

Article

A Physicist's View on Partial 3D Shape Matching

Patrice Koehl ^{1,*}  and Henri Orland ² 

¹ Department of Computer Science, University of California, Davis 95616, CA, USA

² Institut de Physique Théorique, CEA, CNRS, Université Paris-Saclay, 91191 Gif-sur-Yvette, France; henri.orland@cea.fr

* Correspondence: koehl@cs.ucdavis.edu

Abstract: A new algorithm is presented to compute nonrigid, possibly partial comparisons of shapes defined by unstructured triangulations of their surfaces. The algorithm takes as input a pair of surfaces with each surface given by a distinct and unrelated triangulation. Its goal is to define a possibly partial correspondence between the vertices of the two triangulations, with a cost associated with this correspondence that can serve as a measure of the similarity of the two shapes. To find this correspondence, the vertices in each triangulation are characterized by a signature vector of features. We tested both the LD-SIFT signatures, based on the concept of spin images, and the wave kernel signatures obtained by solving the Schrödinger equation on the triangulation. A cost matrix C is constructed such that $C(k, l)$ is the norm of the difference of the signature vectors of vertices k and l . The correspondence between the triangulations is then computed as the transport plan that solves the optimal transport or optimal partial transport problem between their sets of vertices. We use a statistical physics approach to solve these problems. The presentation of the proposed algorithm is complemented with examples that illustrate its effectiveness and manageable computing cost.

Keywords: optimal transport; shape matching; statistical physics



Citation: Koehl, P.; Orland, H. A Physicist's View on Partial 3D Shape Matching. *Algorithms* **2023**, *16*, 346. <https://doi.org/10.3390/a16070346>

Academic Editor: Frank Werner

Received: 13 June 2023

Revised: 5 July 2023

Accepted: 6 July 2023

Published: 18 July 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Our perception and analysis of the world around us have become mostly digital. Many scientific disciplines clearly benefit from this digital revolution. Scientists have created a wide variety of digital sensors that enable them to report on their research subjects at various time and length scales. This has led them to generate extensive quantitative and visual information, and the burden has now been placed into extracting knowledge from this information. In the case of 3D visual data, this amounts to identifying the shapes they contain. Computational geometry, computer vision, and computer graphics all face the challenge of developing effective algorithms to define, quantify, and compare those shapes. The advancements in machine learning and computational geometry have been extremely beneficial to those algorithms. This paper provides another source of evidence in support of enhancing such algorithms. We demonstrate that we can generate a potentially partial mapping between 3D shapes using statistical physics approaches. In our method, the cost of the correspondence acts as a gauge of the shapes' similarity. We demonstrate the efficacy of this strategy on both simulated and actual anatomical data.

Methods that compare shapes fall under the general framework of morphometrics, the study of form, a concept that includes size and shape. While morphometrics is most often associated with biology and natural sciences, its techniques apply to any shape matching problems. Two types of such techniques can usually be identified, those based on global measures of the forms and those based on computing correspondences, or maps between the shapes. We review these two approaches briefly below.

Traditionally, morphometrics rely on measurements of lengths, widths, areas, masses, ratios, and/or angles that are then compared to assess the similarities between shapes [1].

A significant drawback of such an approach is that the measurements are usually correlated and therefore include a significant amount of redundancy. Recent improvements within this approach include computation of geometric moments or Zernike moments over the shape [2–5], or of spherical harmonics over its surface [6]. All those methods are based on global properties of the forms under study, thereby preventing partial matching between such forms.

An alternate approach to shape comparison is to first build a correspondence between the shapes. Finding correspondences, or maps between shapes, is a common problem in geometric processing with a wide range of applications (for reviews, see, for example, [7–10]). Here, we briefly discuss correspondence methods that are pertinent to our study. A promising method for such shape correspondence is to view the shapes as metric spaces. Two shapes are then equal if they are isometric. Otherwise, a map is built between the two shapes that is as isometric as possible, such that the difference between the two shapes is measured as the distance between the map and the isometry. For discrete shapes equipped with discrete metrics, the difference with an isometry is measured by computing the distortion in distances between pairs of points on the surface of the shapes, where the distance can be Euclidean or geometric. This idea has led to the concepts of Gromov–Hausdorff and Gromov–Wasserstein distances between shapes (see, for example, [11–17] in the context of shape comparison). Unfortunately, computing such distance amounts to solving a quadratic assignment problem, which is NP-hard. Despite recent progress (see, for example, [18] for computing the Gromov–Wasserstein distance), efforts have focused on alternate approaches to finding shape correspondence. The first such approach proceeds by mapping the shapes into a common parameterizing metric domain, so that it is possible to directly compute a distance between points on different shapes. Methods in this category usually proceed in three steps. They first define a set of well-chosen landmarks or keypoints on the surfaces of the shapes, then assign “signatures” to those keypoints (i.e., their coordinates in the metric parameterizing domain), and finally determine a correspondence between these points, using the similarities of their signatures. Such a strategy has become standard for comparing 2D images. Methods such as SIFT [19], SURF [20] or ORB [21] are commonly used to detect keypoints within 2D images and assign them signatures. Those keypoints are then matched using techniques such as RANSAC or the iterative closest point (ICP) algorithm. The problem of keypoint detection and signature assignment is harder for 3D objects. Many methods for assigning a signature to a keypoint have been proposed [22], such as those that extrapolate the 2D signatures for images by building multiscale representations of the neighborhood of the keypoint [23–26], those that rely on the properties of the Laplace–Beltrami operator defined on the surface of interest [27–29], and those that rely on conformal mapping to a standard domain [14,30,31]. Matching the points based on their signatures is performed using the concept of “bag-of-features” to search for similar shapes, following the idea of Google search [32], the concept of shape distributions [33], by directly comparing top correspondences between shapes [29], or by using optimal transport [14,31,34–36]. The second approach consists of relaxing the requirements that the correspondence be point-wise by considering soft correspondence [16,37–39]. Finally, we briefly mention the data-driven approaches that take advantage of modern machine learning frameworks. We refer readers interested in those techniques to recent papers and surveys [22,40–45].

Our aim in this paper is to provide an alternate framework for shape comparison that falls under the shape correspondence category, but that uses methods derived from statistical physics to measure the similarities between shapes, with a special focus on partial matching (see Figure 1 for an overview). We consider shapes that are defined by their surface, usually represented by a triangulated mesh characterized with vertices, edges, and triangles. We consider all vertices in a mesh as keypoints. We test two types of signatures for the vertices, one based on a representation of their neighborhood, and one based on the properties of the Laplace Beltrami operator for the mesh. The former is based on the idea of scale-invariant spin images adapted to triangular meshes, the LD-SIFT signatures [46]. The

latter is based on the idea of solving the Schrödinger equation on the surface to characterize how waves travel on this surface and therefore capture its geometry. The corresponding signatures are referred to as wave kernel (WK) signatures [29]. The mapping between the vertices is generated from the transport plan that solves either the optimal balanced transport (OT) problem or the optimal unbalanced transport (UOT) problem. We consider the unbalanced versions of the OT problem as it is expected to solve partial matching problems. We use a statistical physics approach to solve these OT problems [47,48]. The cost associated with the optimal plan defines a distance between the two shapes. This paper does not stand on its own. The concepts of spin images and WK signatures for meshes have been proposed before. The idea of using optimal transport to compute correspondences between points describing shapes has been described in detail in the pioneering work of [14], and applied in different settings [31,34,35]. The novelty of this paper is to integrate those different components into a global physics-based approach for solving the full and partial shape registration problems. Our report should not be expected to be exhaustive: we limit ourselves to two shape signatures and two optimal transport techniques, but provide in-depth analyses of their strengths and weaknesses.

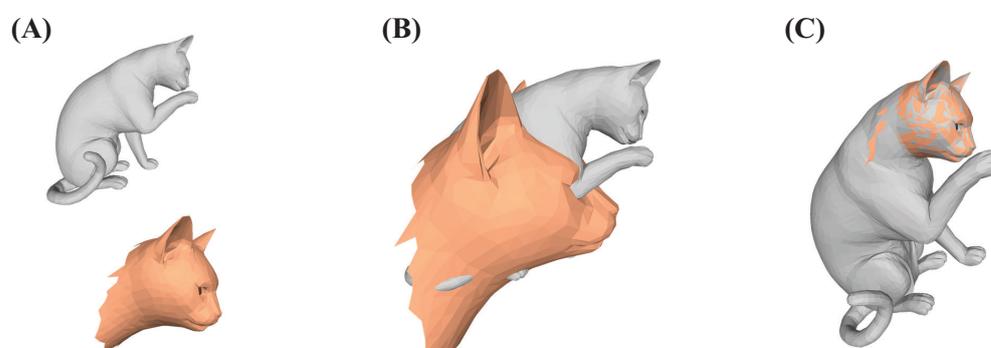


Figure 1. Comparing shapes is a central problem in computational geometry and computer graphics. This problem is exacerbated when those shapes cannot be matched partially. We propose an algorithm that combines multiple techniques from statistical physics from physics-based shape signature to alignment based on optimal transport to solve this problem. When registering the head of a cat (orange) (A) to the full corresponding cat using shape signatures and either balanced optimal transport (B), or unbalanced optimal transport (C), our algorithm enables either full or partial matching, respectively.

We are well aware that with the increase of computing power and the number of shape datasets available, deep learning techniques dominate the domain of shape comparisons (see, for example, [49–52]). Applications of deep learning, however, are contingent on the access to large datasets of shapes that are relevant to the shapes under study. It is not our intent to compete with such approaches. Instead, we focus on a physics-based approach that provides an alternate framework for solving partial shape matching. Ultimately, our formalism should prove useful for developing better loss functions for machine learning.

The paper is organized as follows. In the next section, we introduce the different elements of our framework, namely, signatures of vertices on surfaces and unbalanced optimal transport. In the results section, we compare the two types of signatures considered, as well as the two types of OT solutions proposed for registering those signatures, on nonrigid full and partial 3D matching examples, as well as on anatomical datasets. For the full shape matching, we provide comparisons with other methods based on the SHREC19 benchmark [53]. The following section includes a discussion on the differences between the two signatures we have considered, as well as on the differences between the two OT frameworks. The summary and conclusion section highlights possible future developments.

2. A Physics Approach to Shape Comparison

2.1. Basic Ideas

Let S_1 and S_2 be two surfaces of genus zero, represented by the meshes \mathcal{M}_1 and \mathcal{M}_2 , respectively. Both meshes are taken to be triangular, with $\mathcal{M} = (V, E, T)$ and $V = \{v_i\}$, $E = \{e_{ij}\}$, and $T = \{t_{ijk}\}$ denoting the vertices, edges and triangles, respectively. Our objective is to create a potentially partial correspondence between the vertices of the two meshes, along with a cost that can be used to gauge how similar the two surfaces are. To find this correspondence, we characterize the vertices either with their LD-SIFT signatures [46], or with their wave kernel signatures [29]. A cost matrix C is then constructed such that $C(k, l)$ is the L1 or L2 norm between the feature vectors of vertices k and l in \mathcal{M}_1 and \mathcal{M}_2 , respectively. The correspondence is then computed as the transport plan that solves the optimal transport (OT) or optimal unbalanced transport (OUT) problem between \mathcal{M}_1 and \mathcal{M}_2 . We use a statistical physics approach to solve these problems [47,54].

Mathematical details for the different steps mentioned above are provided in the next subsections.

2.2. LD-SIFT: Geometric Signatures for Meshes

We characterize the vertices of a mesh using a variant of spin images for shapes, namely, LD-SIFT invariants [46]. For sake of completeness, we briefly outline their constructions; more details can be found in the original paper. Note that we do not include the keypoints detection scheme proposed by the authors.

