes distance d_{P} falls below a given threshold. The cumulative distribution of GP-BD distances most closely resembles that of GT-BD which utilizes the ground truth observer landmarks. It falls short of GT^2-BD and the Observer-Determined Landmarks Procrustes (ODLP), both of which rely additionally on the "ground truth" correspondences between observer landmarks. In terms of producing more relatively smaller distances between the shape pairs (which is favorable by the correlation between smaller distances and enhanced morphometric interpretability reported in [31]; also suggested by the cumulative distribution curves of ODLP, GT^2-BD, and GT-BD), GP-BD outperforms the baseline continuous Procrustes maps (CPM), as well as variants of Gaussian process landmarks with alternative kernel constructions, GP$_{\mathrm{nW}}$-BD (Gaussian process landmarks with a non-weighted kernel) and GP$_{\mathrm{Euc}}$-BD (Gaussian process landmarks with a standard Euclidean heat kernel).

determined landmarks used for computing the ODLP distances; their high correlation with ODLP justifies the use of bounded distortion maps for capturing biologically meaningful and corresponding geometric features. GP-BD outperforms CPM, which is consistent with our observation in Figure 9 that a majority (over 80%) of the distance values computed from GP-BD is more similar in distribution to the distances computed with ground truth landmarks. The advantageously higher correlation of GP-BD over GP$_{\mathrm{Euc}}$ and GP$_{\mathrm{nW}}$ illustrates the importance of reweighting in the kernel construction.

In addition, we perform *Permutational Multivariate Analysis of Variance* (PERMANOVA) [4] for the distance matrices computed in this paper, with the taxonomic groups shown in the ordination plots (Figure 17, see detailed explanations in Appendix C) as the treatment effects for the molar dataset; for the first metatarsal and radius datasets, we use family and species groups as treatment effects, respectively. The results are presented in Table 3. The purpose of this test is to quantitatively compare the power of differentiating taxonomic groups for these distance matrices. The pseudo F-statistics in PERMANOVA is a properly normalized ratio between among-group and within-group sum of squared distances. The statistical significance is then calculated from the fraction—among sufficiently many shuffles of the rows and columns

	GT	GT^2	CPM	GP	GP_{Euc}	GP_{nW}
Molars	**0.7042**	**0.7563**	0.6647	0.6870	0.6135	0.6257
First Metatarsals	**0.7371**	**0.8326**	0.4887	0.7318	0.6607	0.7117
Radii	**0.3273**	**0.4909**	0.1775	0.3231	0.2510	0.2746

Table 2: Correlation coefficients output from Mantel correlation analysis for the various distances computed in this paper versus ODLP distances, for three different datasets. See Table 1 for the list of abbreviations; note that we omitted "-BD" for the sake of space. The relatively high correlations of GT-BD and GT^2-BD versus ODLP is not surprising due to their direct dependence on the observer-determined landmarks; the first 2 columns can thus be viewed as performance upper bounds when "oracle landmark correspondences" are provided. For all three datasets, the distance matrix computed from GP-BD correlates better with ODLP than CPM, as well as almost all of its variants, which is consistent with the CDF plot Figure 9 and the ordination plots in Figure 17 and speaks of the advantages of matching anatomical surfaces with Gaussian process landmarks. The P-values of all results are < 0.01.

of the distance matrix—of the permutation instances that produce a higher pseudo F-ratio. While all distance matrices demonstrate statistically significant pseudo F-ratios in this test, the GT-BD distance matrix leads the board of pseudo F-ratios, indicating its superior ability of separating taxonomic groups; the GP-BD distance matrix and its two variants consistently outperform CPM in terms of pseudo F-ratios, verifying again the improved quality of pairwise anatomical surface registrations as well as the induced shape Procrustes distances.

	ODLP	GT	GT^2	CPM	GP	GP_{Euc}	GP_{nW}
Molars	12.26	**16.70**	13.90	9.42	**16.26**	11.15	12.99
First Metatarsals	20.30	**29.87**	18.12	6.38	**32.85**	19.17	21.82
Radii	**9.19**	**11.89**	9.03	5.59	9.09	6.35	8.04

Table 3: Pseudo F-ratios output from PERMANOVA for various distance matrices computed in this paper for three different datasets. See Table 1 for the list of abbreviations; note that we omitted "-BD" for the sake of space. The number of groups for the molar dataset is 30, equaling to the number of polygonal regions in the ordination plot in Figure 17; the numbers of groups for first metatarsals and radius are 8 (families) and 4 (genuses), respectively. The first column refers to observer-determined landmarks Procrustes (ODLP) distances calculated in [13] as a baseline; the first 3 columns can thus be viewed as performance upper bounds when "oracle landmark correspondences" are provided. For all three datasets, the pseudo-F ratio of the GP-BD distance matrix is better or at least comparable to the oracles. All distance matrices are statistically significant in terms of their powers of separating species groups: the P-values of all results are < 0.01.

5. Discussions. In this paper, we apply the uncertainty-based landmark generating algorithm proposed in [30] to three-dimensional geometric morphometrics. The algorithmically produced landmarks resemble biologists' landmarks selected with expert knowledge, providing adequate coverage for both geometric features and semantically "uncertain" regions on the anatomical surfaces of practical interest. We tested the applicability of this landmarking

procedure for various tasks on real datasets.

The Gaussian process landmarking algorithm presented in this paper takes one anatomical surface as input at one time, which is not exactly consistent with the methodology of geometric morphometricians. In fact, these biologists do not simply place landmark on individual surface; rather, they are trained to take into consideration an entire collection of shapes, so as to consistently place landmarks that are in joint correspondence. In standard practice, landmarking a new anatomical surface typically involve repeated comparisons with all the other surfaces in the collection, and the already placed landmarks on a surface are still subject to change upon future knowledge acquired from landmarking more surfaces. It would be highly interesting to adapt our algorithm to accommodate for this type of group-wise comparison strategies as well.

Though geometric approaches in general lack the ability to recognize Type 1 landmarks (recall Bookstein's typology reviewed in Subsection 2.1), we think that the similarities between Gaussian process landmarks and user selected Type 2 and Type 3 landmarks call into question user-based landmarks. In particular, we found the ordination plots to be as successful, and in some ways superior, to the user-based landmarks in reflecting previous ideas about shape affinities. While F-ratios were highest in the user-determined sample, we note that it is highly likely that, for any underdetermined landmarks (e.g., the Type 1 landmarks, Type 3 landmarks, or when trying to place Type 2 landmarks on eroded features like blunted cusp tips) biologists will tend to minimize variance within species. In other words, if there is not enough geometry to allow consistent placement of a point, users are likely to unconsciously choose a positioning that visually maximizes similarity to other members of the species. Given the likely bias towards minimizing within group error by biologists during landmarking, it is quite remarkable that the Gaussian process approach comes so close to the "ground truth", and exceeds by so much the other automated methods in terms of F-ratio. In user-placed landmarks, there is also the question of how many and which landmarks were chosen, not to mention whether different users are as accurate in placing the same landmarks. Finally, there is the limitation that any traditional landmark used must be present in every specimen of the sample. Given the challenges with the traditional user based approach, and the demonstrated ability of Gaussian process to emulate qualitatively and quantitatively Type 2 and Type 3 landmarks, we believe the Gaussian process landmarking algorithm has great potential to be an alternative, automated approach for such landmarks.

Software. `MATLAB` code for the surface registration algorithm is publicly available at https://github.com/shaharkov/GPLmkBDMatch.

Acknowledgments. The authors thank Peng Chen, Chen-Yun Lin, Shaobo Han, Sayan Mukherjee, Rob Ravier, and Shan Shan for inspirational discussions; Ethan Fulwood, Bernadette Perchalski, Julia Winchester, Arianna Harrington, Robert Ravier, and Shan Shan for assistance with collecting observer landmarks; the anonymous reviewers for many constructive feedback.

Appendix A. Proof of Theorem 3.1.

