Paretian similarity for partial comparison of non-rigid objects

Alexander M. Bronstein[1], Michael M. Bronstein[1], Alfred M. Bruckstein[1], and Ron Kimmel[1]

Department of Computer Science, Technion – Israel Institute of Technology, Haifa 32000, Israel
{bron, mbron, freddy, ron}@cs.technion.ac.il

Abstract. In this paper, we address the problem of partial comparison of non-rigid objects. We introduce a new class of set-valued distances, related to the concept of Pareto optimality in economics. Such distances allow to capture intrinsic geometric similarity between parts of non-rigid objects, obtaining semantically meaningful comparison results. The numerical implementation of our method is computationally efficient and is similar to GMDS, a multidimensional scaling-like continuous optimization problem.

1 Introduction

Analysis of non-rigid objects is an important field emerging in the pattern recognition community [16, 25, 14, 28]. Such problems arise, for example, in face recognition [6, 7], matching of articulated objects [27, 21, 30, 24, 5], image segmentation [20], texture mapping and morphing [9, 3]. A central problem is defining a meaningful criterion of similarity between non-rigid objects. Such a criterion should be invariant to deformations, have metric properties and allow for consistent discretization and efficient computation.

Theoretically, many natural deformations of objects can be modeled as near-isometric (distance preserving) transformations. The problem in this setting is translated into finding intrinsic geometric similarity between the objects. Early attempts of approximate isometry-invariant comparison were presented by Elad and Kimmel [16]. The authors proposed representing the intrinsic geometry of objects in a common metric space with simple geometry, thereby allowing to undo the degrees of freedom resulting from isometries. Representations obtained in such a way were called the bending-invariant canonical forms and were computed using *multidimensional scaling* (MDS) [1]. The Elad-Kimmel method is not exactly isometry-invariant because of the inherent error introduced by such an embedding.

Mémoli and Sapiro [25] used the Gromov-Hausdorff distance, introduced in [18] for the comparison of metric spaces. This distance has appealing theoretical properties but its computation is NP-hard. The authors proposed an algorithm that approximates the Gromov-Hausdorff distance in polynomial time by computing a different distance related to it by a probabilistic bound. In follow-up works, Bronstein *et al.* showed a different approach, according to which the computation of the Gromov-Hausdorff distance is formulated as a continuous MDS-like problem and solved efficiently using a local minimization algorithm [10, 8]. This numerical framework was given the name of *generalized MDS* (GMDS).

Here, we address an even more challenging setting of the non-rigid object analysis problem – *partial comparison* of non-rigid objects. In this setting, we need to find similarity between non-rigid objects having similar subsets. Such a situation is very common in practice, for example, in three-dimensional face recognition, where due to imperfect data acquisition, use of eyeglasses, or changes in the facial hair, parts of the objects may be missing or differ substantially [2]. Attempts to cope with such artifacts were presented in [4]. In two-dimensional shape recognition, partial comparison is an underlying problem of many shape similarity methods, attempting to divide the objects into meaningful parts, compare the parts separately and then integrate the partial similarities together [26]. Psychophysical research suggests there is strong evidence that such a "recognition by parts" mechanism is employed by the human vision [19]. Unfortunately, we do not have a clear understanding of how our brain partitions the objects we see into meaningful parts, and therefore, cannot give a precise definition of a part [27]. The recent work of Latecki *et al.* [23] allows to avoid the ambiguous definition of a part by finding a simplification of shapes which minimizes some criterion of similarity.

The main contribution of this paper is a new class of set-valued distances, related to the concept of Pareto optimality in economics. Such distances allow to capture intrinsic geometric similarity between parts of non-rigid objects. We show that the Paretian similarity can be efficiently computed using numerics resembling the GMDS. This paper is organized as follows. In Section 2, we present the theoretical background and briefly overview the properties of the Gromov-Hausdorff distance. In Section 3 we introduce the Paretian similarity, and in Section 4 show how to efficiently compute it. Section 5 demonstrates some experimental results. Though we deal with three-dimensional objects, the method is generic and can be applied to two-dimensional non-rigid objects as well (see for example [24, 5]). Section 6 concludes the paper. Due to space limitations, we do not prove our results here. The proofs will be published in the extended version of this paper.

2 Theoretical background

We model the objects as two-dimensional smooth compact connected and complete Riemannian manifolds (surfaces), possibly with boundary. We will denote the space of such objects by $\mathbb{M}$. An object $\mathcal{S} \in \mathbb{M}$ is equipped with the metric $d_{\mathcal{S}} : \mathcal{S} \times \mathcal{S} \to \mathbb{R}$, induced by the Riemannian structure; $d_{\mathcal{S}}(s, s')$ is referred to as the *geodesic distance* between the points s, s'. The Riemannian structure of the surface also defines a measure $\mu_{\mathcal{S}}(\mathcal{S}')$, which measures the area of the set $\mathcal{S}' \subset \mathcal{S}$. We denote the corresponding σ-algebra (a collection of subsets of $\mathcal{S}$ closed under countable union and complement, on which the measure $\mu_{\mathcal{S}}$ is defined) by $\Sigma_{\mathcal{S}}$.

A property will be said to hold *almost everywhere* (abbreviated as *a.e.*) on $\mathcal{S}$ if it holds on a subset $\mathcal{S}' \subseteq \mathcal{S}$ with $\mu_{\mathcal{S}}(\mathcal{S}'^c) = \mu_{\mathcal{S}}(\mathcal{S} \setminus \mathcal{S}') = 0$.