Let $v \in V$ be a vertex of the triangulated mesh, and let $N(v)$ be a neighborhood of v , i.e., the set of vertices in \mathcal{V} that are within a distance $S(v)$ of v . $S(v)$ is a “scale” for v , defined as $S(v) = KD(v)$, where K is a constant independent of v , and D is a local measure of the density around v , computed as

$$D(v) = \frac{1}{|O(v)|} \sum_{w \in O(v)} |v - w|, \tag{1}$$

where $O(v)$ is the set of all vertices w such that $(vw) \in E$ (i.e., $O(v)$ is the one-ring of v), and $|v - w|$ is the Euclidean distance between the two vertices v and w . Let \mathcal{C} be the covariance matrix over all vertices of $N(v)$. From its two leading eigenvectors e_1 and e_2 , we compute the vector $n = e_1 \times e_2$, where \times stands for cross product. The vector n defines a normal to the mesh at v , and the plan $P(v)$ that is perpendicular to n with v in it is referred to as the *dominant plane* of v . In general, P is the same as the tangent plane to the mesh at v . We then project all vertices in V onto $P(v)$ to generate a depth map on this plane. The square section of this map centered on v with size $S(v)$ defines a 2D image. We compute the SIFT descriptor for this image at its center (i.e., at the position of v); this descriptor corresponds to the LD-SIFT descriptor of v in the mesh. More details on the actual computation will be provided in the section Implementation.

2.3. WKS: Spectral Signatures for Meshes

The wave kernel signature (WKS) was described in detail in [29]. Here, we present its main concept and describe its implementation for surfaces represented as triangulated meshes.

The WKS is based on solving the time-dependent Schrödinger wave equation in the absence of an external potential for a nonrelativistic particle on the surface of the shape of interest. The evolution of such a particle is described by its wave function $\psi(x, t)$, which is the solution of

$$\frac{\partial \psi}{\partial t}(x, t) = i\Delta\psi(x, t), \tag{2}$$

where Δ is the Laplace–Beltrami (LB) operator. Note that this equation is complex due to the presence of i . We will write $E_0 \leq E_1 \leq \dots \leq E_k \dots$ the eigenvalues of the LB operator,

and ϕ_k the eigenvector associated to the eigenvalue E_k . For a given energy probability distribution f_E^2 with expectation value E , the wave function of the particle is given by

$$\psi_E(x, t) = \sum_k e^{iE_k t} \phi_k(x) f_E(E_k). \tag{3}$$

The probability to measure the particle at position x and time t on the surface is then $|\psi_E(x, t)|^2$. As the time parameter has no straightforward relevance to the geometry of the shape, the WKS is defined as the average probability over time to measure the particle at position s :

$$WKS(E, x) = \sum_k \phi_k^2(x) f_E^2(E_k). \tag{4}$$

Assuming that deformation of shapes are independently distributed, Aubry et al. assumed a log-normal distribution for f_E^2 , leading to the following definition for the WKS [29].

Definition 1. *The wave kernel signature at a point x of the surface considered is a real-valued function in the logarithmic energy scale $e = \log(E)$ defined as*

$$WKS(x, e) = C_e \sum_k \phi_k^2(x) e^{-\frac{e - \log(E_k)}{2\sigma^2}}, \tag{5}$$

where $C_e = \left(\sum_k e^{-\frac{e - \log(E_k)}{2\sigma^2}} \right)^{-1}$ is a normalization constant and σ is a parameter corresponding to the width of the normal distribution.

It is possible to compute the WKS $p(x) = (WKS(x, e_1), \dots, WKS(x, e_n))$ for each point of the shape, where e_{min}, \dots, e_{max} are some fixed energy values.

2.4. The Optimal Balanced Transport Problem

A full description of this formalism is provided in [47,54]. We outline the method here as it is key to the framework proposed in this paper.

The sets of vertices of the two meshes are relabeled as S_1 of size N_1 and S_2 of size N_2 . Each point k in S_1 (resp S_2) is allocated a “mass” $m_1(k)$ (respectively, $m_2(k)$). We ensure that $\sum_k m_1(k) = \sum_l m_2(l)$, namely, that we have balance. For convenience, we set those two sums to be 1. We encode the cost of transport between S_1 and S_2 as a positive matrix $C(k, l)$ with $k \in \{1 \dots, N_1\}$ and $l \in \{1 \dots, N_2\}$. $C(k, l)$ is set to the L1 or L2-norm of the difference of the signatures of the points k and l on \mathcal{M}_1 and \mathcal{M}_2 , respectively (see Implementations for the choice of the norm). The OT problem can then be stated as finding a transportation matrix G between S_1 and S_2 that minimizes the transport cost U defined as

$$U(G) = \sum_{k,l} G(k, l) C(k, l), \tag{6}$$

where the summations extend over all $k \in \{1 \dots, N_1\}$ and $l \in \{1 \dots, N_2\}$. The minimum of U is to be found for the values of $G(k, l)$ that satisfy the following constraints:

$$\begin{aligned} \forall (k, l), \quad & G(k, l) \geq 0, \\ \forall k, \quad & \sum_l G(k, l) = m_1(k), \\ \forall l, \quad & \sum_k G(k, l) = m_2(l). \end{aligned} \tag{7}$$

The solution to the OT problem provides an optimal transport plan G_{opt} and the corresponding minimum transport cost $U_{min} = U(G_{opt})$.

We adopt a reasoning common in statistical physics. A system in thermal equilibrium at a fixed temperature will sample a large number of states. The likelihood that this system will exist in any certain condition is represented by the associated Gibbs distribution. The state with the least energy is the one that is most likely. Therefore, the problem of minimizing an energy function can be recast as the problem of locating the most likely state of the system it characterizes. This system corresponds to the polytope $\mathcal{G}(S_1, S_2)$ of all transportation plans that are in compliance with the constraints defined in Equation (7). To find the most probable state, we consider the partition function $Z_\beta(S_1, S_2)$ computed over all states in this polytope, given by

$$Z_\beta(S_1, S_2) = e^{-\beta \mathcal{F}_\beta(S_1, S_2)} = \int_{G \in \mathcal{G}(S_1, S_2)} e^{-\beta U(G)} d\mu_{12}. \tag{8}$$

Note that $\beta = 1/k_B T$, with k_B the Boltzmann constant and T the temperature. $d\mu_{12}$ is an integration variable that can be assimilated to the Lebesgue measure for the space of transportation plans. $\mathcal{F}_\beta(S_1, S_2)$ is the free energy of the system. Unfortunately, it is not possible to explicitly calculate this free energy. Instead, we suggest a method that estimates the extremum of this free energy using the concept of saddle point approximation.

Considering the constraints defined in Equation (7) and the fact that $0 \leq G(k, l) \leq 1$, $Z_\beta(S_1, S_2)$ can be expressed as

$$Z = \int_0^1 \prod_{kl} dG(k, l) e^{-\beta \sum_{kl} C(k, l) G(k, l)} \prod_k \delta\left(\sum_l G(k, l) - m_1(k)\right) \prod_l \delta\left(\sum_k G(k, l) - m_2(l)\right). \tag{9}$$

We can represent the delta functions using the Fourier transform, adding new auxiliary variables $\lambda(k)$ and $\mu(l)$. Ignoring some multiplicative constants, the partition function can then be expressed as

$$Z = \int_{-i\infty}^{+i\infty} \prod_k d\lambda(k) \int_{-i\infty}^{+i\infty} \prod_l d\mu(l) \int_0^1 \prod_{k,l} dG(k, l) e^{-\beta \sum_{k,l} G(k, l) (C(k, l) + \lambda(k) + \mu(l)) + \beta \left(\sum_k \lambda(k) m_1(k) + \sum_l \mu(l) m_2(l)\right)}. \tag{10}$$

Note that the $\lambda(k)$ and $\mu(l)$ are imaginary complex numbers. Performing the integration over the variables $G(k, l)$, we obtain

$$Z = \int_{-i\infty}^{+i\infty} \prod_k d\lambda(k) \int_{-i\infty}^{+i\infty} \prod_l d\mu(l) e^{-\beta F_{eff}(\beta, \lambda, \mu)}, \tag{11}$$

where F_{eff} is a functional, or effective free energy, defined by

$$F_{eff}(\beta, \lambda, \mu) = - \left(\sum_k \lambda(k) m_1(k) + \sum_l \mu(l) m_2(l) \right) - \frac{1}{\beta} \sum_{kl} \ln \left[\frac{1 - e^{-\beta(C(k, l) + \lambda(k) + \mu(l))}}{\beta(C(k, l) + \lambda(k) + \mu(l))} \right]. \tag{12}$$

The saddle point approximation (SPA) is performed in the complex plane of $\{\lambda_k\}$ and $\{\mu_l\}$ and identifies the most probable state, i.e., the most probable transport plan G from

the extrema of the effective free energy. Such an extremum is expressed with respect to the variables λ and μ :

$$\frac{\partial F_{eff}(\beta, \lambda, \mu)}{\partial \lambda_k} = 0 \quad \text{and} \quad \frac{\partial F_{eff}(\beta, \lambda, \mu)}{\partial \mu_l} = 0. \tag{13}$$

After some rearrangements, those two equations can be written as

$$\begin{aligned} \forall k, \quad \sum_l G^{MF}(k, l) &= m_1(k), \\ \forall l, \quad \sum_k G^{MF}(k, l) &= m_2(l), \end{aligned} \tag{14}$$

where,

$$G^{MF}(k, l) = \phi \left[\beta \left(C(k, l) + \lambda^{MF}(k) + \mu^{MF}(l) \right) \right], \tag{15}$$

and

$$\phi(x) = \frac{e^{-x}}{e^{-x} - 1} + \frac{1}{x}. \tag{16}$$

The superscript *MF* refers to “mean field”, i.e., the solution at the SPA.

In [54], we showed that the Hessian of the effective energy $F_{eff}(\beta, \lambda, \mu)$ is negative semidefinite with $(N_1 + N_2 - 1)$ strictly negative eigenvalues and one zero eigenvalue. Furthermore, the eigenvector corresponding to the zero eigenvalue is $(1, \dots, 1, -1, \dots, -1)$ (with N_1 coordinates equal to 1, and N_2 equal to -1). Setting one of the parameters λ_k^{MF} or μ_k^{MF} to zero, the free energy function on this restricted parameter space is strictly concave.

The characteristics of the free energy functional draw attention to a number of benefits of the suggested framework, which recasts the optimal transport problem as a process that is temperature-sensitive. The optimal transport problem is first transformed into a strongly concave problem with a single solution for each temperature. Equation (12) can be used to identify the maximum of the free energy functional thanks to its concavity. Second, the modified problem specifies an optimal transport plan for each temperature that converges to the actual optimal transport plan when $T \rightarrow 0$. Finally, the convergence is monotone as a function of temperature.

We associate the transport plan at the maximum of the free energy functional with an optimum mean field energy $U_\beta^{MF}(S_1, S_2)$ and free energy $F_\beta^{MF}(S_1, S_2)$. These energies satisfy some important properties. Namely, they are monotonic decreasing functions of the parameter β that converge to the balanced transport distance between S_1 and S_2 . In addition, they are metric [54]. The quantity $U_\beta^{MF}(S_1, S_2)$ is then assimilated to a distance between the two sets of points.

2.5. The Optimal Unbalanced Transport Problem

A full description of this formalism is provided elsewhere [48]. It is outlined here for sake of clarity.