Proof of Theorem 3.1. Recall from [89] (or [22, Lemma 8], [90, Lemma B.10]) that

$$\begin{aligned}(\text{A.1}) \qquad & \frac{1}{(2\pi\epsilon)^{\frac{d}{2}}}\int_M \exp\left(-\frac{\|x-y\|^2}{2\epsilon}\right) f(y)\ \mathrm{dVol}(y) \\ & \qquad = f(x) + \frac{\epsilon}{2}\left[E(x) f(x) + \Delta f(x)\right] + O\left(\epsilon^{\frac{3}{2}}\right) \qquad \text{as } \epsilon \to 0\end{aligned}$$

where E is a scalar function of the curvature of M at $x \in M$. Thus as $\epsilon \to 0$

$$\begin{aligned}(\text{A.2}) \qquad & \exp\left[-V(z)\right]\int_M \exp\left(-\frac{\|z-y\|^2}{2\epsilon}\right) f(y)\ \mathrm{dVol}(y) \\ & \qquad = (2\pi\epsilon)^{\frac{d}{2}} \exp\left[-V(z)\right]\left\{f(z) + \frac{\epsilon}{2}\left[E(z) f(z) + \Delta f(z)\right] + O\left(\epsilon^{\frac{3}{2}}\right)\right\}.\end{aligned}$$

Replacing the function f in (A.1) with the right hand side of (A.2), we have

$$\begin{aligned}&\int_M\int_M \exp\left(-\frac{\|x-z\|^2}{2\epsilon}\right)\exp\left[-V(z)\right]\exp\left(-\frac{\|z-y\|^2}{2\epsilon}\right) f(y)\ \mathrm{dVol}(z)\ \mathrm{dVol}(y) \\ &= (2\pi\epsilon)^d\left\{\exp\left[-V(x)\right]\left\{f(x) + \frac{\epsilon}{2}\left[E(x) f(x) + \Delta f(x)\right] + O\left(\epsilon^{\frac{3}{2}}\right)\right\}\right. \\ &\quad + \frac{\epsilon}{2}E(x)\exp\left[-V(x)\right]\left\{f(x) + \frac{\epsilon}{2}\left[E(x) f(x) + \Delta f(x)\right] + O\left(\epsilon^{\frac{3}{2}}\right)\right\} \\ &\quad \left. + \frac{\epsilon}{2}\Delta\left(\exp\left[-V(x)\right]\left\{f(x) + \frac{\epsilon}{2}\left[E(x) f(x) + \Delta f(x)\right] + O\left(\epsilon^{\frac{3}{2}}\right)\right\}\right) + O\left(\epsilon^{\frac{3}{2}}\right)\right\} \\ &= (2\pi\epsilon)^d\left\{f(x)\exp\left[-V(x)\right] + \frac{\epsilon}{2}\left\{\left[E(x) f(x) + \Delta f(x)\right]\exp\left[-V(x)\right]\right.\right. \\ &\quad \left.\left. + f(x) E(x)\exp\left[-V(x)\right] + \Delta\Big(f(x)\exp\left[-V(x)\right]\Big)\right\} + O\left(\epsilon^{\frac{3}{2}}\right)\right\} \qquad \text{as } \epsilon \to 0.\end{aligned}$$

Plugging the right hand side above into the numerator of the left hand side of (3.4), and replacing the denominator with the same expression as the numerator except for setting $f \equiv 1$

in (3.4), we obtain, as $\epsilon \to 0$,

$$\begin{aligned}
&\frac{\int_M \int_M \exp\left(-\frac{\|x-z\|^2}{2\epsilon}\right) \exp\left[-V(z)\right] \exp\left(-\frac{\|z-y\|^2}{2\epsilon}\right) f(y)\, \mathrm{dVol}(z)\, \mathrm{dVol}(y)}{\int_M \int_M \exp\left(-\frac{\|x-z\|^2}{2\epsilon}\right) \exp\left[-V(z)\right] \exp\left(-\frac{\|z-y\|^2}{2\epsilon}\right) \mathrm{dVol}(z)\, \mathrm{dVol}(y)} \\
&= \left\{ f(x) + \frac{\epsilon}{2}\left[2E(x) f(x) + \Delta f(x) + \exp\left[V(x)\right] \Delta\Big(f(x) \exp\left[-V(x)\right]\Big)\right] \right. \\
&\qquad \left. + O\left(\epsilon^{\frac{3}{2}}\right) \right\} \times \left\{ 1 - \frac{\epsilon}{2}\Big(2E(x) + \exp\left[V(x)\right] \Delta \exp\left[-V(x)\right]\Big) + O\left(\epsilon^{\frac{3}{2}}\right) \right\} \\
&= f(x) + \frac{\epsilon}{2}\Delta f(x) + \frac{\epsilon}{2} \exp\left[V(x)\right] \Delta\left(f(x) \exp\left[-V(x)\right]\right) \\
&\qquad - \frac{\epsilon}{2} f(x) \exp\left[V(x)\right] \Delta \exp\left[-V(x)\right] + O\left(\epsilon^{\frac{3}{2}}\right) \\
&= f(x) + \epsilon\Big(\Delta f(x) - \nabla f(x) \cdot \nabla V(x)\Big) + O\left(\epsilon^{\frac{3}{2}}\right).
\end{aligned}$$

This completes the proof. ∎

Appendix B. Numerical Experiments on Additional Datasets.

In this section we present the results of the surface registration algorithm in Subsection 4.2 on two other anatomical datasets, one consisting of 57 proximal first metatarsals of prosimian primates, New and Old World monkeys (Figure 11), and the other consisting of 45 distal radii of apes and humans (Figure 10). Detailed descriptions of these two additional datasets can be found in [13].

We provide additional figures (Figure 12, Figure 13, Figure 14) illustrating the advantageous coverage properties of the Gaussian process landmarks, in terms of their geodesic distances to observer landmarks. Each figure corresponds to one of the three datasets (molars, first metatarsals, and radii) and contains three subplots — one for each of three levels of quantile (25%, 50%, 75%). See captions of these figures (also that of Figure 5) for more detailed explanations.

We also provide in Figure 15 and Figure 16 additional illustrations of multiple sets of expert landmarks manually placed following the same protocol in [13]. This is supplementary to Figure 6 in the main text. We refer to the captions of the these figures for more detailed information.

Appendix C. Ordination and Comparative Biological Interpretation for Shape Collections.

This section provides a qualitative but visually more straightforward way to compare the shape distance matrices obtained from the various pairwise registration methods in Table 1 on a collection of primate molars. The results presented here suggests that shape distances obtained by matching Gaussian process landmarks can potentially capture more morphometrical information and provide better characterizations of the shape variation across species groups.

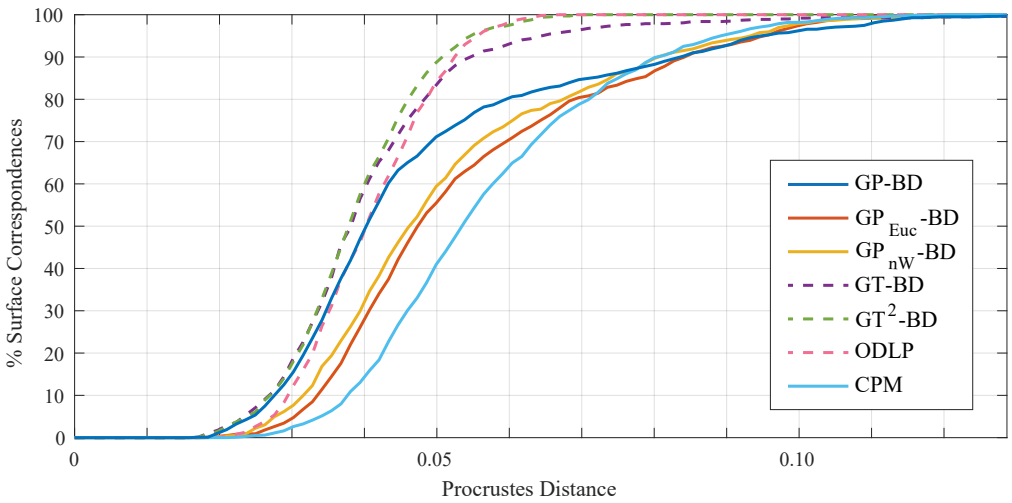

Figure 10: Comparing GP-BD to the alternative methods listed in Table 1 for the dataset of distal radii of apes and humans. The general trend is consistent with our observation in Figure 9: GT-BD outperforms CPM and is comparable to variants of Gaussian process landmarks with alternative kernel constructions, whiling falling short of, but not far from, GT^2-DB and ODLP which rely additionally on the oracle "ground truth" correspondences between observer landmarks.