The pair $(\mathcal{S}, d_{\mathcal{S}})$ can be thought of as a metric space. In a broad sense, we refer to the distance structure of $\mathcal{S}$ as to its *intrinsic geometry*, to distinguish it from the way in which the surface is embedded into the ambient space, which is called the *extrinsic geometry*. Given a subset $\mathcal{S}' \subset \mathcal{S}$, we have two meaningful ways to define a metric on it. One possibility is to restrict the metric $d_{\mathcal{S}}$ to $\mathcal{S}'$, i.e., $d_{\mathcal{S}}|_{\mathcal{S}'}(s, s') = d_{\mathcal{S}}(s, s')$ for all s, s' in $\mathcal{S}^r$. Such metric is called the *restricted metric*. Another possibility is to derive the metric from the Riemannian structure of $\mathcal{S}'$; we call it the *induced metric* and denote it by $d_{\mathcal{S}'}$. $d_{\mathcal{S}'}$ coincides with $d_{\mathcal{S}}|_{\mathcal{S}'}$ if $\mathcal{S}'$ is *geodesically convex*. A subset $\mathcal{S}^r \subset \mathcal{S}$ with the restricted metric $d_{\mathcal{S}}|_{\mathcal{S}^r}$ is called an *r-covering* of $\mathcal{S}$ if $\mathcal{S} = \bigcup_{s \in \mathcal{S}^r} B_{\mathcal{S}}(s, r)$, where $B_{\mathcal{S}}(s_0, r) = \{s \in \mathcal{S} : d_{\mathcal{S}}(s, s_0) < r\}$ is a ball of radius r around s_0 in $\mathcal{S}$. In practical applications, finite coverings of $\mathcal{S}$ are of particular interest. Such coverings always exist assuming that $\mathcal{S}$ is compact.

Two objects $\mathcal{S}$ and $\mathcal{Q}$ are said to be *isometric* if they are identical from the point of view of their intrinsic geometry. This implies the existence of a bijective bi-Lipschitz continuous distance preserving map called an *isometry*. In practice, genuine isometries rarely exist, and objects encountered in the real life may be only nearly-isometric. A map $f : \mathcal{S} \to \mathcal{Q}$ is said to have *distortion* ϵ if

$$\mathrm{dis}\, f = \sup_{s, s' \in \mathcal{S}} |d_{\mathcal{S}}(s, s') - d_{\mathcal{Q}}(f(s), f(s'))| = \epsilon. \tag{1}$$

We say that $\mathcal{S}$ is *ϵ-isometrically embeddable* into $\mathcal{Q}$ if $\mathrm{dis}\, f \le \epsilon$. Such f is called an *ϵ-isometric embedding*. If in addition f is *ϵ-surjectivie* (i.e. $d_{\mathcal{Q}}(q, f(\mathcal{S})) \le \epsilon$ for all $q \in \mathcal{Q}$, where the set-to-point distance is defined as $d_{\mathcal{Q}}(q, f(\mathcal{S})) = \inf_{s \in \mathcal{S}} d_{\mathcal{Q}}(q, f(s))$), it is called an *$\epsilon$-isometry*, and $\mathcal{S}$ and $\mathcal{Q}$ are called *ϵ-isometric*. In [18], Mikhail Gromov introduced a criterion of similarity between metric spaces, commonly known today as the *Gromov-Hausdorff distance*. For compact spaces, it can be written in the following way:

$$d_{\mathrm{GH}}(\mathcal{Q}, \mathcal{S}) = \frac{1}{2} \inf_{\substack{f:\mathcal{S} \to \mathcal{Q} \\ g:\mathcal{Q} \to \mathcal{S}}} \max\{\mathrm{dis}\, f, \mathrm{dis}\, g, \mathrm{dis}\,(f, g)\}, \tag{2}$$

where $\mathrm{dis}\,(f, g) = \sup_{s \in \mathcal{S}, q \in \mathcal{Q}} |d_{\mathcal{S}}(s, g(q)) - d_{\mathcal{Q}}(q, f(s))|$. The Gromov-Hausdorff distance is a metric on the quotient space $\mathbb{M} \setminus \mathrm{Iso}(\mathbb{M})$, the space in which a point represents the equivalence class of all the self-isometries of an object. Another property of d_{GH} will be of fundamental importance for us: (i) if $d_{\mathrm{GH}}(\mathcal{S}, \mathcal{Q}) \le \epsilon$, then $\mathcal{S}$ and $\mathcal{Q}$ are 2ϵ-isometric; (ii) if $\mathcal{S}$ and $\mathcal{Q}$ are ϵ-isometric, then $d_{\mathrm{GH}}(\mathcal{S}, \mathcal{Q}) \le 2\epsilon$. Particularly, for $\epsilon = 0$ we have the *isometry invariance* property: $d_{\mathrm{GH}}(\mathcal{S}, \mathcal{Q}) = 0$ if and only if $\mathcal{S}$ and $\mathcal{Q}$ are isometric [13].

It may happen that $\mathcal{S}$ and $\mathcal{Q}$ are not ϵ-isometric, but parts of them are. To describe such a situation, we introduce the notion of *(λ, ϵ)-isometry*: $\mathcal{S}$ and $\mathcal{Q}$ are said to be (λ, ϵ)-isometric if there exist $\mathcal{S}' \subseteq \mathcal{S}$ and $\mathcal{Q}' \subseteq \mathcal{Q}$ with $\max\{\mu_{\mathcal{S}}(\mathcal{S}'^c), \mu_{\mathcal{Q}}(\mathcal{Q}'^c)\} \le \lambda$, such that $(\mathcal{S}', d_{\mathcal{S}'}|_{\mathcal{S}})$ and $(\mathcal{Q}', d_{\mathcal{Q}'}|_{\mathcal{Q}})$ are ϵ-isometric.