We are still concerned with two sets of points S_1 of size N_1 and S_2 of size N_2 . Each point k in S_1 (resp S_2) is assigned a mass $m_1(k)$ (resp $m_2(k)$), but this time the masses are variable, to allow for partial transport. Simply removing the constraint of fixed masses leads to a trivial solution: any solver would concentrate the mass on one point k for S_1 and l for S_2 , chosen such as $C(k, l)$ is as a minimum among all values in the cost matrix C , and set the transport to be 1 between those two points, and to 0 for all other pairs of points. This is guaranteed to achieve a minimum for the energy U , but this minimum is uninteresting. Following the idea of a penalization on the mass constraints (see [55] for a discussion on this approach), we define the discrete optimal unbalanced transport problem as finding

a transport matrix G as well as the masses m_1 and m_2 of those points that minimize the functional U defined by

$$U(S) = \sum_{k,l} G(k,l)C(k,l) + D(m_1, \rho_1) + D(m_2, \rho_2), \tag{17}$$

where $S = (G, m_1, m_2)$ is a “state” describing possibly partial mass transport between S_1 and S_2 , ρ_1 and ρ_2 are reference probability measures, and D is a distance between the variable masses m and those reference distributions ρ . There are many possible choices for the distance D . For example, Georgiou et al. [56] set D to be the total variance between the distributions, i.e.,

$$D(m, \rho) = \sum_k |m(k) - \rho(k)|.$$

Another option is to use ϕ divergences [55,57,58], such as $\phi(t) = t \log(t) - t + 1$, in which case D is the Kullback-Leibler (KL) divergence:

$$D(m, \rho) = \sum_k m(k) \left(\log \left(\frac{m(k)}{\rho(k)} \right) - m(k) + \rho(k) \right).$$

We instead used a Pearson χ^2 divergence, for reasons that will become clear below:

$$D(m, \rho) = \sum_k \frac{m(k)^2}{\rho(k)^2}.$$

The energy in Equation (17) can then be rewritten as

$$U = \sum_{k,l} C(k,l)G(k,l) + \sum_k \alpha_1(k)m_1^2(k) + \sum_l \alpha_2(l)m_2^2(l)$$

where we set $\alpha_1(k) = 1/\rho_1^2(k)$ and $\alpha_2(k) = 1/\rho_2^2(k)$.

The minimum of U is to be found for the values of $G(k,l)$, $m_1(k)$ and $m_2(l)$ that satisfy the following constraints:

$$\begin{aligned} \forall(k,l), \quad G(k,l) &\geq 0, \\ \forall k, \quad \sum_l G(k,l) &= m_1(k), \\ \forall l, \quad \sum_k G(k,l) &= m_2(l), \\ \sum_{k,l} G(k,l) &= 1. \end{aligned} \tag{18}$$

For the balanced OT problem, we reformulate the problem of minimizing the energy function given by Equation (17) as the problem of finding the most probable state of the system it defines. We define the set of all those states as \mathcal{S} . The partition function computed over this set is given by

$$Z_\beta(\mathcal{S}) = e^{-\beta\mathcal{F}_\beta(\mathcal{S})} = \int_{S \in \mathcal{S}} e^{-\beta U(S)} d\mu_{12}. \tag{19}$$

In this equation, $\mathcal{F}_\beta(\mathcal{S})$ is the free energy of the system. Taking into account the constraints defined in Equation (18), this partition function can be written as

$$\begin{aligned}
 Z_\beta(S_1, S_2) &= \int_0^1 \prod_{kl} dG(k, l) \int_{-\infty}^{+\infty} \prod_k dm_1(k) \\
 &\int_{-\infty}^{+\infty} \prod_l dm_2(l) e^{-\beta(\sum_{kl} C(k, l)G(k, l) + \sum_k \alpha_1(k)m_1^2(k) + \sum_l \alpha_2(l)m_2^2(l))} \\
 &\prod_k \delta\left(\sum_l G(k, l) - m_1(k)\right) \prod_l \delta\left(\sum_k G(k, l) - m_2(l)\right) \\
 &\delta\left(\sum_{kl} G(k, l) - 1\right).
 \end{aligned} \tag{20}$$

The integrals are over all variables $G(k, l)$, $m_1(k)$, and $m_2(l)$. We use the Fourier representation of the delta functions, thereby introducing new auxiliary variables λ_k , with $k \in \{1 \dots, N_1\}$, μ_l , with $l \in \{1 \dots, N_2\}$, and x . The partition function can then be written as (up to a multiplicative constant)

$$\begin{aligned}
 Z_\beta(S_1, S_2) &= \int_{-i\infty}^{+i\infty} \prod_k d\lambda(k) \int_{-i\infty}^{+i\infty} \prod_l d\mu(l) \int_{-i\infty}^{+i\infty} dx \\
 &\left[e^{\beta x} \int_0^1 \prod_{kl} dG(k, l) e^{-\beta \sum_{kl} G(k, l)(C(k, l) + \lambda(k) + \mu(l) + x)} \right. \\
 &\int_{-\infty}^{+\infty} \prod_k dm_1(k) e^{-\beta(\alpha_1(k)m_1^2(k) - \lambda(k)m_1(k))} \\
 &\left. \int_{-\infty}^{+\infty} \prod_l dm_2(l) e^{-\beta(\alpha_2(l)m_2^2(l) - \mu(l)m_2(l))} \right].
 \end{aligned} \tag{21}$$

Note again that we integrated the complex i from the Fourier representation into λ , μ , and x . Performing the integrations over the variables $G(k, l)$, $m_1(k)$, and $m_2(l)$ (note that the latter are analytical, due to our choice of the Pearson χ_2 divergence for D), we obtain

$$\begin{aligned}
 Z_\beta(S_1, S_2) &= \int_{-i\infty}^{+i\infty} \prod_k d\lambda(k) \int_{-i\infty}^{+i\infty} \prod_l d\mu(l) \int_{-i\infty}^{+i\infty} dx \\
 &e^{-\beta F_\beta(\lambda, \mu, x)},
 \end{aligned} \tag{22}$$

where we defined F_β as

$$\begin{aligned}
 F_\beta(\lambda, \mu, x) &= -x - \frac{1}{4} \sum_k \frac{\lambda(k)^2}{\alpha_1(k)} - \frac{1}{4} \sum_l \frac{\mu(l)^2}{\alpha_2(l)} \\
 &- \frac{1}{\beta} \sum_{kl} \ln \left[\frac{1 - e^{-\beta(C(k, l) + \lambda(k) + \mu(l) + x)}}{\beta(C(k, l) + \lambda(k) + \mu(l) + x)} \right].
 \end{aligned} \tag{23}$$

We derive a saddle point approximation (SPA) of the most probable state in \mathcal{S} by looking for extrema of this effective free energy with respect to the variables λ , μ , and x :

$$\frac{\partial F_\beta}{\partial \lambda(k)} = 0, \quad \frac{\partial F_\beta}{\partial \mu(l)} = 0, \quad \text{and} \quad \frac{\partial F_\beta}{\partial x} = 0. \tag{24}$$

After some rearrangements, those equations can be written as

$$\forall k, \sum_l H(k, l) = \frac{\lambda(k)}{2\alpha_1(k)}, \tag{25a}$$

$$\forall l, \sum_k H(k, l) = \frac{\mu(l)}{2\alpha_2(l)}, \tag{25b}$$

$$\sum_{k,l} H(k, l) = 1, \tag{25c}$$

where,

$$H(k, l) = \phi[\beta(C(k, l) + \lambda(k) + \mu(l) + x)], \tag{26}$$

and $\phi(x)$ is the same function that was used for the balanced OT problem (Equation (16)). The Hessian of F_β is negative definite and therefore the free energy function is strictly concave: it has a unique optimum. Setting $G(k, l) = H^{MF}(k, l)$ where H^{MF} is the value at this optimum, G forms an optimal transport plan between S_1 and S_2 . We can associate this transport plan with an optimum mean field energy $U_\beta^{MF}(S_1, S_2)$ and free energy $F_\beta^{MF}(S_1, S_2)$. These energies satisfy some important properties. Namely, there are monotonic decreasing functions of the parameter β that converge to the unbalanced transport distance between S_1 and S_2 . In contrast, however, with the corresponding values for the balanced OT problem [47,54], they are not metric as they do not satisfy the triangular inequality. This is a common issue with partial matching. Note also that U^{MF} and F^{MF} are not even divergences, as they are not of value zero when comparing a set of points with itself. This can easily be corrected by following an idea proposed by Peyré and collaborators [55,59] and defining

$$SU_\beta(S_1, S_2) = U_\beta^{MF}(S_1, S_2) - 0.5(U_\beta^{MF}(S_1, S_1) + U_\beta^{MF}(S_2, S_2)), \tag{27}$$

with the same correction for the free energy.

Similar to the balanced OT problem, the proposed framework for solving the unbalanced OT problem has some important advantages. The OT problem is first transformed into a strongly concave problem with a single solution for each temperature. This problem has a linear complexity in the number of variables. Simple algorithms can be used to discover the extremum of the free energy functional due to its concavity. Second, the convergence of the mean field energy to the actual OT distance is monotonic.

3. Implementation

To reach a robust implementation of the method outlined above that is fast enough to be practical requires addressing a number of concerns. The sections that follow describe how we attended to those concerns.

3.1. Computing the LD-SIFT Signature

Let \mathcal{M} be a triangular mesh, with $\mathcal{M} = (V, E, T)$, and $V = \{v\}$, $E = \{e\}$, and $T = \{t\}$ denoting its vertices, edges, and triangles, respectively. Our implementation directly follows the method described in [46], with the following practical details. The scale $S(v)$ of a vertex v in \mathcal{M} is given by Equation (1), i.e., it is the mean edge length over all edges attached to v , multiplied by a constant K which we set to 3. $N(v)$ is the set of all vertices of V that are within the distance $S(v)$ of v . Let $N = |N(v)|$ and let w be one of those vertices,

and let r_w be its position in the ambient space \mathbb{R}^3 . To compute a normal of the surface at v , we first compute a covariance matrix over $N(v)$:

$$C = \frac{1}{N} \sum_{w \in N(v)} r_w r_w^T - \frac{1}{N^2} \left(\sum_{w \in N(v)} r_w \right) \left(\sum_{w \in N(v)} r_w \right)^T. \tag{28}$$

The normal n at v is then computed as the cross product of the two dominant eigenvectors of C . The whole mesh \mathcal{M} is then rendered onto the plane $P(v)$ that is normal to n and passing through v , using a method adapted from Patch Software Render by Dirk-Jan Kroon [60]. The corresponding image size is set to 21×21 pixels, with each of its dimensions covering $2 \times S(v)$ on $P(v)$, centered at v . SIFT descriptors are then computed for the point at the center of the image (which corresponds to v) using the method implemented in VLFeat [61]. Two sets of descriptors are computed, for the image and for image rotated by 180° , and averaged out. This leads to a vector of 128 descriptors that forms the LD-descriptor of v .

Given LD-SIFT signatures for all vertices in a mesh, the distance between two such vertices is computed as the L2-norm of the differences between their signatures. Note that under this framework, vertices are represented in the same feature spaces; therefore, the same procedure applies for computing the distance between two vertices of two different meshes.

3.2. Computing the Wave Kernel Signatures

The wave kernel signatures are computed from the eigenpairs of the Laplace–Beltrami (LB) operator for the triangular mesh on the surface of a shape. Many schemes have been proposed to approximate the LB operator on a discrete surface represented by a triangular mesh [62]. We use the so-called cotangent scheme. In this scheme, we start with defining an operator LB on the mesh. This operator is given as a matrix L of size $N \times N$, where N is the number of vertices in the mesh. L is set to $M^{-1}W$, where M is a diagonal matrix whose element $M(i, i)$ corresponds to the area associated with vertex i and W is the symmetric matrix of cotangent weights, as defined in [63], to account for possible boundaries within the mesh. The general eigenproblem $W\phi = \lambda M\phi$ has real eigenvalues as solutions. We solve this problem by considering the standard eigenvalue problem $Av = \lambda v$, where $A = M^{-1/2}WM^{-1/2}$ [64]. The eigenvalues of the latter are the same as the eigenvalues of the generalized problem, and its eigenvectors v are associated with the eigenvectors ϕ of the former according to $\phi = M^{-1/2}v$.