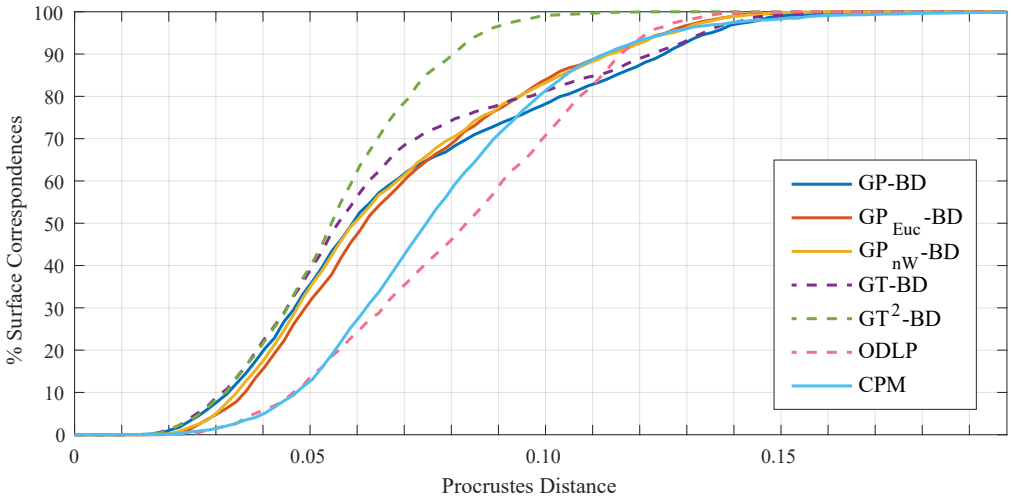

Figure 11: Comparing GP-BD to the alternative methods listed in Table 1 for the dataset of first metatarsals of prosimian primates, New and Old World monkeys. The general trend is consistent with our observation in Figure 9: GT-BD outperforms CPM and variants of Gaussian process landmarks with alternative kernel constructions, falls short of, but not far from, GT^2-DB and ODLP which rely additionally on the oracle "ground truth" correspondences between observer landmarks.

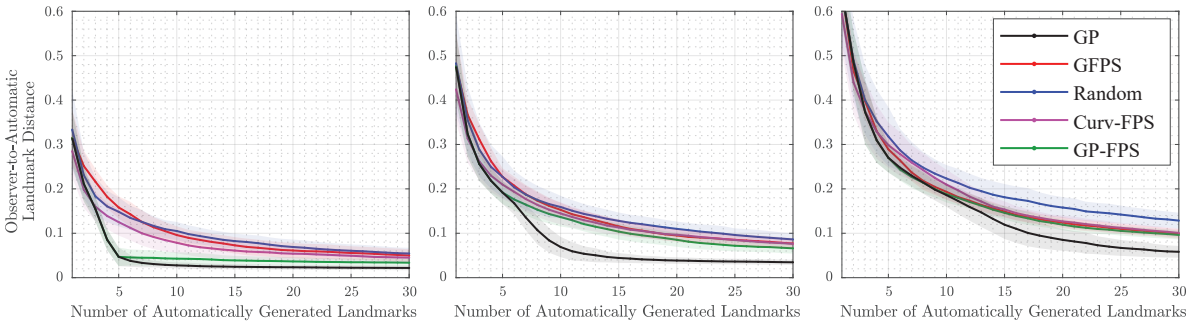

(a) 25% percentile (b) 50% percentile (c) 75% percentile

Figure 12: Geodesic distance from an observer landmark to its nearest automatically generated landmark, with respect to different numbers of automatic landmarks, on a collection of 116 second mandibular molars of prosimian primates and closely related non-primates (see [13] for a more detailed description of this dataset and the observer landmark acquisition). Each point on any of the three solid curves is obtained by averaging the 25%, 50% (median) and 75% percentiles of observer-to-automatic landmark distances computed for each surface in this dataset; the transparent bands represent confidence intervals of one standard deviation. The types of automatic landmarks are Gaussian process landmarks (GP), geodesic farthest point sampling landmarks (GFPS), random landmarks uniformly selected from the vertices of each triangular mesh, and geodesic farthest point sampling landmarks initialized by either critical points of curvature (Curv-FPS) or the first few Gaussian process landmarks (GP-FPS). The random landmarks are only used as a baseline in this experiment.

Comparative biologists, ecologists, and other natural and social scientists frequently rely on *ordination* or *gradient analysis* to summarize and emphasize patterns in multivariate datasets [19, 20, 4]. The equivalent but more familiar term for statisticians and applied mathematicians is *dimension reduction*. Generally speaking, it is assumed or hypothesized that groups that are distantly related or have different ecologies will display some physical differences related to these variables. Nonetheless, it is difficult to know which physical traits reflect phylogenetic and ecological differences *a priori*, or whether certain traits are independent of each other. Ordination methods such as *Principal Components Analysis* (PCA), *Multi-Dimensional Scaling* (MDS), and *Correspondence Analysis* (CA) reveal the nature of those physical differences by summarizing multivariate datasets in lower dimensional graphical form.

In Figure 17 we visualize ordinations induced from four different distance matrices using classical multi-dimensional scaling (MDS). The dataset consists of 116 second mandibular molars of prosimian primates and non-primate close relatives, first published in [13]. The distances include the continuous Procrustes distance [2] and three BD-filtering methods; see Table 1 for the abbreviations. In each ordination MDS plot, points corresponding to the same species group are enclosed in a polygonal, and the color pattern of the polygons indicates shape similarity among generic specimens across species groups, based on visual inspection and traditional comparative analyses. Similar ordination plots have been used for comparing different algorithms designed for automatically quantifying the geometric similarity of anatomical surfaces; see e.g., [14, 97, 31]. This type of plots can also be considered as visualizations

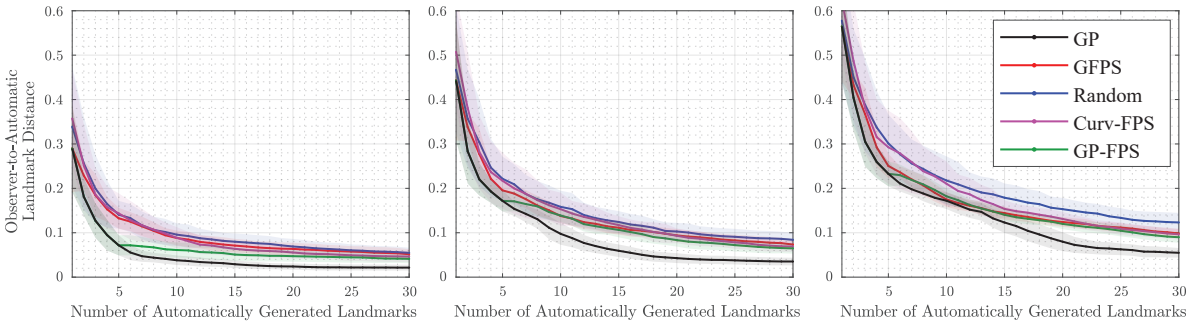

(a) 25% percentile (b) 50% percentile (c) 75% percentile

Figure 13: Geodesic distance from an observer landmark to its nearest automatically generated landmark, with respect to different numbers of automatic landmarks, on a collection of 45 distal radii of apes and humans (see [13] for a more detailed description of this dataset and the observer landmark acquisition). Each point on any of the three solid curves is obtained by averaging the 25%, 50% (median) and 75% percentiles of observer-to-automatic landmark distances computed for each surface in this dataset; the transparent bands represent confidence intervals of one standard deviation. The types of automatic landmarks are Gaussian process landmarks (GP), geodesic farthest point sampling landmarks (GFPS), random landmarks uniformly selected from the vertices of each triangular mesh, and geodesic farthest point sampling landmarks initialized by either critical points of curvature (Curv-FPS) or the first few Gaussian process landmarks (GP-FPS). The random landmarks are only used as a baseline in this experiment.

of *morphospaces* in evolutionary and developmental biology [63].