3 Paretian similarity

We can define partial similarity by saying that two objects are partially similar if they have large similar parts. What is implied by the words "similar" and "large" is a semantic question. Formally, we define a part $\mathcal{S}'$ as a subset of $\mathcal{S}$ belonging to the σ-algebra $\Sigma_{\mathcal{S}}$ (this condition is necessary in order for the part to be measurable). In our problem, it is natural to use intrinsic geometric similarity of parts, quantified by the Gromov-Hausdorff distance. The part size is quantified by the absolute or the normalized measure on the surface. We can give a more precise definition to partial similarity in the following way: two objects $\mathcal{S}$ and $\mathcal{Q}$ are partially similar if they have parts $\mathcal{S}' \in \Sigma_{\mathcal{S}}$ and $\mathcal{Q}' \in \Sigma_{\mathcal{Q}}$ of large measure $\mu_{\mathcal{S}}(\mathcal{S}')$ and $\mu_{\mathcal{Q}}(\mathcal{Q}')$, such that $(\mathcal{S}', d_{\mathcal{S}}|_{\mathcal{S}'})$ and $(\mathcal{Q}', d_{\mathcal{Q}}|_{\mathcal{Q}'})$ are nearly isometric. Note that we use the *restricted* metric on $\mathcal{S}'$ and $\mathcal{Q}'$; this fact will allow precompute the distances only ones and not recompute them every time for each $\mathcal{S}'$ and $\mathcal{Q}'$.

We will denote by $\epsilon(\mathcal{S}', \mathcal{Q}') = d_{\mathrm{GH}}(\mathcal{S}', \mathcal{Q}')$ the *similarity* of $\mathcal{S}'$ and $\mathcal{Q}'$, and by $\lambda(\mathcal{S}', \mathcal{Q}') = \max\{\mu_{\mathcal{S}}(\mathcal{S}'^c), \mu_{\mathcal{Q}}(\mathcal{Q}'^c)\}$ the *partiality*, representing the size of the region we crop off from the objects.[1] The computation of partial similarity can be formulated as a *multicriterion* optimization problem: among all the possible pairs $(\mathcal{S}', \mathcal{Q}') \in \Sigma_{\mathcal{S}} \times \Sigma_{\mathcal{Q}}$, find one that *simultaneously minimizes* ϵ and λ. In this formulation, our approach can be seen as a generalization of [23]. Obviously, in most cases it is impossible to bring both criteria to zero because they are competing. Each $(\mathcal{S}', \mathcal{Q}')$ can be represented as a point $(\lambda(\mathcal{S}', \mathcal{Q}'), \epsilon(\mathcal{S}', \mathcal{Q}'))$ in the plane. At certain points, improving one criterion inevitably compromises the other. Such solutions, representing the best tradeoff between the criteria, are called *Pareto optimal* in economics. This notion is closely related to rate-distortion analysis in information theory [15] and to receiver operation characteristics in pattern recognition [17]. We say that $(\mathcal{S}^*, \mathcal{Q}^*)$ is a Pareto optimum if at least one of the following holds,

$$\begin{aligned} &\epsilon(\mathcal{S}^*, \mathcal{Q}^*) \leq \epsilon(\mathcal{S}', \mathcal{Q}'); \quad \text{or,} \\ &\lambda(\mathcal{S}^*, \mathcal{Q}^*) \leq \lambda(\mathcal{S}', \mathcal{Q}'), \end{aligned} \tag{3}$$

for all $\mathcal{S}' \subseteq \mathcal{S}$ and $\mathcal{Q}' \subseteq \mathcal{Q}$. The set of all the Pareto optimal solutions is called the *Pareto frontier* and can be visualized as a planar curve (see Figures 1–2). Solutions below this curve do not exist.

The fundamental difference between the Paretian similarity and similarity in the traditional sense (which can be quantified by a scalar "distance" value), is the fact that we have a multitude of similarities, each corresponding to a Pareto optimum. We can think of the Pareto frontier as of a generalized, *set-valued distance*, which is denoted here by d_{P}. Set-valued distance requires a redefinition of notions commonly associated with scalar-valued distances. For instance, it is usually impossible to establish a full order relation between the distances $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S})$ and $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{R})$, since they may be mutually incompatible. We

[1] Partiality can be defined in other ways, for example, $\lambda(\mathcal{S}', \mathcal{Q}') = \mu_{\mathcal{S}}(\mathcal{S}'^c) + \mu_{\mathcal{Q}}(\mathcal{Q}'^c)$.

can only define point-wise order relations in the following way: if (λ_0, ϵ_0) is above $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S})$, we will write $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}) < (\lambda_0, \epsilon_0)$; other strong and weak inequalities are defined in a similar way. The notation $(\lambda_0, \epsilon_0) \in d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S})$ will be used to say that (λ_0, ϵ_0) is a Pareto optimum. Using this definition, we can summarize the properties of Paretian similarity as follows:

Theorem 1 (Properties of d_{P}). *The distance d_{P} satisfies:*

(P1) Non-negativity: $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}) \subseteq [0, \infty) \times [0, \infty)$.
(P2) Symmetry: $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}) = d_{\mathrm{P}}(\mathcal{S}, \mathcal{Q})$.
(P3) Monotonicity: *If $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}) \leq (\lambda, \epsilon)$, then $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}) \leq (\lambda', \epsilon)$ for every $\lambda' \geq \lambda$, and $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}) \leq (\lambda, \epsilon')$ for every $\epsilon' \geq \epsilon$.*
(P4) Partial similarity: *(i) If $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}) \leq (\lambda, \epsilon)$, then $\mathcal{S}$ and $\mathcal{Q}$ are $(\lambda, 2\epsilon)$-isometric; (ii) if $\mathcal{S}$ and $\mathcal{Q}$ are (λ, ϵ)-isometric, then $d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}) \leq (\lambda, 2\epsilon)$.*
(P5) Consistency to sampling: *If $\mathcal{S}^r$ and $\mathcal{Q}^r$ are finite r-coverings of two shapes of bounded curvature $\mathcal{S}$ and $\mathcal{Q}$, respectively, then $\lim_{r \to 0} d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S}^r) = d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S})$.*

Properties (P1)-(P5) follow from the properties of the Gromov-Hausdorff distance (see e.g. [8]). Due to space limitations, we do not give a formal proof of this theorem. $(0, 0) \in d_{\mathrm{P}}(\mathcal{Q}, \mathcal{S})$ if and only if there exists an a.e. isometry between $\mathcal{S}$ and $\mathcal{Q}$.