In all the numerical experiments presented in the next section, the parameters were fixed. We computed $N = 500$ eigenvalues of the LB operator and we evaluated at $M = 200$ values of e . We used $e_{min} = \log(E_1) + 2\sigma$, where E_1 is the smallest eigenvalue different from 0 and $e_{max} = \log(E_N) - 2\sigma$, where E_N is the N -th eigenvalue. The increment δ in energy was set to $\frac{e_{max} - e_{min}}{M}$, and the variance σ (see Equation (5)) was set to 7δ .

To compare the WKS at the point x on a shape S_1 and a point y on a shape S_2 , we follow the method proposed by Aubry et al. [29] and define a distance using the L_1 norm of the normalized signature difference:

$$d_{WKS}(x, y) = \sum_{k=1}^M \left| \frac{WKS(x, e_k) - WKS(y, e_k)}{WKS(x, e_k) + WKS(y, e_k)} \right|. \tag{29}$$

Note that this implies that the same energy spectrum is used when computing the WKS for both shapes.

3.3. Solving the Balanced and Unbalanced Optimal Transport Problems

Implementations of the physics-based approaches to solving the balanced and unbalanced OT problems are similar, and follow Algorithm 1. We note first that the procedures described in Sections 2.4 and 2.5 provide a scheme for computing the transportation plan

between two sets of points at a given finite temperature T (or equivalently at a given β value, where $\beta = 1/T$). This temperature can be seen as a regularization parameter, and the actual optimal transportation plan is obtained when $\beta \rightarrow +\infty$. It is possible to run the procedure directly at a low temperature. We found, however, that it is best to use an annealing procedure, starting with small β and gradually increasing this parameter [47,54]. The solution at a given β serves as input for the following β value. This annealing procedure is efficient as we know that the mean field energy values are monotonic functions of β .

Algorithm 1 A temperature-dependent framework for solving the balanced or unbalanced optimal transport problem.

Input: The two sets of points S_1 and S_2 , the cost matrix C between S_1 and S_2 . Initial value β_0 for β . For the balanced OT problem, the masses m_1 and m_2 associated with S_1 and S_2 , respectively

Initialize: Initialize arrays λ and μ to 0, and initialize $x = 0$ (unbalanced OT). Set $\beta^0 = \beta_0/STEP$, where $STEP$ is set to $\sqrt{10}$.

for $k = 1, \dots$ until convergence **do**

(1) Initialize $\beta^k = STEP * \beta^{k-1}$.

(2) Solve the system of non linear Equations (Equation (14) for balanced, Equation (25) for unbalanced OT) for λ, μ and x at saddle point

(3) Compute optimal transport plan G_β^{MF} , and $U^{MF}(\beta^k)$. For unbalanced OT, compute masses m_1^{MF}, m_2^{MF} .

(4) Check for convergence: if $|U^{MF}(\beta^k) - U^{MF}(\beta^{k-1})|/U^{MF}(\beta^{k-1}) < TOL$, stop

end for

Output: The converged transport plan $G_\beta^{MF}(k,l)$, the corresponding transport cost $U^{MF}(\beta)$. For unbalanced OT, the optimized masses m_1^{MF}, m_2^{MF} .

At each β value, the nonlinear systems of equations defined by Equations (14) and (25) for the balanced and unbalanced OT problem, respectively, are solved using an iterative Newton–Raphson method [47,54]. Once the system is solved, the optimal transportation plan G_β^{MF} and the corresponding transport energy $U^{MF}(\beta)$ are computed. If the latter has converged, i.e., does not differ from the value at the previous β value more than a tolerance TOL , the program is terminated as the method is thought to have converged.

The primary computational cost of this algorithm is associated with solving the nonlinear systems of equations, which needs to be repeated at each value of β . We use an iterative Newton–Ralphson method that solves a linearized system of equations at each iteration. We tested both a direct solver based on LU decomposition and an iterative conjugate gradient (CG) method to solve this system. This will be discussed in the results section, under the computing time considerations.

4. Experimental Results

In this section, we present experimental results highlighting the relative advantages of using either the LD-SIFT or WKS vertex signatures combined with balanced or partial OT to find the correspondences between vertices, for full and partial 3D shape comparisons. We perform experiments focused on shape similarity and correspondence.

4.1. Full Shape Similarity

4.1.1. Synthetic Data: TOSCA

We tested our different procedures first on the synthetic TOSCA dataset [65,66], available at http://tosca.cs.technion.ac.il/book/resources_data.html, accessed on 1 March 2019. We used eight classes of objects from this dataset (see Figure 2: cat, centaur, horse, dog, seahorse, two male figures (Michael and David), and one female figure (Victoria)). Each class consists of the same shapes under different poses and/ or different triangulations. The differences between the poses echo nonrigid transformations within objects (see [65,66])

for details). In addition, the poses within a class are represented with different topologies (i.e., different triangulations), and may even have different genera. The corresponding dataset, TOSCA8, includes six instances in each class, i.e., a total of 48 classes. Each pose is represented with a triangulated surface mesh with approximately 3400 vertices and 6600 faces, with the exception of the seahorse meshes that include approximately 2100 vertices and 4200 faces.

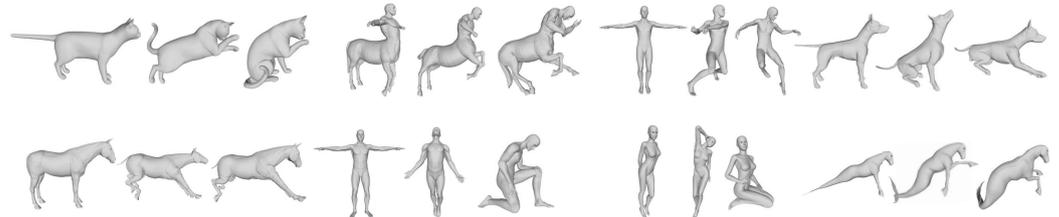


Figure 2. A few representative shapes in our TOSCA8 dataset that includes 8 classes (from left to right and top to bottom): cat, centaur, David, dog, horse, Michael, Victoria, and seahorse.

In the first set of experiments, we compared four different approaches for shape similarity, with two different signatures for representing the vertices of the shape, LD-SIFT and WKS, combined with two different methods to compute the correspondences between those vertices, namely, balanced (with fixed masses) optimal transport, which we will refer to as simply OT, and unbalanced (with variable masses) optimal transport, which we will name Partial OT. In each experiment, a pair of shapes is represented with their sets of vertices, S_1 and S_2 , and the cost matrix C between those vertices, such that $C(k, l)$ between a vertex k on shape 1 and a vertex l on shape 2 is equal to the distance between their signature vectors (see previous section for the choice of distance for the two types of signatures considered). For balanced OT, the masses of the vertices are set as uniform, while for unbalanced OT, we allow for mass creation and deletion. We computed a set of matrices $D(\beta)$ for β ranging between 1000 and 10^{11} , such that $D(\beta)(k, l)$ is the optimized transport energy $U_{\beta}^{MF}(S_k, S_l)$ (balanced OT), or the regularized optimized transport energy $SU_{\beta}^{MF}(S_k, S_l)$ (unbalanced OT), i.e., the temperature-based distance between the sets of vertices S_k and S_l of the shapes k and l . Note that for the balanced OT, $U_{\beta}^{MF}(S_k, S_l)$ is a true distance, while for the unbalanced OT, $SU_{\beta}^{MF}(S_k, S_l)$ satisfies the first two properties of a metric (identity and symmetry), but not the triangle inequality. We also computed $D(\infty)$, the matrix of distances at convergence in β . See Figure 3 for graphical representations of $D(\infty)$ for the four sets of experiments.

As observed in Figure 3, the four distance matrices are visually different. To provide a more quantitative assessment of these differences, we performed a set of classification experiments. We started with randomly selecting one shape from each of the eight classes to build a training set. We then performed 1-nearest neighbor classification experiment over the remaining shapes, where “nearest” is defined based on the distance value in one of the distance matrices. The difference between the predicted class of a shape with the actual class of that shape allows us to estimate the probability of correct classification $P(\infty)$, given a distance matrix $D(\infty)$. Results for the four sets of experiments are given in Table 1.

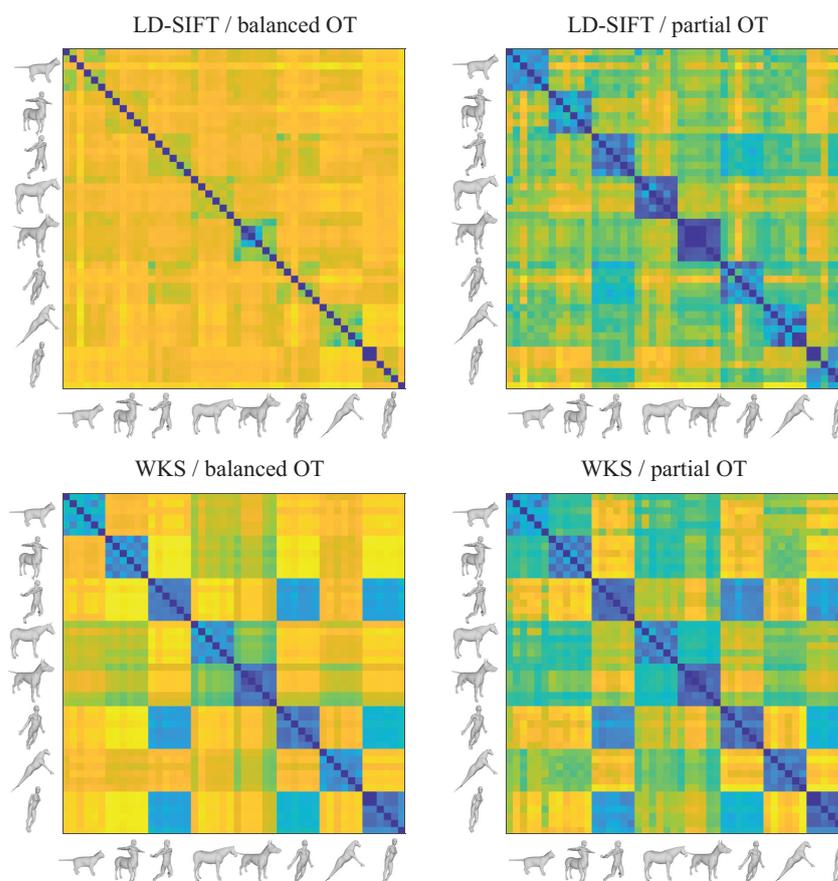


Figure 3. Distance matrices for shape similarity within the TOSCA8 dataset using the optimal transport framework with LD-SIFT (top) and WKS signatures (bottom) to represent the vertices on the shapes, and with balanced, i.e., fixed masses (left), and unbalanced, i.e., variable masses (right), optimal transport optimization to find the correspondence between vertices of two different shapes. Blue colors represent small distances (high similarity), while yellow colors represent large distances (low similarity).

Table 1. Success rates of shape classification experiments over the TOSCA8 dataset.

	3D-SIFT	WKS
OT (balanced)	83.3 ^a	100.0
OT (Partial)	97.3	98.6

^a Percentage of correctly classified shapes, computed over 10,000 experiments (see text for details).