An important criterion for good ordination is the extent to which specimens belonging to the same species group cluster near each other, and specimens from different species groups are separated from each other. A popular summary statistic is to measure the ratio between within-group distances and among-group distances, which we quantitatively calculate and present in Table 3. Nevertheless, it is already evident from qualitatively inspecting Figure 17 that bounded distortion matching (Figure 17a, Figure 17b, Figure 17d) in general better distinguishes species groups than continuous Procrustes maps (Figure 17c). Considering taxonomic group membership and previous interpretations regarding shape affinities of different taxa, the ordination induced from GP-BD distances (Figure 17a) is visually most appealing. We explain some compelling examples in detail, referencing previous observations by biologists and paleontologists.

(i) Figure 17a has a relatively low within/between-group distance ratio (see Table 3 for more details), with remarkable distinctiveness of the Indriidae, *Tupaia*, and *Cynocephalus* groups from the other groups, which is a faithful reflection of the shape dissimilarity based on visual inspections;

(ii) Among all different ordinations presented in Figure 17, Figure 17a best separates out fossil (the brown-colored groups — Plesiadapoidea, *Plesiolestes*, *Purgatorius*, *Leptacodon*, and *Altanius*) and living non-primates (the yellow-colored groups — *Tupaia* and *Ptilocercus* (Scandentia or "treeshrews") and *Cynocephalus* (Dermoptera)) from strepsirrhine primates and most early fossil primates (light and dark groups of lemurs

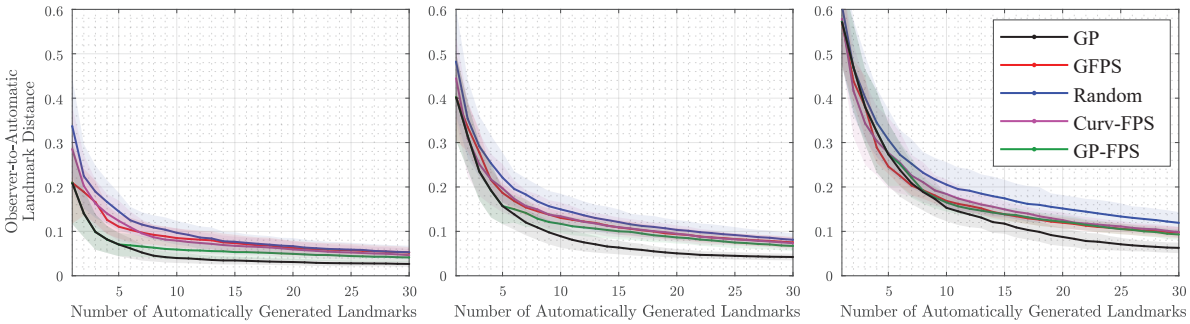

(a) 25% percentile (b) 50% percentile (c) 75% percentile

Figure 14: Geodesic distance from an observer landmark to its nearest automatically generated landmark, with respect to different numbers of automatic landmarks, on a collection of 57 first metatarsals of prosimian primates (see [13] for a more detailed description of this dataset and the observer landmark acquisition). Each point on any of the three solid curves is obtained by averaging the 25%, 50% (median) and 75% percentiles of observer-to-automatic landmark distances computed for each surface in this dataset; the transparent bands represent confidence intervals of one standard deviation. The types of automatic landmarks are Gaussian process landmarks (GP), geodesic farthest point sampling landmarks (GFPS), random landmarks uniformly selected from the vertices of each triangular mesh, and geodesic farthest point sampling landmarks initialized by either critical points of curvature (Curv-FPS) or the first few Gaussian process landmarks (GP-FPS). The random landmarks are only used as a baseline in this experiment.

— *Eulemur*, *Lemur*, *Varecai*, *Cheirogaleus*, *Hapalemur*, *Prolemur*, and Cheirogaleidae; light and dark green groups of Lorisidae — *Loris*, *Galago*, *Arctocebus*, *Nycticebus*, and *Perodicticus*; the red group of the extinct adapiforms and omomyiforms — *Adapis*, *Cantius*, and *Donrussellia*);

(iii) The thought-provoking patterns of overlapping reflected in Figure 17a are in accordance with views established in existing comparative biological literature. For instance, *Tarsius* (a haplorhine primate) and plesiadapoid nonprimates overlap, reflecting the idea by previous authors that they might be united in a common group called "Plesiotarsiiformes" [33, 34]. As another example, *Teilhardina*, *Donrussellia*, and *Cantius* all overlap, reflecting the fact that all of them represent very primitive members of primates: *Donrussellia* was originally thought to be a new species of *Teilhardina* [84]; *Cantius* was originally thought to be an omomyiform like *Teilhardina* instead of an adapiform [88], and [36] observed additional features of the premolar teeth uniquely relating *Cantius torresi* to *Teilhardina*. Though *Cantius* is now reocognised as an adapiform, the wide separation between it and the other sampled adapiform (*Adapis*) also matches previous qualitative discussions and analyses of independent datasets. In particular, the fact that *Adapis* overlaps *Lemur*, *Eulemur*, and *Lepilemur* is reminiscent of arguments by [35] that *Adapis* has special affinities to the living strepsirrhines to the exclusion of other adapiforms, as well as that *Lepilemur* belongs to a group called "Megaladapidae" in part because of the claim that some members of the group have teeth extremely similar to those of *Adapis* ([86]). Anaylsis of other skeletal re-

gions such as the ankle bones [87] also link *Adapis* more closely to living lemurs than to other adapiforms. Finally, the overlap of Eosimias with treeshrews matches the suggestion by [37] that this taxon, otherwise thought of as an anthropoid, is in fact treeshrew-like in its dentition. Again analyses of ankle bones return a similar pattern; see e.g., [104, 15].

REFERENCES

[1] D. C. Adams, F. J. Rohlf, and D. E. Slice, *A Field Comes of Age: Geometric Morphometrics in the 21st Century*, Hystrix, 24 (2013), pp. 7–14.

[2] R. Al-Aifari, I. Daubechies, and Y. Lipman, *Continuous Procrustes Distance Between Two Surfaces*, Communications on Pure and Applied Mathematics, 66 (2013), pp. 934–964, https://doi.org/10.1002/cpa.21444, http://dx.doi.org/10.1002/cpa.21444.

[3] P. Alliez, D. Cohen-Steiner, O. Devillers, B. Lévy, and M. Desbrun, *Anisotropic Polygonal Remeshing*, ACM Trans. Graph., 22 (2003), pp. 485–493, https://doi.org/10.1145/882262.882296, http://doi.acm.org/10.1145/882262.882296.

[4] M. J. Anderson, *A New Method for Non-Parametric Multivariate Analysis of Variance*, Austral Ecology, 26 (2001), pp. 32–46.

[5] M. Aubry, U. Schlickewei, and D. Cremers, *The Wave Kernel Signature: A Quantum Mechanical Approach to Shape Analysis*, in Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on, IEEE, 2011, pp. 1626–1633.

[6] T. F. Banchoff, *Critical Points and Curvature for Embedded Polyhedral Surfaces*, The American Mathematical Monthly, 77 (1970), pp. 475–485.

[7] C. Beard, *Teilhardina. The International Encyclopedia of Primatology. 1–2.*, 2017. DOI: 10.1002/9781119179313.wbprim0444.

[8] M. Belkin and P. Niyogi, *Laplacian Eigenmaps for Dimensionality Reduction and Data Representation*, Neural Comput., 15 (2003), pp. 1373–1396, https://doi.org/10.1162/089976603321780317, http://dx.doi.org/10.1162/089976603321780317.

[9] T. Berry and J. Harlim, *Variable Bandwidth Diffusion Kernels*, Applied and Computational Harmonic Analysis, 40 (2016), pp. 68–96, https://doi.org/10.1016/j.acha.2015.01.001, http://www.sciencedirect.com/science/article/pii/S1063520315000020.

[10] P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova, and P. Wojtaszczyk, *Convergence Rates for Greedy Algorithms in Reduced Basis Methods*, SIAM Journal on Mathematical Analysis, 43 (2011), pp. 1457–1472.

[11] F. L. Bookstein, *Morphometric Tools for Landmark Data: Geometry and Biology*, Cambridge University Press, 1991.

[12] D. M. Boyer, L. Costeur, and Y. Lipman, *Earliest Record of Platychoerops (Primates, Plesiadapidae), a new species from Mouras Quarry, Mont de Berru, France*, American Journal of Physical Anthropology, 149 (2012), pp. 329–346, https://doi.org/10.1002/ajpa.22119, http://dx.doi.org/10.1002/ajpa.22119.