3.1 Converting set-valued distances into scalar-valued distances

In order to be able to compare similarities, we need to convert the set-valued distance into a traditional, scalar-valued one. The easiest way to do it is by considering a single point on the Pareto frontier. For example, we can fix the value of λ and use the corresponding distortion ϵ as the distance. We obtain a scalar-valued distance, to which we refer as the *λ-Gromov-Hausdorff distance*:

$$d_{\mathrm{GH}}^{\lambda}(\mathcal{Q}, \mathcal{S}) = \frac{1}{2} \inf_{\substack{f : \mathcal{S}' \to \mathcal{Q}' \\ g : \mathcal{Q}' \to \mathcal{S}' \\ \mathcal{S}' : \mu_{\mathcal{S}}(\mathcal{S}'^c) \leq \lambda \\ \mathcal{Q}' : \mu_{\mathcal{Q}}(\mathcal{Q}'^c) \leq \lambda}} \max\{\operatorname{dis} f, \operatorname{dis} g, \operatorname{dis}(f, g)\}.$$

The particular case of $d_{\mathrm{GH}}^{0}(\mathcal{Q}, \mathcal{S})$ can be thought of as an a.e. Gromov-Hausdorff distance. Alternatively, we can fix the value of ϵ; a scalar distance obtained this way may be useful in a practical situation when we know *a priory* the accuracy of surface acquisition and distance measurement. A third possibility is to take the area under the Pareto frontier as a scalar-valued distance.

We should note, however, that both of the above choices are rather arbitrary. A slightly more motivated selection of a single point out of the set of Pareto optimal solutions was proposed by Salukwadze [29] in the context of multicriterion optimization problems arising in control theory. Salukwadze suggested choosing a Pareto optimum, which is the closest (in sense of some norm) to some optimal, usually non-achievable point. In our case, such an optimal point is $(0,0)$. Given a Pareto frontier $d_{\mathrm{P}}(\mathcal{S},\mathcal{Q})$, we define the *Salukwadze similarity* as

$$d_{\mathrm{SAL}}(\mathcal{Q},\mathcal{S}) = \inf_{(\lambda,\epsilon)\in d_{\mathrm{P}}(\mathcal{S},\mathcal{Q})} \|(\lambda,\epsilon)\|. \tag{4}$$

Depending on the choice of the norm $\|\cdot\|$ in (4), we obtain different solutions, some of which have an explicit form. For instance, choosing the L_p-norm, we can define the L_p-*Salukwadze distance* as follows:

$$d^p_{\mathrm{SAL}}(\mathcal{Q},\mathcal{S}) = \inf_{\substack{f:\mathcal{S}'\to\mathcal{Q}'\\ g:\mathcal{Q}'\to\mathcal{S}'\\ \mathcal{S}',\mathcal{Q}'}} \left\{ \frac{1}{2^p}\max\{\mathrm{dis}\, f, \mathrm{dis}\, g, \mathrm{dis}\,(f,g)\}^p + \max\{\mu_{\mathcal{S}}(\mathcal{S}'^{\mathrm{c}}), \mu_{\mathcal{Q}}(\mathcal{Q}'^{\mathrm{c}})\}^p \right\}.$$

This formulation is a very intuitive interpretation of the multicriterion optimization problem: we are simultaneously minimizing $d_{\mathrm{GH}}(\mathcal{S}',\mathcal{Q}')$ and maximizing the measures of $\mathcal{S}'$ and $\mathcal{Q}'$. In order to avoid scaling ambiguity between the distortion and the measure, a normalization factor $\alpha\cdot d_{\mathrm{GH}}(\mathcal{S}',\mathcal{Q}')$ (where α has units of distance) can be used.

3.2 Relaxed Paretian similarity

The computation of $d_{\mathrm{P}}(\mathcal{Q},\mathcal{S})$ requires optimization over $\Sigma_{\mathcal{S}}\times\Sigma_{\mathcal{Q}}$ and is impractical, since in the discrete case it gives rise to a combinatorial optimization problem with complexity growing exponentially in the sample size of $\mathcal{S}$ and $\mathcal{Q}$. However, the problem can be relaxed by resorting to *fuzzy* representation of parts as continuous *membership functions* $m_{\mathcal{S}}:\mathcal{S}\to[0,1]$ and $m_{\mathcal{Q}}:\mathcal{Q}\to[0,1]$. The value of $m_{\mathcal{S}}$ measures the degree to which a point belongs to the part of $\mathcal{S}$ (zero implies exclusion, one implies inclusion). Instead of $\lambda(\mathcal{S}',\mathcal{Q}')$, we define the *fuzzy partiality*

$$\tilde{\lambda}(m_{\mathcal{S}},m_{\mathcal{Q}}) = \max\left\{\int_{\mathcal{S}}(1-m_{\mathcal{S}}(s))d\mu_{\mathcal{S}}, \int_{\mathcal{Q}}(1-m_{\mathcal{Q}}(q))d\mu_{\mathcal{Q}}\right\} \tag{5}$$

and instead of $\epsilon(\mathcal{S}',\mathcal{Q}')$, use a fuzzy version of the Gromov-Hausdorff distance,