There are a few observations we can make based on Figure 3 and Table 1. First, the WKS signatures outperform the LD-SIFT signatures for full shape comparison. Under both OT frameworks (balanced or partial), shape comparisons based on WKS capture the similarities of the shapes within their own classes (the blocks along the diagonals on the distance matrices in the bottom row of Figure 3), as well as similarities between those classes. From the WKS/balanced OT distance matrix, we observed that the four-legged shapes (cats, centaur, horse, dogs) and the human-like shapes (David, Michael, and Victoria) show similarities between themselves, but are still distinguishable from each other. The seahorses stand on their own. We assign the difference between LD-SIFT and WKS to the fact that, by design, the WKS signatures capture the geometry of the mesh representing a shape by analyzing the Laplace–Beltrami operator on that mesh, while the LD-SIFT signatures are more focused on point distributions. Second, the discriminative power of the LD-SIFT signatures is significantly improved when those signatures are used in association with

partial OT for computing correspondence. This is indicative of the importance of partial matching when those signatures are used, and will be discussed in more detail below.

4.1.2. Anatomical Data

Our second test considers three anatomical datasets, representing three different parts of the skeletons of primates. The first dataset includes 61 shapes corresponding to the proximal first metatarsals (MT1) of prosimian primates, the second dataset includes 45 shapes of the radii of apes and humans, and the third dataset includes 116 molars of prosimian primates and nonprimate close relatives (see Figure 4 for examples in each dataset) [14].

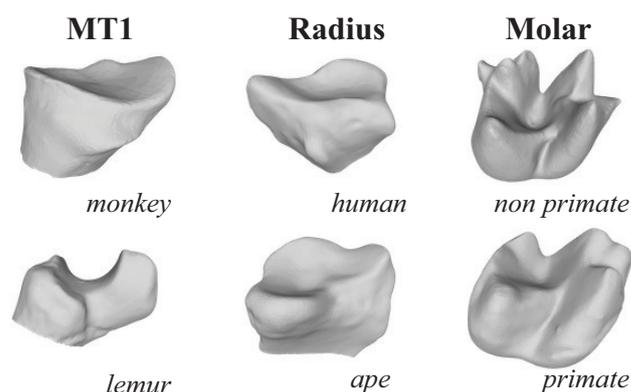


Figure 4. A few representative shapes in the metatarsal (MT1), radius, and molar anatomical datasets.

For each dataset, we performed an all-against-all comparison experiment. We used either the LD-SIFT or the WKS signatures in conjunction with unbalanced OT to find the correspondences between vertices. To evaluate the performance of our approaches, we compared the outcomes to those determined by Boyer et al. [14]. We ran two classification analyses on each dataset with a “leave one out” procedure: each surface (treated as unknown) is assigned to the taxonomic group of its nearest neighbor among the remainder of the surface in the dataset (treated as known). Tables 2–4 list success rates (in percentage) for our approaches compared to those of Boyer et al. for the three datasets.

Table 2. Success rates (percentage) of leave-one-out classification experiments for the First Metatarsal dataset.

Dataset	First Metatarsal (MT1)				
	Classification	N ^a	cP ^b	cWn ^b	LD-SIFT ^c
Genera	13	76	51	34	72
Families	9	84	69	38	88
Superfamilies	2	100	98	100	100

^a Number of groups at the taxonomic level considered. ^b Results from [14]. ^c Our results, using either LD-SIFT or WKS as signatures, and unbalanced OT for computing the correspondence.

Table 3. Success rates (percentage) of leave-one-out classification experiments for the Radius anatomical dataset.

Dataset	Radius				
	Classification	N ^a	cP ^b	cWn ^b	LD-SIFT ^c
Genera	4	76	51	34	72

^a Number of groups at the taxonomic level considered. Note that only genera information is available for this dataset. ^b Results from [14]. ^c Our results, using either LD-SIFT or WKS as signatures, and unbalanced OT for computing the correspondence.

Table 4. Success rates (percentage) of leave-one-out classification experiments for the Molar anatomical dataset.

Dataset	Molars					
	Classification	N ^a	cP ^b	cWn ^b	LD-SIFT ^c	WKS ^c
Genera		24	91	68	24	80
Families		17	92	75	28	85
Superfamilies		5	95	83	69	94

^a Number of groups at the taxonomic level considered. ^b Results from [14]. ^c Our results, using either LD-SIFT or WKS as signatures, and unbalanced OT for computing the correspondence.

By construction, the conformal Wasserstein neighborhood distance, cWn, is the closest distance to the ones we have introduced in this study, with two significant differences. First, cWn uses a different dissimilarity measure between vertices of the two surfaces under consideration, computed as follows. The procedure starts by projecting the surfaces onto the planar disk using conformal maps. The distance between two vertices, v and v' , on the two different surfaces is then computed by comparing their conformal factors (namely, parameters associated with the conformal flattening of the surface) in their neighborhood on the plane, modulo a Möbius transform (see Ref. [14] for details). All of those distances form a cost matrix between the two surfaces. Second, cWn is obtained by solving the balanced OT problem based on this cost matrix, in opposition to the unbalanced OT formalism we introduced. We found that the distance SU with the WKS signature introduced in this study outperforms cWn on all three datasets, at all phylogenetic classification levels.

Interestingly, we found that the WKS signatures outperform the LD-SIFT signatures for the full shape comparison on those anatomical datasets. We assign the difference between LD-SIFT and WKS to the fact that, by design, the WKS signatures capture the actual geometry of the mesh representing a shape by analyzing its Laplace–Beltrami operator, while the LD-SIFT signatures are more focused on local point distributions and, as such, miss the overall geometry. We will see below, however, that this may not always be a limitation.

4.1.3. Comparisons with Other Shape Matching Tools: SHREC19

We consider shape correspondence, namely, the identifications of corresponding points between two (or more) 3D shapes. Those correspondences are embedded in the optimal transport plan.

We used the SHREC19 benchmark [53] to gauge how well our algorithms can retrieve correspondences. Each 3D shape in this benchmark is given as a triangular mesh of its surface. These shapes were created using 3D scans of actual objects. Each object was captured in several poses that each corresponded to a different type of deformation. The deformations are divided into four categories, or “groups”, for classification purposes. These four categories are articulated (group 0), isometric (group 1), nonisometric (group 2), and topologic/geometric (group 3) deformations, respectively. Example of shapes for each test set are provided in Figure 5. We used the low-resolution version of this benchmark, with each shape represented with a mesh with approximately 10,000 vertices and 20,000 triangles.

SHREC19 then consists of 76 pairs of shapes regrouped in four test sets (see [53] for details). For each pair of shapes (X, Y) in the benchmark, the ground truth correspondence between the vertices of X and the vertices of Y is known.

To evaluate the quality of correspondence recovery with our algorithms, we considered normalized geodesics between the ground truth and the predicted correspondence. Briefly, for a pair of shapes (X, Y) , if p_i is a point on a shape X , y_i is its predicted correspondence

on shape Y , and z_i the ground truth position of p_i on Y . The normalized geodesic error $\epsilon(p_i)$ is computed as follows:

$$\epsilon(p_i) = \frac{d_Y(y_i, z_i)}{\text{area}(Y)^{1/2}},$$

where $d_Y(y_i, z_i)$ is the geodesic distance between y_i and z_i on the surface of Y . We use the algorithm from [67] to compute this geodesic distance.

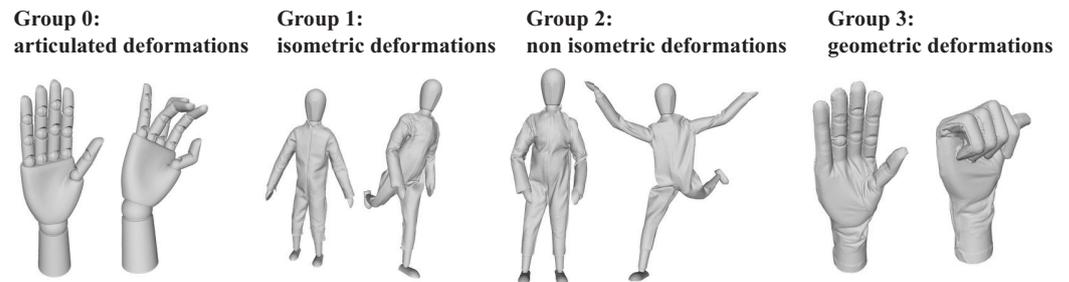


Figure 5. A few shapes for each group of the SHREC19 benchmark.

We compared the LD-SIFT and WKS vertex signatures combined with either balanced or partial OT in their abilities to define those correspondences. In each experiment, a pair of shapes i and j is represented with their complete sets of vertices, S_i and S_j . For each vertex x on i , its correspondence y on j is set to the index of the maximum value on the row corresponding to x in the OT transport plan. Figure 6 shows the corresponding cumulative geodesic errors for the four algorithms for the four test sets in SHREC19. As observed with the TOSCA8 dataset as well as with the anatomical dataset, the combination of WKS signatures with balanced OT leads to the best results for all four SHREC19 test sets.

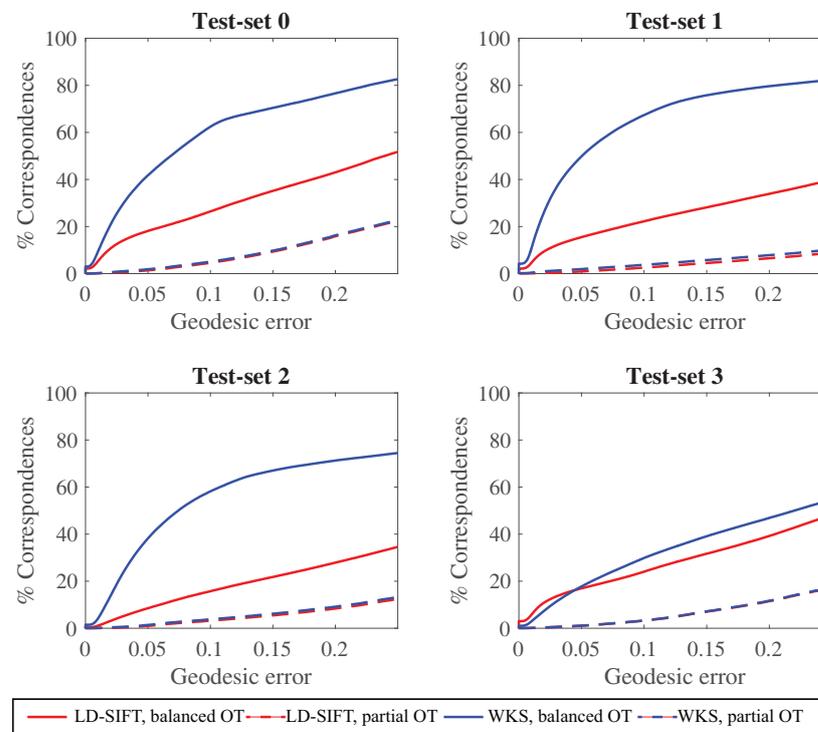


Figure 6. Cumulative distribution functions for the geodesic errors for all four methods (see text for details). Results are shown for all four test sets in SHREC19, which consider articulated deformations (test set 0), isometric deformations (test set 1), nonisometric deformations (test set 2), and topological or geometric deformations (test set 3).

The Workshop on 3D Object Retrieval, which took place in Genova, Italy, in May 2019, included a competition on comparing 3D shape registration. The SHREC19 dataset was initially utilized as a benchmark for that competition [53]. We contrast the outcomes of the five best techniques that competed with the outcomes based on the OT framework. We use the areas under the curve (AUCs) for the cumulative distribution functions of the geodesic normalized errors to evaluate the differences between the different methods: the larger the AUC value, the better the method. Results are presented in Table 5 separately for each test set in SHREC19, as well as when the test sets are combined.

Table 5. Computing correspondences between 3D shapes from the SHREC19 benchmark.