[13] D. M. Boyer, Y. Lipman, E. St. Clair, J. Puente, B. A. Patel, T. Funkhouser, J. Jernvall, and I. Daubechies, *Algorithms to Automatically Quantify the Geometric Similarity of Anatomical Surfaces*, Proceedings of the National Academy of Sciences, 108 (2011), pp. 18221–18226, https://doi.org/10.1073/pnas.1112822108, http://www.pnas.org/content/108/45/18221.abstract, https://arxiv.org/abs/http://www.pnas.org/content/108/45/18221.full.pdf+html.

[14] D. M. Boyer, J. Puente, J. T. Gladman, C. Glynn, S. Mukherjee, G. S. Yapuncich, and I. Daubechies, *A New Fully Automated Approach for Aligning and Comparing Shapes*, The Anatomical Record, 298 (2015), pp. 249–276.

[15] D. M. Boyer, S. Toussaint, and M. Godinot, *Postcrania of the Most Primitive Euprimate and Implications for Primate Origins*, Journal of Human Evolution, 111 (2017), pp. 202–215.

[16] A. M. Bronstein, *Spectral Descriptors for Deformable Shapes*, arXiv preprint arXiv:1110.5015, (2011).

[17] U. Castellani, M. Cristani, S. Fantoni, and V. Murino, *Sparse Points Matching by Combining*

3D Mesh Saliency with Statistical Descriptors, Computer Graphics Forum, 27 (2008), pp. 643–652.
[18] K. N. Chaudhury, Y. Khoo, and A. Singer, *Global Registration of Multiple Point Clouds using Semidefinite Programming*, SIAM Journal on Optimization, 25 (2015), pp. 468–501.
[19] K. Clarke and R. Green, *Statistical Design and Analysis for a 'Biological Effects' Study*, Marine Ecology Progress Series, (1988), pp. 213–226.
[20] K. R. Clarke, *Non-Parametric Multivariate Analyses of Changes in Community Structure*, Austral Ecology, 18 (1993), pp. 117–143.
[21] D. Cohen-Steiner and J.-M. Morvan, *Restricted Delaunay Triangulations and Normal Cycle*, in Proceedings of the nineteenth annual symposium on Computational geometry, ACM, 2003, pp. 312–321.
[22] R. R. Coifman and S. Lafon, *Diffusion Maps*, Applied and Computational Harmonic Analysis, 21 (2006), pp. 5–30, https://doi.org/10.1016/j.acha.2006.04.006, http://www.sciencedirect.com/science/article/pii/S1063520306000546. Special Issue: Diffusion Maps and Wavelets.
[23] W. Czaja and M. Ehler, *Schroedinger Eigenmaps for the Analysis of Biomedical Data*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 35 (2013), pp. 1274–1280.
[24] R. DeVore, G. Petrova, and P. Wojtaszczyk, *Greedy Algorithms for Reduced Bases in Banach Spaces*, Constructive Approximation, 37 (2013), pp. 455–466.
[25] I. L. Dryden and K. V. Mardia, *Statistical Shape Analysis*, vol. 4, John Wiley & Sons New York, 1998.
[26] D. S. Fry, *Shape Recognition Using Metrics on the Space of Shapes*, PhD thesis, Harvard University, Cambridge, MA, USA, 1993. UMI Order No. GAX94-12337.
[27] T. Gao, *Hypoelliptic Diffusion Maps and Their Applications in Automated Geometric Morphometrics*, PhD thesis, Duke University, 2015.
[28] T. Gao, *The Diffusion Geometry of Fibre Bundles*, submitted, (2016). arXiv:1602.02330.
[29] T. Gao, J. Brodzki, and S. Mukherjee, *The Geometry of Synchronization Problems and Learning Group Actions*, arXiv preprint arXiv:1610.09051, (2016).
[30] T. Gao, S. Z. Kovalsky, and I. Daubechies, *Gaussian Process Landmarking on Manifolds*, 2018, https://arxiv.org/abs/1802.03479.
[31] T. Gao, G. S. Yapuncich, I. Daubechies, S. Mukherjee, and D. M. Boyer, *Development and Assessment of Fully Automated and Globally Transitive Geometric Morphometric Methods, With Application to a Biological Comparative Dataset With High Interspecific Variation*, The Anatomical Record, (2017), pp. n/a–n/a. Submitted.
[32] L. Z. Garamszegi, *Modern Phylogenetic Comparative Methods and Their Application in Evolutionary Biology*, Concepts and Practice. London, UK: Springer, (2014).
[33] J. W. Gidley, *Paleocene Primates of the Fort Union, with Discussion of Relationships of Eocene Primates*, Proceedings of the United States National Museum, 63 (1923), pp. 1–38.
[34] P. D. Gingerich, *Cranial Anatomy and Evolution of Early Tertiary Plesiadapidae (Mammalia, Primates)*, University of Michigan Papers on Paleontology, 15 (1976), pp. 1–141.
[35] P. D. Gingerich, *Radiation of Eocene adapidae in Europe*, Geobios, 10 (1977), pp. 165–182.
[36] P. D. Gingerich, *Early Eocene Cantius torresi—oldest Primate of Modern Aspect from North America*, Nature, 319 (1986), pp. 319–321.
[37] M. Godinot, *Primate Origins: A Reappraisal of Historical Data Favoring Tupaiid Affinities*, Primate Origins: Adaptations and Evolution. Springer, New York, (2007), pp. 83–142.
[38] T. F. Gonzalez, *Clustering to Minimize the Maximum Intercluster Distance*, Theoretical Computer Science, 38 (1985), pp. 293–306.
[39] P. I. Good, *Permutation, Parametric, and Bootstrap Tests of Hypotheses (Springer Series in Statistics)*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2004.
[40] J. C. Gower, *Generalized Procrustes Analysis*, Psychometrika, 40 (1975), pp. 33–51, https://doi.org/10.1007/BF02291478.
[41] J. C. Gower and G. B. Dijksterhuis, *Procrustes Problems*, vol. 3 of Oxford Statistical Science Series, Oxford University Press Oxford, 2004.
[42] B. Güneysu, *The Feynman-Kac Formula for Schrödinger Operators on Vector Bundles over Complete Manifolds*, Journal of Geometry and Physics, 60 (2010), pp. 1997–2010.
[43] E. Harjunmaa, A. Kallonen, M. Voutilainen, K. Hämäläinen, M. L. Mikkola, and J. Jern-

VALL, *On the Difficulty of Increasing Dental Complexity*, Nature, 483 (2012), pp. 324–327.

[44] E. HARJUNMAA, K. SEIDEL, T. HÄKKINEN, E. RENVOISÉ, I. J. CORFE, A. KALLONEN, Z.-Q. ZHANG, A. R. EVANS, M. L. MIKKOLA, I. SALAZAR-CIUDAD, ET AL., *Replaying Evolutionary Transitions From the Dental Fossil Record*, Nature, 512 (2014), pp. 44–48.

[45] B. R. HASSETT AND T. LEWIS-BALE, *Comparison of 3D Landmark and 3D Dense Cloud Approaches to Hominin Mandible Morphometrics Using Structure-From-Motion*, Archaeometry, 59 (2017), pp. 191–203, https://doi.org/10.1111/arcm.12229, http://dx.doi.org/10.1111/arcm.12229. ARCH-05-0070-2015.R2.

[46] B. HELFFER AND J. SJÖSTRAND, *Puits Multiples en Mecanique Semi-Classique iv Etude du Complexe de Witten*, Communications in partial differential equations, 10 (1985), pp. 245–340.

[47] S. C. JOSHI AND M. I. MILLER, *Landmark Matching via Large Deformation Diffeomorphisms*, IEEE Transactions on Image Processing, 9 (2000), pp. 1357–1370, https://doi.org/10.1109/83.855431.

[48] D. G. KENDALL, *Shape Manifolds, Procrustean Metrics, and Complex Projective Spaces*, Bulletin of the London Mathematical Society, 16 (1984), pp. 81–121.

[49] V. KIM, Y. LIPMAN, AND T. FUNKHOUSER, *Blended Intrinsic Maps*, ACM Transactions on Graphics (Proc. SIGGRAPH), 30 (2011).