$$\tilde{\epsilon}(m_{\mathcal{S}},m_{\mathcal{Q}}) = \frac{1}{2}\inf_{\substack{f:\mathcal{S}\to\mathcal{Q}\\ g:\mathcal{Q}\to\mathcal{S}}} \max\left\{\begin{array}{l} \sup\limits_{s,s'\in\mathcal{S}} m_{\mathcal{S}}(s)m_{\mathcal{S}}(s')|d_{\mathcal{S}}(s,s')-d_{\mathcal{Q}}(f(s),f(s'))| \\ \sup\limits_{q,q'\in\mathcal{Q}} m_{\mathcal{Q}}(q)m_{\mathcal{Q}}(q')|d_{\mathcal{Q}}(q,q')-d_{\mathcal{S}}(g(q),g(q'))| \\ \sup\limits_{s\in\mathcal{S},q\in\mathcal{Q}} m_{\mathcal{S}}(s)m_{\mathcal{Q}}(q)|d_{\mathcal{S}}(s,g(q))-d_{\mathcal{Q}}(f(s),q)| \\ \sup\limits_{s\in\mathcal{S}} D\,(1-m_{\mathcal{Q}}(f(s)))m_{\mathcal{S}}(s) \\ \sup\limits_{q\in\mathcal{Q}} D\,(1-m_{\mathcal{S}}(g(q)))m_{\mathcal{Q}}(q) \end{array}\right\},$$

where D is some large constant. The computation of the relaxed partial similarity requires minimization of $(\tilde{\lambda}(m_{\mathcal{S}},m_{\mathcal{Q}}),\tilde{\epsilon}(m_{\mathcal{S}},m_{\mathcal{Q}}))$ on all the pairs of membership functions $(m_\mathcal{S}, m_\mathcal{Q})$, which is computationally tractable, as will be described in Section 4. The Pareto optimum of this problem is defined in the same way as in equation (3); we will henceforth denote the relaxed Pareto frontier by $\tilde{d}_\mathrm{P}$. The following relation between d_P and $\tilde{d}_\mathrm{P}$ holds:

Theorem 2 (Relation of d_P and $\tilde{d}_\mathrm{P}$). *Let $D = \max\{\mathrm{diam}\,\mathcal{S}, \mathrm{diam}\,\mathcal{Q}\}/\delta(1-\delta)$, for some $0<\delta<1$. Then, $\tilde{d}_\mathrm{P}(\mathcal{S},\mathcal{Q}) \le \left((1-\delta)^{-1}, \delta^{-2}\right)\cdot d_\mathrm{P}(\mathcal{S},\mathcal{Q})$, where the inequality is interpreted in the vector sense.*

The proof is based on the Chebyshev inequality and is not given here due to space limitations.

4 Computational framework

Practical computation of the Paretian similarity is performed on discretized objects. The surface $\mathcal{S}$ is represented as a triangular mesh, whose vertices constitute a finite r-sampling $\mathcal{S}_N = \{s_1, ..., s_N\}$. A point s on $\mathcal{S}$ is represented as a pair (t, u), where t is the index of the triangular face enclosing it, and u is the vector of barycentric coordinates with respect to the vertices of that triangle. The metric on $\mathcal{S}$ is discretized by numerically approximating the geodesic distances between the samples s_i on the triangular mesh, using the *fast marching method* (FMM) [22]. Geodesic distances between two arbitrary points on the mesh are interpolated from $d_\mathcal{S}(s_i, s_j)$'s using the three-point interpolation approach presented in [8]. The measure on $\mathcal{S}$ is discretized as $\{\mu_{\mathcal{S}_N}(s_1), ..., \mu_{\mathcal{S}_N}(s_N)\}$, assigning to each $s_i \in \mathcal{S}_N$ the area of the corresponding Voronoi cell.

Given two discretized surfaces $\mathcal{S}_N$ and $\mathcal{Q}_M$, we compute the relaxed Paretian similarity as the solution to

$$\min_{\substack{q_1',...,q_N' \in \mathcal{Q} \\ s_1',...,s_M' \in \mathcal{S} \\ m_{\mathcal{S}_N}(s_1),...,m_{\mathcal{S}_N}(s_N)\in[0,1] \\ m_{\mathcal{Q}_M}(q_1),...,m_{\mathcal{Q}_M}(q_M)\in[0,1]}} \quad \max \left\{ \begin{array}{l} \max\limits_{i,j} m_{\mathcal{S}_N}(s_i) m_{\mathcal{S}_N}(s_j) |d_\mathcal{S}(s_i,s_j) - d_\mathcal{Q}(q_i', q_k')| \\ \max\limits_{k,l} m_{\mathcal{Q}_M}(q_k) m_{\mathcal{Q}_M}(q_l) |d_\mathcal{Q}(q_k,q_l) - d_\mathcal{S}(s_k', s_l')| \\ \max\limits_{i,k} m_{\mathcal{S}_N}(s_i) m_{\mathcal{Q}_M}(q_k) |d_\mathcal{S}(s_i,s_k') - d_\mathcal{Q}(q_k, q_i')| \\ \max\limits_{i} D\,(1 - m_{\mathcal{S}_N}(s_i'))m_{\mathcal{S}_N}(s_i) \\ \max\limits_{k} D\,(1 - m_{\mathcal{Q}_M}(q_k'))m_{\mathcal{Q}_M}(q_k) \end{array} \right\}$$
$$\text{s.t.} \quad \sum m_{\mathcal{S}_N}(s_i)\mu_{\mathcal{S}_N}(s_i) \ge 1-\lambda$$
$$\sum m_{\mathcal{Q}_M}(q_k)\mu_{\mathcal{Q}_M}(q_k) \ge 1-\lambda, \qquad (6)$$

for a fixed set of values of λ, each λ giving a different point on the Pareto frontier $\tilde{d}_\mathrm{P}$. Here, $m_{\mathcal{S}_N}(s_i)$ and $m_{\mathcal{Q}_M}(q_k)$ denote the discretized membership functions computed on s_i and q_k, respectively. $m_{\mathcal{S}_N}(s)$ and $m_{\mathcal{Q}_M}(q)$ denote the interpolated weights for arbitrary points $s \in \mathcal{S}$ and $q \in \mathcal{Q}$. Note that the minimization over all mappings $f : \mathcal{S} \to \mathcal{Q}$ and $g : \mathcal{Q} \to \mathcal{S}$ is replaced by minimization over the images $q_i' = f(s_i)$ and $s_k' = g(q_k)$, in the spirit of multidimensional scaling.