Method	Test-Set 0	Test-Set 1	Test-Set 2	Test-Set 3	All
RPTS [68]	0.920 ^a	0.926	0.824	0.929	0.899
NRP [69]	0.878	0.899	0.801	0.858	0.862
WRAP	0.853	0.920	0.772	0.870	0.856
KM [70]	0.760	0.865	0.757	0.799	0.804
GISC [71]	0.565	0.659	0.674	NA ^b	NA ^b
LD-SIFT, balanced OT	0.3	0.24	0.18	0.28	0.25
LD-SIFT, partial OT	0.08	0.04	0.05	0.06	0.05
WKS, balanced OT	0.6	0.64	0.55	0.32	0.54
WKS, partial OT	0.09	0.05	0.05	0.06	0.06

^a Area under the curve (AUC) for the cumulative distribution functions of the normalized geodesic errors in the correspondences. Results for the first five methods are derived from Ref. [53]. ^b Not available in Ref. [53].

From Table 5, we see that the correspondences computed with the WKS signatures are better than those computed with the LD-SIFT signature, and that balanced OT performs better than partial OT for full shape correspondence. Furthermore, the four registration-based methods, RTPS, NRP, WRAP, and KM, perform better than the approach that merely computes correspondences. This is probably because all of the deformations in the SHREC19 dataset are based on a mathematical morphing, which means that a mapping function is expected to be able to capture them. Finally, our formalism with the WKS signature and balanced OT performs at least as well as the genetic algorithm implemented in GISC. This genetic algorithm is the closest to the OT formalism in its concept.

4.2. Partial Shape Similarity

For nonrigid partial shape similarity, we considered four classes, the cats and horses from TOSCA8, as well as their isolated heads. The heads are taken directly from the meshes of the corresponding full shapes. Those heads contain approximately 700 vertices and 1400 faces. The corresponding dataset, HEAD, includes six instances for both the cats and horses, and all corresponding heads, for a total of 24 shapes. We ran an all-against-all comparison of all the shapes in that dataset, using the four different approaches for shape similarity described above, namely, based on two types of signatures, LD-SIFT and WKS, and two frameworks for computing the correspondences of the vertices, balanced OT and partial OT. We computed the matrix of distances at convergence in β , $D(\infty)$, for all four approaches. We expect the following results: the distance matrix should highlight high similarity within each class (cats, cat heads, horse, horse heads), as well as high similarity between the cat and the cat heads, as well as between the horse and the horse heads, but low similarities between the cats (either full, or head only) and the horses (again, either full or heads only). The actual distance matrices for the four approaches are represented graphically in Figure 7. Only the combination LD-SIFT/partial OT gives us the expected behavior, in striking difference with the results for full shape matching illustrated in Figure 3. For this combination, we even observe subdiagonals in the blocks representing the similarities between the cats and their heads, as well as for the horses and their heads, indicating that the procedure is able to identify the correct head for each cat and for each horse.

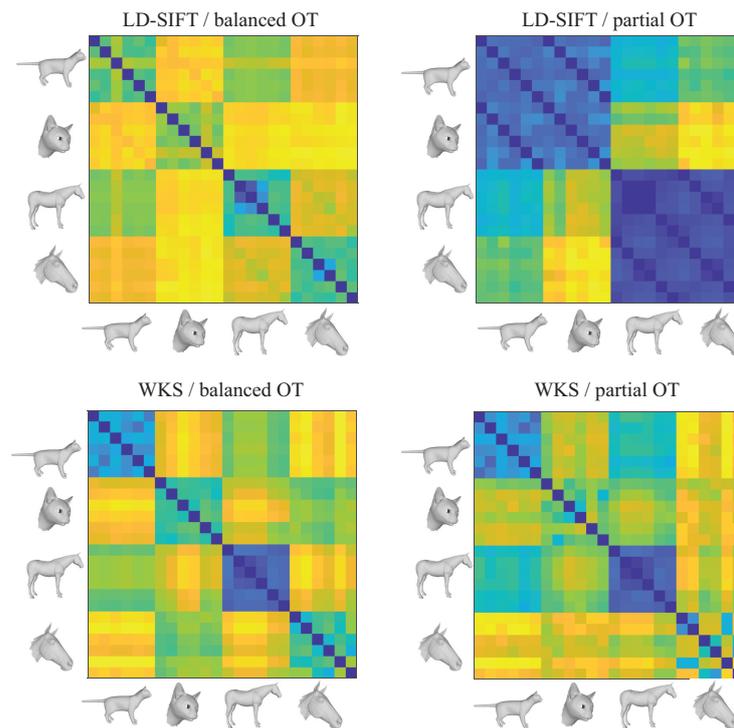


Figure 7. Distance matrices for shape similarity within the HEAD dataset using the optimal transport framework with LD-SIFT (**top**) and WKS signatures (**bottom**) to represent the vertices on the shapes, and with balanced, i.e., fixed masses (**left**), and unbalanced, i.e., variable masses (**right**), optimal transport optimization to find the correspondence between vertices of two different shapes. The color code is the same as in Figure 3.

5. Discussion

5.1. Full vs. Partial Shape Similarity

The difference between the results for full shape comparison and partial shape comparison requires some clarifications. First, the LD-SIFT signatures outperform the WKS signatures for partial shape comparison, while we observe the opposite for full shape comparison. These different behaviors are attributed to the way those signatures are generated. We illustrate the differences in Figure 8. In this figure, we represent the LD-SIFT signatures as well as the WKS signatures at corresponding vertices at the tip on one ear of a cat and at the end of one paw, for two different poses, as well as for the vertices at the ear for a mesh that only includes the head of the cat. As illustrated on the right of the figure, the WKS signatures of the corresponding vertices are very similar for the two poses of the cats. In contrast, the LD-SIFT signatures exhibit more differences. Those signatures are computed from a rendering of the mesh in a plane that is local to the vertex of interest. The renderings for the vertex on the ear are similar, while the renderings for the paw are different. The differences in the images are mitigated by the fact that the SIFT signatures are computed only at the center of the image. It remains that the LD-SIFT signatures are computed from a rendering that relies on extrinsic geometry and as such are sensitive to nonrigid deformation. The WKS signatures differ, as they are computed from the Laplace–Beltrami operator which captures the intrinsic geometry of a mesh, and are therefore less sensitive to nonrigid deformations. Such differences between LD-SIFT and WKS signatures explain the differences observed in Figure 3. Those differences, however, are reverted when we consider partial meshes. The LD-SIFT signatures at the ear of the full cat and of the shape that only includes the head of the cat are very similar (left column, top and bottom row), while the corresponding WKS signatures are very different (right column). The WKS signature is based on the solution of the Schrödinger equation over the

whole mesh describing the shape; as such, it remains a global descriptor. This explains the differences observed in Figure 7.

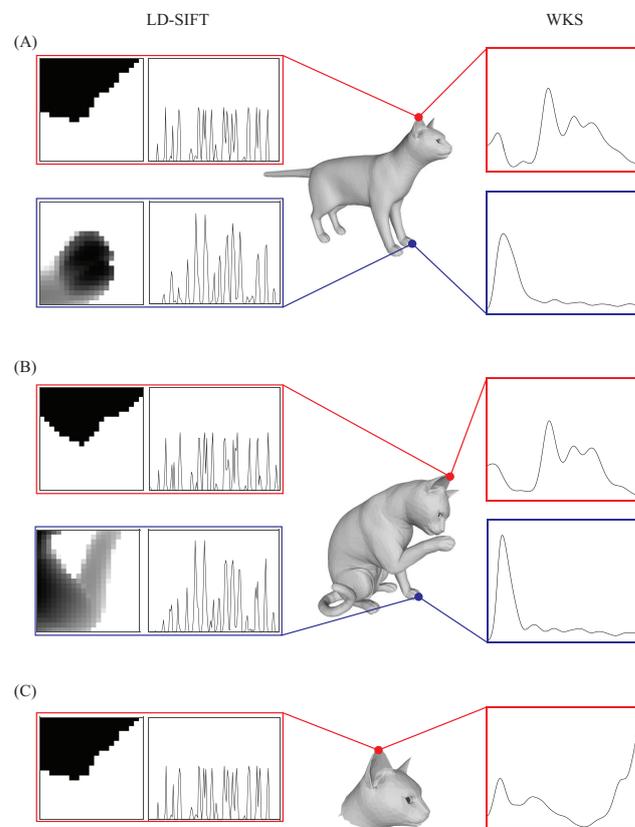


Figure 8. Cats and their WKS and LD-SIFT signatures. We consider two different shapes of the same cat (top (A) and middle row (B), respectively), as well as the head of the cat from the top row (bottom row, (C)). For the complete cats, we represent the LD-SIFT signatures (left) and the WKS signatures of a vertex at the top of their right ear (in red) and of a vertex at the end of their left forward paw (in blue). For the head we only show the signature of the vertex of the top of the right ear. The LD-SIFT signature of a vertex is computed by first rendering the mesh representing the shape onto the plane tangent to the shape at that vertex (an image of this rendering is shown on the left), and then by computing the SIFT features at the center of the corresponding image (the values of these features are shown in the box next to the images).

Second, the balanced OT procedure finds a global correspondence, while the unbalanced OT procedure allows for partial matching. We illustrate the differences in Figure 1. This figure reports on an experiment in which the head of a cat is compared to its corresponding complete cat, using LD-SIFT signatures and either balanced or unbalanced optimal transport to compute the correspondence. At convergence, the transport plan defines matching between vertices of the cat head and of the complete cat. Pairs of vertices (i, j) from the cat head and the complete cat, respectively, are considered as matches if the corresponding converged values $G^{MF}(i, j)$ are larger than $0.9 \times m_1(i)$. All the matching pairs are then used to generate a rigid body transformation (scaling + rotation + translation) between the two meshes. The balanced OT framework impose a global registration of the two sets of points. As a consequence, the final correspondence matrix leads to the head being translated towards the center of mass of the complete cat, and rotated and scaled up such that it provides a maximum coverage of the vertices of the full cat. In contrast, under the unbalanced OT framework, the amount of masses transported from each vertex from the head and received by each vertex from the cat is adapted to minimize the regularized OT energy, leading to good performance and nearly perfect match of the head onto the cat.

5.2. Computing Time

We claimed that the framework we propose in which vertices of two meshes are first characterized using signature vectors and then placed into correspondence using optimal possibly partial transport enables a robust solution to the shape comparison problem. Such a framework, however, would be nonpractical if it proved to be computationally too expensive, a common criticism for methods that rely on optimal transport. We tracked the running times for our technique for the many experiments mentioned above to make sure that this is not the case. As we solve a nonlinear system of equations iteratively at each inverse temperature, we first note that our implementation of optimal transport heavily relies on linear algebra. Each iteration requires the solution of a linear system of equations, the linearized SPA system, which can be solved either directly or iteratively. Thus, it is anticipated that parallelization will have a significant positive impact on the entire algorithm. As a result, we created two versions of our comparison process: one that utilizes GPUs and maybe multiple CPUs. Both largely rely on the processor-specific optimized BLAS and LAPACK libraries. For the direct solver of the SPA system, we used the LAPACK routine `dsysv` on CPU, and the CUSOLVER routines `DnDgetrf` and `DnDgetrs` on GPU, based on CUDA. For the alternate iterative solver, we implemented a conjugate gradient solver with an incomplete LU (ILU) preconditioner. The distributions of the computing times for the different implementations of our framework (direct or iterative solver, and CPU or GPU implementations) are shown in Figure 9, while the means of those distributions are provided in Table 6. As expected, we observe a significant speedup when the comparisons are run on multiple processors: factors of 7.8 and 9.2 for the direct and for the iterative solver on an 8 CPUs/16 threads system, respectively, and factors of 28 and 241 for the same solvers on GPU. The improvement in computing time is much greater for the iterative solver. This is expected, as direct dense solvers are not yet as optimized on GPUs than they are on CPUs. We do observe that the iterative solvers are faster than the direct solver, and we advocate using the former.