[50] C.-W. KO, J. LEE, AND M. QUEYRANNE, *An Exact Algorithm for Maximum Entropy Sampling*, Operations Research, 43 (1995), pp. 684–691.

[51] P. KOEHL AND J. HASS, *Landmark-Free Geometric Methods in Biological Shape Analysis*, Journal of The Royal Society Interface, 12 (2015), p. 20150795.

[52] S. Z. KOVALSKY, M. GALUN, AND Y. LIPMAN, *Accelerated quadratic proxy for geometric optimization*, ACM Transactions on Graphics (TOG), 35 (2016), p. 134.

[53] A. KRAUSE, A. SINGH, AND C. GUESTRIN, *Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies*, Journal of Machine Learning Research, 9 (2008), pp. 235–284.

[54] Y. LIPMAN, *Bounded Distortion Mapping Spaces for Triangular Meshes*, ACM Transactions on Graphics (TOG), 31 (2012), p. 108.

[55] Y. LIPMAN AND I. DAUBECHIES, *Conformal Wasserstein Distances: Comparing Surfaces in Polynomial Time*, Advances in Mathematics, 227 (2011), pp. 1047–1077, https://doi.org/10.1016/j.aim.2011.01.020, http://www.sciencedirect.com/science/article/pii/S0001870811000351.

[56] Y. LIPMAN, J. PUENTE, AND I. DAUBECHIES, *Conformal Wasserstein Distance: II. Computational Aspects and Extensions.*, Math. Comput., 82 (2013), http://dblp.uni-trier.de/db/journals/moc/moc82.html#LipmanPD13.

[57] Y. LIPMAN, S. YAGEV, R. PORANNE, D. W. JACOBS, AND R. BASRI, *Feature Matching with Bounded Distortion*, ACM Transactions on Graphics (TOG), 33 (2014), p. 26.

[58] Y.-S. LIU, M. LIU, D. KIHARA, AND K. RAMANI, *Salient Critical Points for Meshes*, in Proceedings of the 2007 ACM symposium on Solid and physical modeling, ACM, 2007, pp. 277–282.

[59] D. G. LOWE, *Object Recognition from Local Scale-Invariant Features*, in Computer vision, 1999. The proceedings of the seventh IEEE international conference on, vol. 2, Ieee, 1999, pp. 1150–1157.

[60] N. MANTEL, *The Detection of Disease Clustering and a Generalized Regression Approach*, Cancer research, 27 (1967), pp. 209–220.

[61] S. MELZI, E. RODOLÀ, U. CASTELLANI, AND M. M. BRONSTEIN, *Localized Manifold Harmonics for Spectral Shape Analysis*, in Computer Graphics Forum, Wiley Online Library, 2017.

[62] P. MITTEROECKER AND P. GUNZ, *Advances in Geometric Morphometrics*, Evolutionary Biology, 36 (2009), pp. 235–247.

[63] P. MITTEROECKER AND S. M. HUTTEGGER, *The Concept of Morphospaces in Evolutionary and Developmental Biology: Mathematics and Metaphors*, Biological Theory, 4 (2009), pp. 54–67.

[64] C. MOENNING AND N. A. DODGSON, *Fast Marching Farthest Point Sampling*, tech. report, University of Cambridge, Computer Laboratory, 2003.

[65] J. L. MOIGNE, *Introduction to remote sensing image registration*, in 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), July 2017, pp. 2565–2568, https://doi.org/10.1109/IGARSS.2017.8127519.

[66] B. NADLER, S. LAFON, R. R. COIFMAN, AND I. G. KEVREKIDIS, *Diffusion Maps, Spectral Clustering and Reaction Coordinates of Dynamical Systems*, Applied and Computational Harmonic Analysis,

21 (2006), pp. 113–127.
[67] A. Naor, O. Regev, and T. Vidick, *Efficient Rounding for the Noncommutative Grothendieck Inequality*, in Proceedings of the forty-fifth annual ACM symposium on Theory of computing, ACM, 2013, pp. 71–80.
[68] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, *An Analysis of Approximations for Maximizing Submodular Set Functions-I*, Mathematical Programming, 14 (1978), pp. 265–294.
[69] A. Nemirovski, *Sums of Random Symmetric Matrices and Quadratic Optimization under Orthogonality Constraints*, Mathematical programming, 109 (2007), pp. 283–317.
[70] T. E. Nichols and A. P. Holmes, *Nonparametric Permutation Tests for Functional Neuroimaging: A Primer with Examples*, Human brain mapping, 15 (2002), pp. 1–25.
[71] S. Niranjan, A. Krause, S. M. Kakade, and M. Seeger, *Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design*, in Proceedings of the 27th International Conference on Machine Learning, 2010.
[72] M. Ovsjanikov, Q.-X. Huang, and L. Guibas, *A Condition Number for Non-Rigid Shape Matching*, Computer Graphics Forum, 30 (2011), pp. 1503–1512.
[73] O. Özyeşil, N. Sharon, and A. Singer, *Synchronization Over Cartan Motion Groups via Contraction*, SIAM Journal on Applied Algebra and Geometry, 2 (2018), pp. 207–241.
[74] S. Pandey, W. Voorsluys, M. Rahman, R. Buyya, J. Dobson, and K. Chiu, *Brain Image Registration Analysis Workflow for fMRI Studies on Global Grids*, in 2009 International Conference on Advanced Information Networking and Applications, May 2009, pp. 435–442, https://doi.org/10.1109/AINA.2009.13.
[75] E. Paradis, *Analysis of Phylogenetics and Evolution with R*, Springer Science & Business Media, 2011.
[76] F. Pesarin, *Multivariate Permutation Tests: with Applications in Biostatistics*, vol. 240, Wiley Chichester, 2001.
[77] D. L. Peutrec, F. Nier, and C. Viterbo, *Precise Arrhenius Law for p-Forms: The Witten Laplacian and Morse-Barannikov Complex*, Annales Henri Poincaré, 14 (2013), pp. 567–610.
[78] I. Plyusnin, A. R. Evans, A. Karme, A. Gionis, and J. Jernvall, *Automated 3D Phenotype Analysis Using Data Mining*, PLoS One, 3 (2008), p. e1742.
[79] J. Puente, *Distances and Algorithms to Compare Sets of Shapes for Automated Biological Morphometrics*, PhD thesis, Princeton University, 2013.
[80] J. Ramsay and B. Silverman, *Functional Data Analysis*, Springer Series in Statistics, Springer, 2005.
[81] J. O. Ramsay and B. W. Silverman, *Applied Functional Data Analysis: Methods and Case Studies*, vol. 77, Springer New York, 2002.
[82] F. J. Rohlf and F. L. Bookstein, *Proceedings of the Michigan Morphometrics Workshop*, University of Michigan Museum of Zoology, 1990.
[83] V. L. Roth, *On Three-Dimensional Morphometrics, and on the Identification of Landmark Points*, in Contributions to Morphometrics, L. F. Marcus, E. Bello, and G.-V. A., eds., Museo Nacional de Ciencias Naturales, Madrid, 1993, pp. 41–61.
[84] D. E. Russell, P. Louis, and D. E. Savage, *Primates of the French Early Eocene*, University of California Publications in the Geological Sciences, 73 (1967), pp. 1–46.
[85] T. J. Santner, B. J. Williams, and W. I. Notz, *The Design and Analysis of Computer Experiments*, Springer Series in Statistics, Springer Science & Business Media, 2013.
[86] J. H. Schwartz and I. Tattersall, *Evolutionary Relationships of Living Lemurs and Lorises (Mammalia, Primates) and Their Potential Affinities with European Eocene Adapidae.*, Anthropological papers of the AMNH; v. 60, pt. 1, (1985).
[87] E. R. Seiffert, L. Costeur, and D. M. Boyer, *Tarsal Morphology of Caenopithecus, a Large Adapiform Primate from the Middle Eocene of Switzerland*, PeerJ, 3 (2015), p. e1036.
[88] E. L. Simons, *A New Eocene Primate Genus, Cantius, and a Revision of Some Allied European Lemuroids*, vol. 7, British Museum, 1962.
[89] A. Singer, *From Graph to Manifold Laplacian: The Convergence Rate*, Applied and Computational Harmonic Analysis, 21 (2006), pp. 128–134.
[90] A. Singer and H.-T. Wu, *Vector Diffusion Maps and the Connection Laplacian*, Communications on Pure and Applied Mathematics, 65 (2012), pp. 1067–1144, https://doi.org/10.1002/cpa.21395, http://dx.doi.org/10.1002/cpa.21395.