The minimization problem (6) can be solved by alternatingly solving two smaller problems, namely the minimization of (6) with respect to $m_{\mathcal{S}}(s_i)$ and $m_{\mathcal{Q}}(q_k)$ for fixed s'_k and q'_i, which can be cast as the following constrained minimization problem

$$\min_{\substack{\epsilon\geq 0\\ m_{\mathcal{S}_N}(s_1),\ldots,m_{\mathcal{S}_N}(s_N)\in[0,1]\\ m_{\mathcal{Q}_M}(q_1),\ldots,m_{\mathcal{Q}_M}(q_M)\in[0,1]}} \epsilon \text{ s.t. } \begin{cases} m_{\mathcal{S}_N}(s_i)m_{\mathcal{S}_N}(s_j)|d_{\mathcal{S}}(s_i,s_j)-d_{\mathcal{Q}}(q'_i,q'_k)|\leq\epsilon \\ m_{\mathcal{Q}_M}(q_k)m_{\mathcal{Q}_M}(q_l)|d_{\mathcal{Q}}(q_k,q_l)-d_{\mathcal{S}}(s'_k,s'_l)|\leq\epsilon \\ m_{\mathcal{S}_N}(s_i)m_{\mathcal{Q}_M}(q_k)|d_{\mathcal{S}}(s_i,s'_k)-d_{\mathcal{Q}}(q_k,q'_i)|\leq\epsilon \\ D\,(1-m_{\mathcal{S}_N}(s'_i))m_{\mathcal{S}_N}(s_i)\leq\epsilon \\ D\,(1-m_{\mathcal{Q}_M}(q'_k))m_{\mathcal{Q}_M}(q_k)\leq\epsilon \\ \sum m_{\mathcal{S}_N}(s_i)\mu_{\mathcal{S}_N}(s_i)\geq 1-\lambda \\ \sum m_{\mathcal{Q}_M}(q_k)\mu_{\mathcal{Q}_M}(q_k)\geq 1-\lambda, \end{cases} \quad (7)$$

and the minimization of (6) with respect to s'_k and q'_i for fixed $m_{\mathcal{S}}(s_i)$ and $m_{\mathcal{Q}}(q_k)$, which can be formulated as

$$\min_{\substack{\epsilon\geq 0\\ q'_1,\ldots,q'_N\in\mathcal{Q}\\ s'_1,\ldots,s'_M\in\mathcal{S}}} \text{ s.t. } \begin{cases} m_{\mathcal{S}_N}(s_i)m_{\mathcal{S}_N}(s_j)|d_{\mathcal{S}}(s_i,s_j)-d_{\mathcal{Q}}(q'_i,q'_k)|\leq\epsilon \\ m_{\mathcal{Q}_M}(q_k)m_{\mathcal{Q}_M}(q_l)|d_{\mathcal{Q}}(q_k,q_l)-d_{\mathcal{S}}(s'_k,s'_l)|\leq\epsilon \\ m_{\mathcal{S}_N}(s_i)m_{\mathcal{Q}_M}(q_k)|d_{\mathcal{S}}(s_i,s'_k)-d_{\mathcal{Q}}(q_k,q'_i)|\leq\epsilon \\ D\,(1-m_{\mathcal{S}_N}(s'_i))m_{\mathcal{S}_N}(s_i)\leq\epsilon \\ D\,(1-m_{\mathcal{Q}_M}(q'_k))m_{\mathcal{Q}_M}(q_k)\leq\epsilon \end{cases} \quad (8)$$

and solved using the multi-resolution approach proposed in [11, 12, 10, 8] for the GMDS. Another, more efficient approach, is to solve a weighted L_2 approximation to (8) and use iterative re-weighting as a means of approximating the original L_∞ problem.

4.1 Sensitivity to noise

If the accuracy of geodesic distance measurement is δ (in FMM methods, δ is of order of the maximum edge length in the mesh), the accuracy of the Gromov-Hausdorff distance is bounded by 2δ. Using an L_2 criterion instead of the L_∞ is advantageous in the case of noise, since it is less influenced by outliers. Like all the approaches based on the analysis of intrinsic geometry, our method may be sensitive to *topological noise*, or in other words, noise in extrinsic geometry that results in different topology of the surface.

5 Results

We tested our method on a set of partially overlapping objects, created from the Elad-Kimmel database [16]. Five objects (dog, spider, giraffe, man and crocodile) were used; a full version of each object appeared, in addition to four different deformations of which parts were cropped, resulting in five instances per object (total of 25 objects, see Figure 3). The resulting objects were partially overlapping (Figure 2, top). The objects were represented as triangular meshes and comprised between 1500 to 3000 points. The geodesic distances were computed using FMM. Set-valued distances were computed between all the objects using 13 values of λ. We used a multiresolution iteratively re-weighted scheme described in Section 4. Six resolution levels with 50 points at the finest level were used. The algorithms were implemented in MATLAB; the computation of the Pareto frontier took about a minute on a standard Intel Pentium IV computer.

Figure 2 shows the Pareto frontier corresponding to the set-valued distance between two instances of the dog object. Shades of red represent the values of the membership functions. The overlapping regions in the two objects are clearly visible. Figure 1 shows the Pareto frontiers arising from partial comparison of different objects. One can observe that the dog-man and the dog-giraffe comparisons (red) result in curves above those obtained for the comparison of different instances of the dog (black). Figure 3 depicts the L_1-Salukwadze distance (with scaling factor $\alpha = 200$) between the objects, represented as Euclidean similarities. Clusters corresponding to different objects are clearly distinguishable. For comparison, we refer the reader to [8], where the computation of the Gromov-Hausdorff distance between the full versions of the same objects presented here is shown.