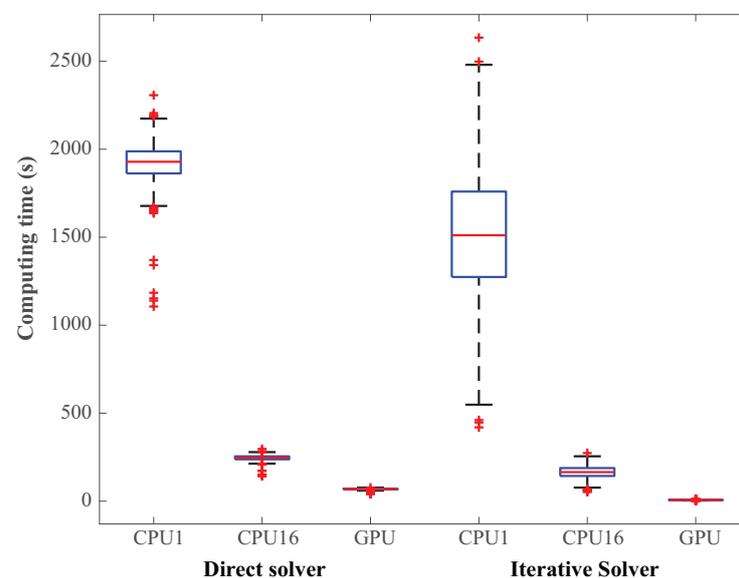


Figure 9. Boxplots of the computing times (wall clock times) of the optimal transport part of our framework for shape comparison. We compare the direct solver and the indirect solver for the linearized SPA solver, on either a single CPU (CPU1), on multiple CPUs (CPU16, i.e., Intel Core i7 processor with 8 cores/16 threads running at 4.00 GHz), and 64 GB of memory, or on a GPU system (GPU, i.e., Linux server, with Xeon Platinum 8168 CPU at 2.7 GHz, and a NVIDIA RT2080 Ti GPU card with 11 GB of memory). We included all pairwise comparisons of the shapes of TOSCA8 that are represented with around 3400 vertices (i.e., all shapes except the seahorses), with the exclusion of self-comparisons that were found to be significantly faster, for a total of 861 comparisons. The mean values for all distributions are given in Table 6.

Table 6. Mean computing times for comparing shapes from the TOSCA8 dataset ^a.

	CPU1	CPU16	GPU
Direct solver	1921 (1) ^b	246 (7.8)	68 (28)
Iterative solver	1496 (1)	163 (9.2)	6.2 (241)

^a See caption of Figure 9 for details. ^b Clock time (in s). The speedup, computed as the ratio of total computing time over clock time, is given in parentheses.

6. Summary and Conclusions

In this paper, we revisited the important problem of nonrigid 3D shape comparison for shapes represented by triangular meshes. We followed the standard framework of first computing signatures (i.e., feature vectors) for the vertices of the meshes to be compared, and then to find an optimal correspondence between those vertices that minimizes a cost matrix computed from the signatures. Our framework differs from other similar frameworks, however, in that we replaced the standard ICP procedure used to find this correspondence with a more elaborate optimal transport strategy. Such a strategy is usually deemed to be too computationally expensive. We rely on our own physics-based approach to solving the optimal transport problem as a means to circumvent this problem. This physics-based approach uses an approximation, much akin to the entropy-regularized OT method that has become popular [72]. In contrast with entropy regularization, we have convergence guarantee for our approach as well as established stability and robustness properties that enable us to use our OT solvers routinely on large systems, with confidence in their ability to generate the actual optimal correspondence. We described how we can approach balanced and unbalanced (i.e., partial) optimal transport problems with our framework, which then translate into complete shape and partial shape comparison solutions.

To find a meaningful correspondence between the vertices of two shapes to be compared requires a good estimate of the cost of associating a vertex to another. In our framework, this cost is based on computing the difference between the 3D signatures of those vertices. We used two different types of signatures, the LD-SIFT signature, which is based on the concept of shape context, i.e., an image that renders the mesh onto the tangent plane to the vertex considered, characterized with its 2D SIFT signature at the center of the image, and the WKS signature, which is based on solving the Schrödinger equation on the surface of the shape. We showed that the latter performs better for whole shape comparison in the presence of nonrigid deformation. This was attributed to the fact that WKS signatures are mostly intrinsic, while LD-SIFT are extrinsic. In contrast, however, we found that LD-SIFT signatures perform well for partial shape comparison, as WKS signatures are more global as they capture properties of the whole mesh. Different signatures can, and need to be, tested within our framework. This is currently under study.

Our implementations of the OT methods were found to be efficient, with nearly optimal use of parallelization, both on CPU and on GPU processors. We acknowledge, however, that there is room and need for improvement. The space complexity of our implementations is $O(N^2)$, as we need to store both the cost matrix and at least one work array of similar size. Those matrices are of size $N \times N$. Such a requirement limits the use of our implementations to problems of size up to a few 10,000, which falls short of the number of vertices observed in actual meshes generated by modern 3D scanners. Handling such large systems will require some redesign of our algorithms and/or the design of efficient methods for selecting a subset of vertices that are representatives of the shapes considered. This is an active area of research, which we will explore in future studies.

Author Contributions: Conceptualization, P.K. and H.O.; methodology, P.K. and H.O.; software, P.K.; formal analysis, P.K. and H.O.; investigation, P.K. and H.O.; writing: original draft preparation, P.K. and H.O. All authors have read and agreed to the published version of the manuscript.

Funding: P.K. acknowledges support from the National Science Foundation (grant no. 1760485).

Data Availability Statement: Not applicable.

Acknowledgments: P.K. acknowledges support from the NSF (grant no.1760485). The work discussed here originated from a visit by P.K. at the Institut de Physique Théorique, CEA Saclay, France, during the fall of 2022. He thanks them for their hospitality and financial support.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Marcus, L. Chapter 4. Traditional morphometrics. In *Proceedings of the Michigan Morphometric Workshop*; Rohlf, J., Bookstein, F., Eds.; The University of Michigan Museum of Zoology: Ann Harbor, MI, USA, 1990; pp. 77–122.
2. Prokop, R.; Reeves, A. A survey of moment-based techniques for unoccluded object representation and recognition. *Graph. Model. Image Process.* **1992**, *54*, 438–460. [[CrossRef](#)]
3. Novotni, M.; Klein, R. 3D Zernike descriptors for content based shape retrieval. In *Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications*, Seattle, WA, USA, 16–20 June 2003; Association for Computing Machinery: New York, NY, USA, 2003; pp. 216–225.
4. Tangelder, J. A survey of content based 3D shape retrieval methods. *Multimed. Tools Appl.* **2008**, *39*, 441–471. [[CrossRef](#)]
5. Pozo, J.; Villa-Uriol, M.C.; Frangi, A. Efficient 3D geometric and Zernike moments computation from unstructured surface meshes. *IEEE Trans. Pattern Anal. Mach. Intell.* **2011**, *33*, 471–484. [[CrossRef](#)] [[PubMed](#)]
6. Venkatraman, V.; Sael, L.; Kihara, D. Potential for protein surface shape analysis using spherical harmonics and 3D Zernike descriptors. *Cell. Biochem. Biophys.* **2009**, *54*, 23–32. [[CrossRef](#)]
7. van Kaick, O.; Zhang, H.; Hamameh, G.; Cohen-Or, D. A survey on shape correspondence. *Comput. Graph. Forum* **2011**, *30*, 1681–1707. [[CrossRef](#)]
8. Tam, G.; Cheng, Z.Q.; Lai, Y.K.; Langbein, F.; Liu, Y.; Marshall, D.; Martin, R.; Sun, X.F.; Rosin, P. Registration of 3D point clouds and meshes: A survey from rigid to nonrigid. *IEEE Trans. Vis. Comput. Graph.* **2013**, *7*, 1199–1217. [[CrossRef](#)]
9. Biasotti, S.; Cerri, A.; Bronstein, A.; Bronstein, M. Recent trends, applications, and perspectives in 3D shape similarity assessment. *Comput. Graph. Forum* **2016**, *35*, 87–119. [[CrossRef](#)]
10. Sahillioğlu, Y. Recent advances in shape correspondence. *Vis. Comput.* **2020**, *36*, 1705–1721. [[CrossRef](#)]
11. Mémoli, F.; Sapiro, G. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.* **2005**, *5*, 313–347. [[CrossRef](#)]
12. Bronstein, A.; Bronstein, M.; Kimmel, R. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proc. Natl. Acad. Sci. USA* **2006**, *103*, 1168–1172. [[CrossRef](#)]
13. Mémoli, F. On the use of Gromov-Hausdorff Distances for Shape Comparison. In *Proceedings of the Eurographics Symposium on Point-Based Graphics*, Prague, Czechia, 2–3 September 2007; pp. 81–90.
14. Boyer, D.; Lipman, Y.; StClair, E.; Puente, J.; Patel, B.; Funkhouser, T.; Jernvall, J.; Daubechies, I. Algorithms to automatically quantify the geometric similarity of anatomical surface. *Proc. Natl. Acad. Sci. USA* **2011**, *108*, 18221–18226. [[CrossRef](#)] [[PubMed](#)]
15. Mémoli, F. Gromov-Wasserstein distances and the metric approach to object matching. *Found. Comput. Math.* **2011**, *11*, 417–487. [[CrossRef](#)]
16. Solomon, J.; Peyré, G.; Kim, V.; Sra, S. Entropic Metric Alignment for Correspondence Problems. *ACM Trans. Graph.* **2016**, *35*, 1–13. [[CrossRef](#)]
17. Peyré, G.; Cuturi, M.; Solomon, J. Gromov-Wasserstein Averaging of Kernel and Distance Matrices. In *Proceedings of the International Conference on Machine Learning ICML'16*, New York City, NY, USA, 19–24 June 2016; pp. 2664–2672.
18. Chowdhury, S.; Miller, D.; Needham, T. Quantized Gromov-Wasserstein. *arXiv* **2021**, arXiv:cs.LG/2104.02013.
19. Lowe, D. Object recognition from local scale-invariant features. In *Proceedings of the International Conference of Computer Vision (ICCV)*, Corfu, Greece, 20–25 September 1999; pp. 1150–1157.
20. Bay, H.; Ess, A.; Tuytelaars, T.; Gool, L.V. SURF: Speeded up robust features. *Comput. Vis. Image Underst.* **2008**, *110*, 346–359. [[CrossRef](#)]
21. Rublee, E.; Rabaud, V.; Konolige, K.; Bradski, G. ORB: An efficient alternative to SIFT or SURF. In *Proceedings of the 2011 IEEE International Conference of Computer Vision (ICCV)*, Barcelona, Spain, 6–13 November 2011; IEEE Computer Society: Washington, DC, USA, 2011; pp. 2564–2571.
22. Georgiou, T.; Liu, Y.; Chen, W.; Lew, M. A survey of traditional and deep learning-based feature descriptors for high dimensional data in computer vision. *Int. J. Multimed. Infor. Retr.* **2020**, *9*, 135–170. [[CrossRef](#)]
23. Chua, C.; Jarvis, R. Point signatures: A new representation for 3d object recognition. *Int. J. Comput. Vision* **1997**, *25*, 63–85. [[CrossRef](#)]
24. Belongie, S.; Malik, J.; Puzicha, J. Shape context: A new descriptor for shape matching and object recognition. In *Proceedings of the NIPS*, Denver, CO, USA 27–30 November 2000; pp. 831–837.
25. Li, X.; Guskov, I. Multi-scale features for approximate alignment of point-based surfaces. In *Proceedings of the Symposium on Geometry Processing*, Vienna, Austria 4–6 July 2005; pp. 217–226.
26. Scovanner, P.; Ali, S.; Shah, M. A 3-Dimensional SIFT Descriptor and Its Application to Action Recognition. In *Proceedings of the 15th International Conference on Multimedia*, Augsburg, Germany, 25–29 September 2007; pp. 357–360.