[91] J. Smith and S. Schaefer, *Bijective parameterization with free boundaries*, ACM Transactions on Graphics (TOG), 34 (2015), p. 70.

[92] O. Smolyanov, H. Weizsäcker, and O. Wittich, *Brownian Motion on a Manifold as Limit of Stepwise Conditioned Standard Brownian Motions*, Stochastic Processes, Physics and Geometry: New Interplays, II, 29 (2000), pp. 589–602.

[93] O. G. Smolyanov, H. v Weizsäcker, and O. Wittich, *Chernoff's Theorem and Discrete Time Approximations of Brownian Motion on Manifolds*, Potential Analysis, 26 (2007), pp. 1–29.

[94] A. M.-C. So, *Moment Inequalities for Sums of Random Matrices and Their Applications in Optimization*, Mathematical Programming, 130 (2011), pp. 125–151.

[95] M. L. Stein, *Interpolation of Spatial Data: Some Theory for Kriging*, Springer Science & Business Media, 2012.

[96] G. K. Tam, Z.-Q. Cheng, Y.-K. Lai, F. C. Langbein, Y. Liu, D. Marshall, R. R. Martin, X.-F. Sun, and P. L. Rosin, *Registration of 3D Point Clouds and Meshes: A Survey from Rigid to Nonrigid*, IEEE Transactions on Visualization and Computer Graphics, 19 (2013), pp. 1199–1217.

[97] K. Turner, S. Mukherjee, and D. M. Boyer, *Persistent homology transform for modeling shapes and surfaces*, Information and Inference: A Journal of the IMA, 3 (2014), pp. 310–344, https://doi.org/10.1093/imaiai/iau011, http://dx.doi.org/10.1093/imaiai/iau011.

[98] O. Van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or, *A Survey on Shape Correspondence*, Computer Graphics Forum, 30 (2011), pp. 1681–1707.

[99] N. Vitek, C. Manz, T. Gao, J. Bloch, S. Strait, and D. M. Boyer, *Semi-supervised Determination of Pseudocryptic Morphotypes Using Observer-free Characterizations of Anatomical Alignment and Shape*. Accepted, 2017.

[100] S. K. Wärmländer, H. Garvin, P. Guyomarc'h, A. Petaros, and S. B. Sholts, *Landmark Typology in Applied Morphometrics Studies: What's the Point?*, The Anatomical Record, (2018).

[101] A. Watanabe, *How Many Landmarks are Enough to Characterize Shape and Size Variation?*, PloS one, 13 (2018), p. e0198341.

[102] J. White, *Geometric Morphometric Investigation of Molar Shape Diversity in Modern Lemurs and Lorises*, The Anatomical Record: Advances in Integrative Anatomy and Evolutionary Biology: Advances in Integrative Anatomy and Evolutionary Biology, 292 (2009), pp. 701–719.

[103] E. Witten, *Supersymmetry and Morse Theory*, J. Differential Geom., 17 (1982), pp. 661–692, https://doi.org/10.4310/jdg/1214437492, https://doi.org/10.4310/jdg/1214437492.

[104] G. S. Yapuncich, E. R. Seiffert, and D. M. Boyer, *Quantification of the Position and Depth of the Flexor Hallucis Longus Groove in Euarchontans, with Implications for the Evolution of Primate Positional Behavior*, American Journal of Physical Anthropology, 163 (2017), pp. 367–406.

[105] M. L. Zelditch, D. L. Swiderski, and H. D. Sheets, *Geometric Morphometrics for Biologists: A Primer*, Academic Press, San Diego, second ed., 2012.

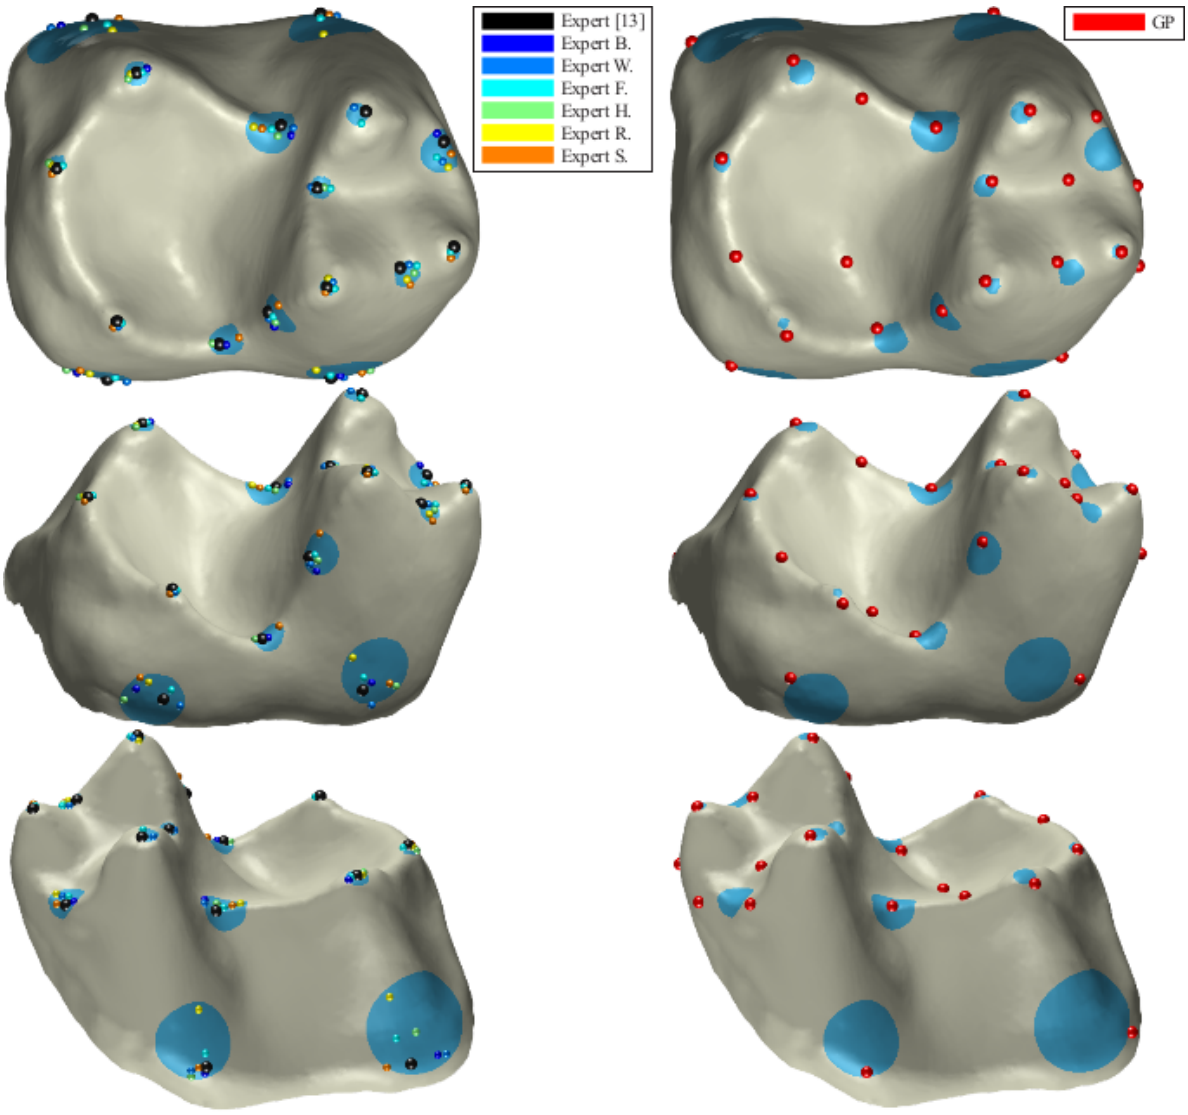

Figure 15: (a) Observer landmarks placed by multiple experts on a fossil molar. Landmarks obtained in [13], shown in black color, were identified manually by an evolutionary anthropologist and used throughout this paper as ground truth. 6 additional sets of landmarks are shown in multiple colors; these landmarks were provided by experts guided to manually place landmarks following the same landmarking scheme, so as to capture the same morphological features; see Figure Figure 16 for a visualization of the landmarks provided by each individual expert. Cyan-shaded areas depict the smallest geodesic disk enclosing each group of observer landmarks. (b) Gaussian process landmarks computed with Algorithm 2.1 are shown in red along the cyan disks representing the distribution of expert-place landmarks. This further exemplifies the potential of the first few Gaussian process landmarks in identifying biologically-informative feature points. This rows of this figure provide additional views of the example provided in Figure 6.