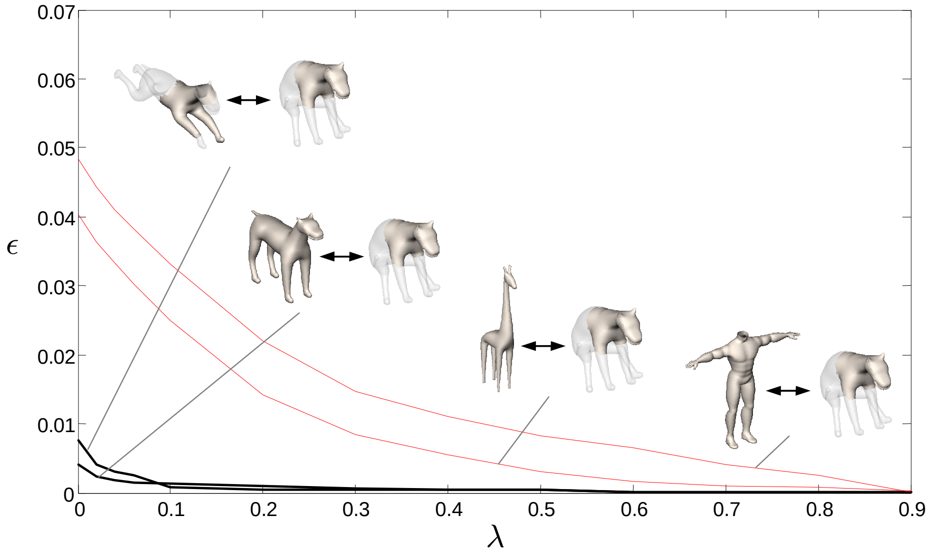

Fig. 1. Pareto similarity between different objects.

6 Conclusions

We presented a method for partial comparison of non-rigid objects. Our approach suggest quantifying partial similarity as a tradeoff between the intrinsic

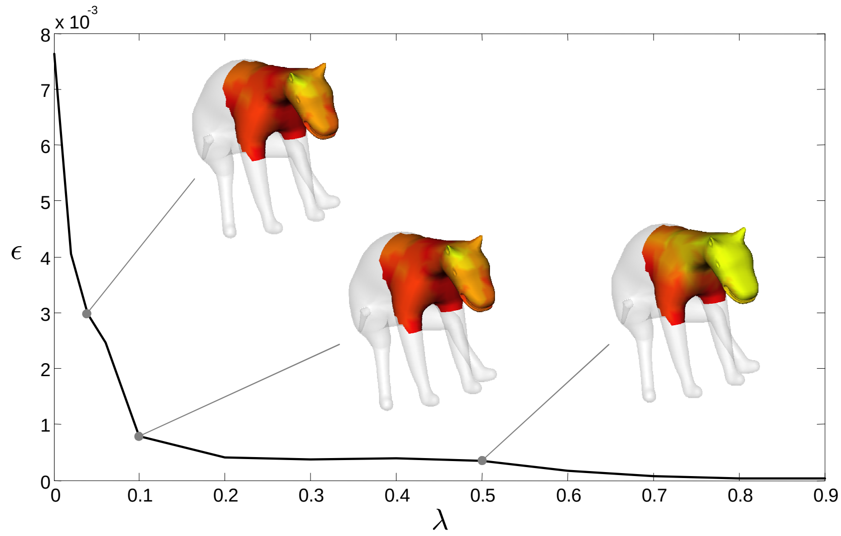

Fig. 2. Example of Paretian similarity. Shown is Pareto frontier corresponding to the set-valued distance between the dog objects. Colors encode the membership functions (red corresponding to 1).

geometric similarity and the area of a subset of the objects, using the formalism of Pareto optimality. Such a construction has a meaningful interpretation; the set-valued distances resulting from it have appealing theoretical and practical properties. For the efficient computation of our similarity criteria, we developed a numerical framework similar to the GMDS algorithm. Experimental results show that our method is able to recognize non-rigid objects even when large parts of them are missing or differ.

Acknowledgement

This research was supported in part by United State-Israel Binational Science Foundation grant No. 2004274, the Israeli Science Foundation (ISF) grant No. 738/04 and the Horowitz fund.

References

1. I. Borg and P. Groenen, *Modern multidimensional scaling - theory and applications*, Springer, 1997.
2. K. W. Bowyer, K. Chang, and P. Flynn, *A survey of 3D and multi-modal 3D+2D face recognition*, Dept. of Computer Science and Electrical Engineering Technical report, University of Notre Dame, January 2004.