27. Rustamov, R. Laplace–Beltrami eigenfunctions for deformation invariant shape representation. In Proceedings of the 5th Eurographics Symposium on Geometry Processing, Barcelona, Spain, 4–6 July 2007; Eurographics Association: Goslar, Germany, 2007; pp. 225–233.
28. Sun, J.; Ovsjanikov, M.; Guibas, L. A concise and provably informative multi-scale signature based on heat diffusion. *Comput. Graph. Forum* **2009**, *28*, 1383–1392. [[CrossRef](#)]
29. Aubry, M.; Schlickewei, U.; Cremers, D. The wave kernel signature: A quantum mechanical approach to shape analysis. In Proceedings of the 2011 IEEE International Conference of Computer Vision (ICCV), Barcelona, Spain, 6–13 November 2011; pp. 1626–1633.
30. Lipman, Y.; Al-Aifari, R.; Daubechies, I. The continuous Procrustes distance between two surfaces. *Commun. Pure Appl. Math.* **2013**, *66*, 934–964.
31. Su, Z.; Wang, Y.; Shi, R.; Zeng, W.; Sun, J.; Luo, F.; Gu, X. Optimal Mass Transport for Shape Matching and Comparison. *IEEE Trans. Pattern Anal. Mach. Intell.* **2015**, *37*, 2246–2259. [[CrossRef](#)]
32. Ovsjanikov, M.; Bronstein, A.; Bronstein, M.; Guibas, L. Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Trans. Graph.* **2011**, *30*, 1–20. [[CrossRef](#)]
33. Osada, R.; Funkhouser, T.; Chazelle, B.; Dobkin, D. Shape distributions. *ACM Trans. Graph.* **2002**, *21*, 807–832. [[CrossRef](#)]
34. Luo, W.; Su, Z.; Zhang, M.; Zeng, W.; Dai, J.; Gu, X. Shape signature based on Ricci flow and optimal mass transportation. *Opt. Eng.* **2014**, *53*, 112209. [[CrossRef](#)]
35. Solomon, J.; Goes, F.D.; Peyré, G.; Cuturi, M.; Butscher, A.; Nguyen, A.; Du, T.; Guibas, L. Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Trans. Graph.* **2015**, *34*, 66:1–66:11. [[CrossRef](#)]
36. Xu, L.; Sun, H.; Liu, Y. Learning With Batch-Wise Optimal Transport Loss for 3D Shape Recognition. In Proceedings of the 2019 Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 3328–3337.
37. Ovsjanikov, M.; Ben-Chen, M.; Solomon, J.; Butscher, A.; Guibas, L. Functional maps: A flexible representation of maps between shapes. *ACM Trans. Graph.* **2012**, *31*, 30. [[CrossRef](#)]
38. Solomon, J.; Nguyen, A.; Butscher, A.; Ben-Chen, M.; Guibas, L. Soft maps between surfaces. *Comput. Graph. Forum* **2012**, *31*, 1617–1626. [[CrossRef](#)]
39. Ren, J.; Poulencard, A.; Wonka, P.; Ovsjanikov, M. Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graph.* **2018**, *37*, 248. [[CrossRef](#)]
40. Boscaini, D.; Masci, J.; Rodolà, E.; Bronstein, M. Learning shape correspondence with anisotropic convolutional neural networks. In Proceedings of the NIPS, Barcelona, Spain, 5–10 December 2016; pp. 3189–3197.
41. Monti, F.; Boscaini, D.; Masci, J.; Rodolà, E.; Svoboda, J.; Bronstein, M. Geometric Deep Learning on Graphs and Manifolds Using Mixture Model CNNs. In Proceedings of the 2017 Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 5425–5434.
42. Halimi, O.; Litany, O.; Rodolà, E.; Bronstein, A.; Kimmel, R. Unsupervised Learning of Dense Shape Correspondence. In Proceedings of the 2019 Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 4365–4374.
43. Cosmo, L.; Panine, M.; Rampini, A.; Ovsjanikov, M.; Bronstein, M.; Rodolà, R. Isospectralization, or How to Hear Shape, Style, and Correspondence. In Proceedings of the 2019 Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 7521–7530.
44. Baker, N.; Lu, H.; Erikhman, G.; Kellman, P. Local features and global shape information in object classification by deep convolutional neural networks. *Vision Res.* **2020**, *172*, 46–61. [[CrossRef](#)]
45. Guo, Y.; Wang, H.; Hu, Q.; Liu, H.; Liu, L.; Bennamoun, M. Deep Learning for 3D Point Clouds: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2020**, *43*, 4338–4364. [[CrossRef](#)]
46. Darom, T.; Keller, Y. Scale Invariant Features for 3D Mesh Models. *IEEE Trans. Image Proc.* **2012**, *21*, 2758–2769. [[CrossRef](#)]
47. Koehl, P.; Delarue, M.; Orland, H. A statistical physics formulation of the optimal transport problem. *Phys. Rev. Lett.* **2019**, *123*, 040603. [[CrossRef](#)] [[PubMed](#)]
48. Koehl, P.; Delarue, M.; Orland, H. Physics approach to the variable-mass optimal-transport problem. *Phys. Rev. E* **2021**, *103*, 012113. [[CrossRef](#)] [[PubMed](#)]
49. Ioannidou, A.; Chatzilari, E.; Nikolopoulos, S.; Kompatsiaris, I. Deep Learning Advances in Computer Vision with 3D Data: A Survey. *ACM Comput. Surv.* **2017**, *50*, 1–38. [[CrossRef](#)]
50. Ahmed, E.; Saint, A.; Shabayek, A.; Cherenkova, K.; Das, R.; Gusev, G.; Aouada, D.; Ottersten, B. A survey on Deep Learning Advances on Different 3D Data Representations. *arXiv* **2018**, arXiv:1808.01462.
51. Zhou, S.K.; Greenspan, H.; Davatzikos, C.; Duncan, J.; Van Ginneken, B.; Madabhushi, A.; Prince, J.; Rueckert, D.; Summers, R. A Review of Deep Learning in Medical Imaging: Imaging Traits, Technology Trends, Case Studies With Progress Highlights, and Future Promises. *Proc. IEEE* **2021**, *109*, 820–838. [[CrossRef](#)]
52. Xu, Y.; Zheng, C.; Xu, R.; Quan, Y.; Ling, H. Multi-View 3D Shape Recognition via Correspondence-Aware Deep Learning. *IEEE Trans. Image Process.* **2021**, *30*, 5299–5312. [[CrossRef](#)] [[PubMed](#)]

53. Dyke, R.M.; Stride, C.; Lai, Y.K.; Rosin, P.L.; Aubry, M.; Boyarski, A.; Bronstein, A.M.; Bronstein, M.M.; Cremers, D.; Fisher, M.; et al. Shape Correspondence with Isometric and Non-Isometric Deformations. In Proceedings of the Eurographics Workshop on 3D Object Retrieval; Genova, Italy, 5–6 May 2019; Biasotti, S., Lavoué, G., Veltkamp, R., Eds.; The Eurographics Association: Eindhoven, The Netherlands, 2019.
54. Koehl, P.; Delarue, M.; Orland, H. Finite temperature optimal transport. *Phys. Rev. E* **2019**, *100*, 013310. [CrossRef]
55. Séjourné, T.; Freydy, J.; Vialard, F.X.; Trounev, A.; Peyré, G. Sinkhorn divergences for unbalanced optimal transport. *arXiv* **2019**, arXiv:1910.12958.
56. Georgiou, T.; Karlsson, J.; Takyar, S. Metrics for power spectra: An axiomatic approach. *IEEE Trans. Signal Proc* **2009**, *57*, 859–867. [CrossRef]
57. Chizat, L.; Peyré, G.; Schmitzer, B.; Vialard, F.X. Unbalanced optimal transport: Dynamic and Kantorovich formulations. *J. Funct. Anal.* **2018**, *274*, 3090–3123. [CrossRef]
58. Chizat, L.; Peyré, G.; Schmitzer, B.; Vialard, F.X. An interpolating distance between optimal transport and Fisher-Rao metrics. *Found. Comput. Math.* **2018**, *18*, 1–44. [CrossRef]
59. Genevay, A.; Cuturi, M.; Peyré, G.; Bach, F. Stochastic Optimization for Large-scale Optimal Transport. In *Advances in Neural Information Processing Systems 29*; Curran Associates, Inc.: Red Hook, NY, USA 2016; pp. 3440–3448.
60. Koon, D.J. Patch Software Renderer. 2020. Available online: <https://www.mathworks.com/matlabcentral/fileexchange/27084-patch-software-renderer> (accessed on 1 March 2020).
61. Vedaldi, A.; Fulkerson, B. VLFeat: An Open and Portable Library of Computer Vision Algorithms. 2008. Available online: <http://www.vlfeat.org/> (accessed on 1 March 2019).
62. Reuter, M.; Biasotti, S.; Giorgi, D.; Patané, G.; Spagnuolo, M. Discrete Laplace-Beltrami operators for shape analysis and segmentation. *Comput. Graph.* **2009**, *33*, 381–390. [CrossRef]
63. Rodolà, E.; Cosmo, L.; Bronstein, M.M.; Torsello, A.; Cremers, D. Partial Functional Correspondence. *Comput. Graph. Forum* **2017**, *36*, 222–236. [CrossRef]
64. Vallet, B.; Lévy, B. Spectral geometry processing with manifold harmonics. *Comput. Graph. Forum* **2008**, *27*, 251–260. [CrossRef]
65. Bronstein, A.; Bronstein, M.; Kimmel, R. Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Sci. Comput.* **2006**, *28*, 1812–1836. [CrossRef]
66. Bronstein, A.; Bronstein, M.; Kimmel, R. Calculus of non-rigid surfaces for geometry and texture manipulation. *IEEE Trans. Vis. Comput. Graph* **2007**, *13*, 902–913. [CrossRef]
67. Mitchell, J.; Mount, D.; Papadimitriou, C. The discrete geodesic problem. *SIAM J. Comput.* **1987**, *16*, 647–668. [CrossRef]
68. Li, K.; Yang, J.; Lai, Y.K.; Guo, D. Robust non-rigid registration with reweighted position and transformation sparsity. *IEEE Trans. Visual. Comput. Graphics* **2018**, *25*, 2255–2269. [CrossRef]
69. Dyke, R.; Lai, Y.K.; Rosin, P.; Tam, G. Non-rigid registration under anisotropic deformations. *Comput. Aided Geom. Des.* **2019**, *71*, 142–156. [CrossRef]
70. Vestner, M.; Löhner, Z.; Boyarski, A.; Litany, O.; Slossberg, R.; Remez, T.; Rodolà, E.; Bronstein, A.; Bronstein, M.; Kimmel, R.; et al. Efficient deformable shape correspondence via kernel matching. In Proceedings of the 2017 International Conference on 3D Vision (3DV), Qingdao, China, 10–12 October 2017; pp. 517–526.
71. Sahillioğlu, Y. A genetic isometric shape correspondence algorithm with adaptive sampling. *ACM Trans. Graph. (ToG)* **2018**, *37*, 1–14. [CrossRef]
72. Cuturi, M. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems 26*; Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2013; pp. 2292–2300.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.