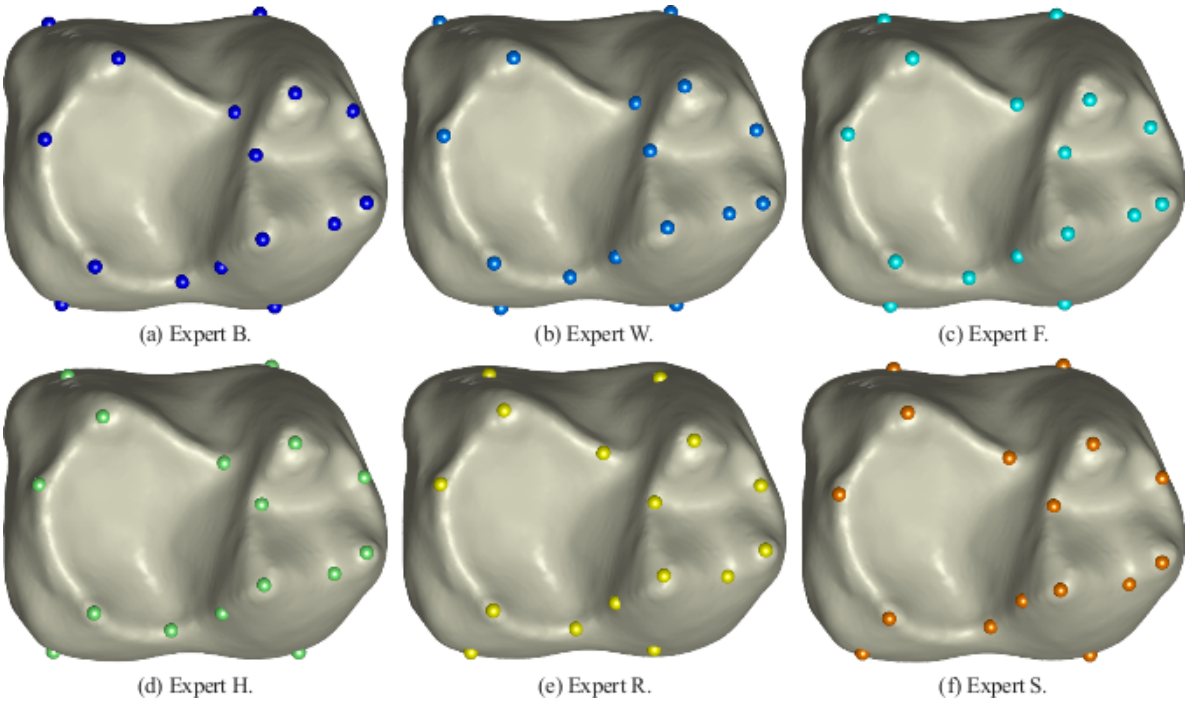

Figure 16: Multiple sets of observer landmarks placed by experts on a fossil molar. Independent experts were instructed to place landmarks following the protocol of [13] (pre-fixed ordering and/or number of landmarks of each type in the typology), the same protocol used for obtaining the landmarks used throughout this paper as ground truth. See Figure 6 and Figure 15 for a comparison of observer landmarks and Gaussian process landmarks automatically computed with Algorithm 2.1.

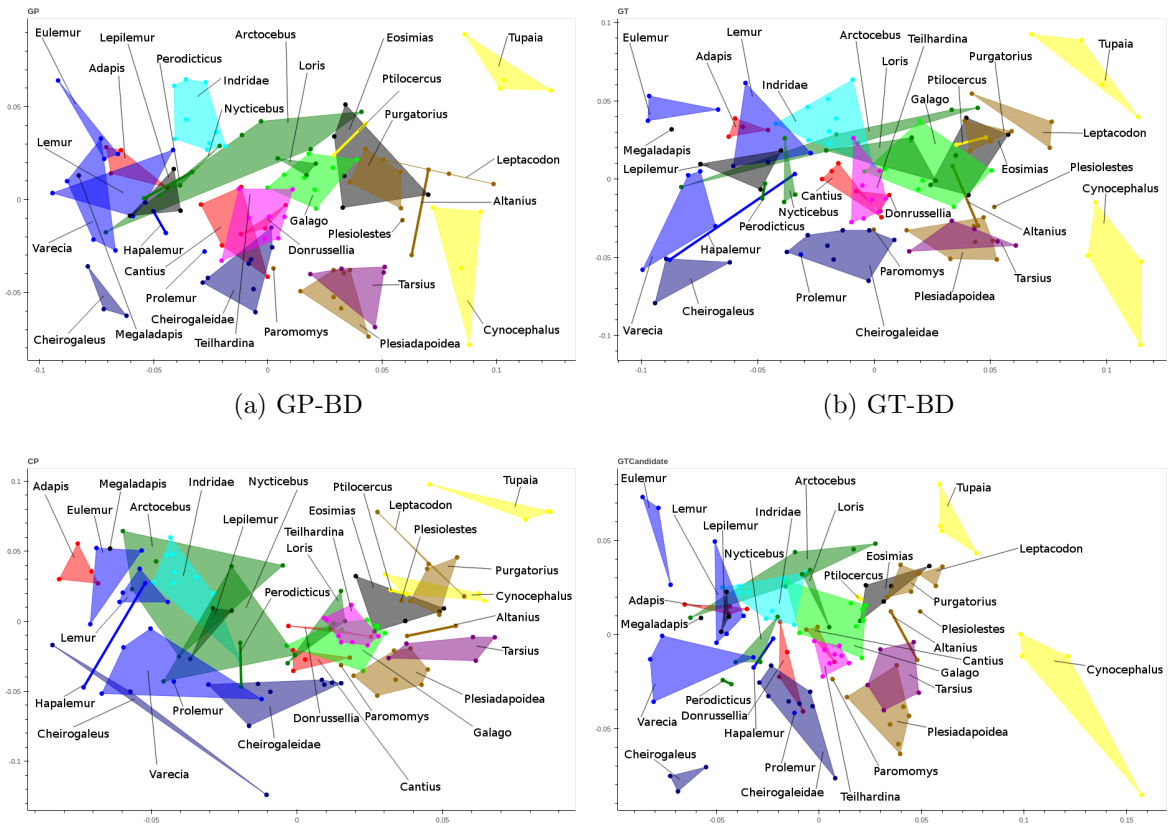

(a) GP-BD (b) GT-BD

(c) CPM (d) GT^2-BD

Figure 17: Multi-Dimensional Scaling (MDS) plots visualizing shape similarities characterized by four alternative distance matrices for a collection of 116 second mandibular molars of prosimian primates and non-primate close relatives. The distances investigated here include **(a)** Procrustes distance induced by GP-BD maps that match as many pairs of Gaussian process landmarks as possible within distortion bound $K = 1.5$; **(b)** Procrustes distance induced by GT-BD maps that match as many pairs of observer placed landmarks as possible within distortion bound $K = 1.5$, but without "ground truth" correspondences between observer placed landmarks; **(c)** continuous Procrustes distance [2]; **(d)** Procrustes distance induced by GT^2-BD maps that match as many observed placed landmarks as possible within distortion bound $K = 1.5$, with "ground truth" correspondences between observer placed landmarks as initial matching candidates. Each polygon encloses a collection of individual specimens of closely related species expected to be similar based on visual inspection and traditional comparative analyses. Similar ordination plots for GP_{Euc} or GP_{nW} are omitted in this figure because the ordination shown in those two plots are much worse than any of the plots (a)–(d). According to the PERMANOVA test results in Table 3, pairwise surface registration with GP-BD achieves comparable power of separating taxonomic groups as matching "ground truth" landmarks placed by professional comparative biologists.