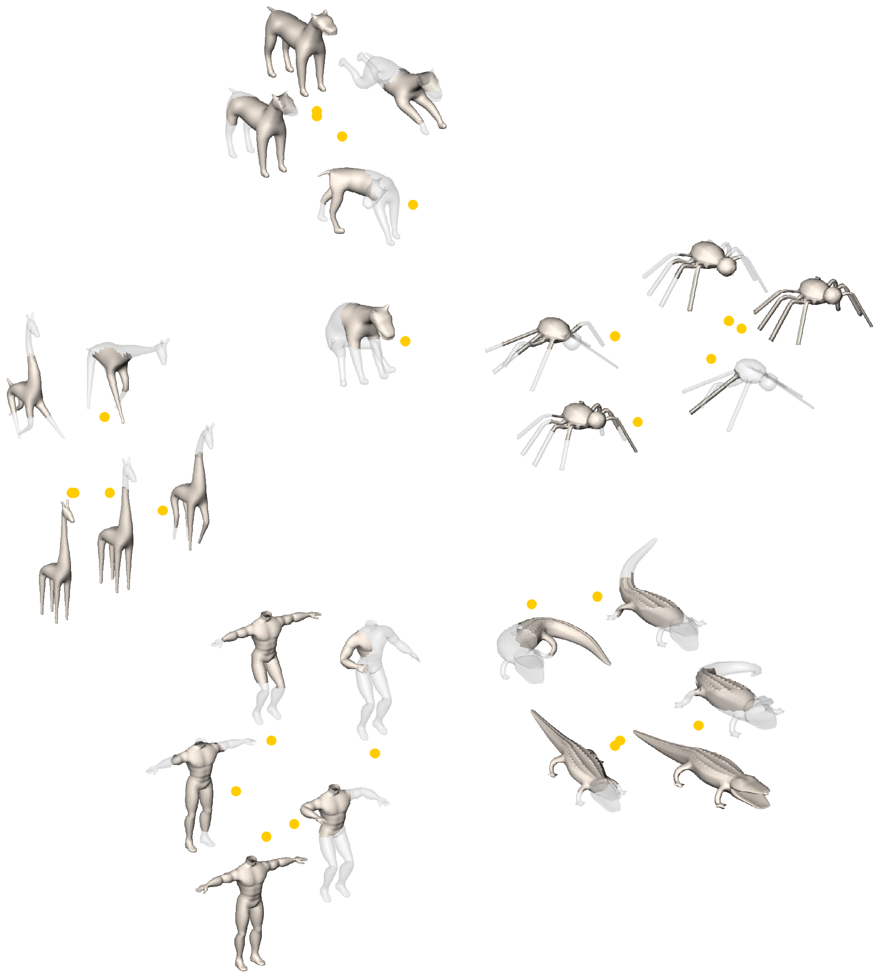

Fig. 3. L_1-Salukwadze distance between partially missing objects, represented as Euclidean similarities.

3. A. M. Bronstein, A. M. Bronstein, and R. Kimmel, *Calculus of non-rigid surfaces for geometry and texture manipulation*, IEEE Trans. Visualization and Computer Graphics (2006), in press.
4. ______, *Robust expression-invariant face recognition from partially missing data*, Proc. ECCV, 2006, pp. 396–408.
5. A. M. Bronstein, M. M. Bronstein, A. M. Bruckstein, and R. Kimmel, *Matching two-dimensional articulated shapes using generalized multidimensional scaling*, Proc. AMDO, 2006, pp. 48–57.
6. A. M. Bronstein, M. M. Bronstein, and R. Kimmel, *Expression-invariant 3D face recognition*, Proc. AVBPA, Lecture Notes on Computer Science, no. 2688, Springer, 2003, pp. 62–69.

7. ———, *Three-dimensional face recognition*, IJCV **64** (2005), no. 1, 5–30.
8. ———, *Efficient computation of isometry-invariant distances between surfaces*, SIAM Journal Scientific Computing **28** (2006), 1812–1836.
9. ———, *Face2face: an isometric model for facial animation*, Proc. AMDO, 2006, pp. 38–47.
10. ———, *Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching*, PNAS **103** (2006), no. 5, 1168–1172.
11. M. M. Bronstein, A. M. Bronstein, R. Kimmel, and I. Yavneh, *A multigrid approach for multidimensional scaling*, Copper Mountain Conf. Multigrid Methods, 2005.
12. ———, *Multigrid multidimensional scaling*, Numerical Linear Algebra with Applications **13** (2006), 149–171.
13. D. Burago, Y. Burago, and S. Ivanov, *A course in metric geometry*, Graduate studies in mathematics, vol. 33, American Mathematical Society, 2001.
14. G. Charpiat, O. Faugeras, and R. Keriven, *Approximations of shape metrics and application to shape warping and empirical shape statistics*, Found. Comput. Math. **5** (2005), no. 1, 1–58.
15. S. de Rooij and P. Vitanyi, *Approximating rate-distortion graphs of individual data: Experiments in lossy compression and denoising*, IEEE Trans. Information Theory (2006), submitted.
16. A. Elad and R. Kimmel, *Bending invariant representations for surfaces*, Proc. CVPR, 2001, pp. 168–174.
17. R. M. Everson and J. E. Fieldsend, *Multi-class ROC analysis from a multi-objective optimization perspective*, Pattern Recognition Letters **27** (2006), no. 8, 918–927.
18. M. Gromov, *Structures métriques pour les variétés riemanniennes*, Textes Mathématiques, no. 1, 1981.
19. D. Hoffman and W. Richards, *Visual cognition*, ch. Parts of recognition, MIT Press, Cambridge, 1984.
20. B.W. Hong, E. Prados, S. Soatto, and L. Vese, *Shape representation based on integral kernels: Application to image matching and segmentation*, Proc. CVPR, 2006, pp. 833–840.
21. D. Jacobs, D. Weinshall, and Y. Gdalyahu, *Class representation and image retrieval with non-metric distances*, IEEE Trans. PAMI **22** (2000), 583–600.
22. R. Kimmel and J. A. Sethian, *Computing geodesic on manifolds*, PNAS, vol. 95, 1998, pp. 8431–8435.
23. L.J. Latecki, R. Lakaemper, and D. Wolter, *Optimal partial shape similarity*, Image and Vision Computing **23** (2005), 227–236.
24. H. Ling and D. Jacobs, *Using the inner-distance for classification of articulated shapes*, Proc. CVPR, 2005.
25. F. Mémoli and G. Sapiro, *A theoretical and computational framework for isometry invariant recognition of point cloud data*, Foundations of Computational Mathematics **5** (2005), no. 3, 313–347.
26. A. Pentland, *Recognition by parts*, Proc. ICCV, 1987, pp. 612–620.
27. D. Geiger R. Basri, L. Costa and D. Jacobs, *Determining the similarity of deformable shapes*, Vision Research **38** (1998), 2365–2385.
28. M. Reuter, F.-E. Wolter, and N. Peinecke, *Laplace-Beltrami spectra as shape-DNA of surfaces and solids*, Computer-Aided Design **38** (2006), 342–366.
29. M. E. Salukwadze, *Vector-valued optimization problems in control theory*, Academic Press, 1979.
30. J. Zhang, R. Collins, and Y. Liu, *Representation and matching of articulated shapes*, Proc. CVPR, vol. 2, June 2004, pp. 342 – 349.