

A Condition Number for Non-Rigid Shape Matching

Maks Ovsjanikov Qi-Xing Huang Leonidas Guibas
Stanford University

Abstract

Despite the large amount of work devoted in recent years to the problem of non-rigid shape matching, practical methods that can successfully be used for arbitrary pairs of shapes remain elusive. In this paper, we study the hardness of the problem of shape matching, and introduce the notion of the shape condition number, which captures the intuition that some shapes are inherently more difficult to match against than others. In particular, we make a connection between the symmetry of a given shape and the stability of any method used to match it while optimizing a given distortion measure. We analyze two commonly used classes of methods in deformable shape matching, and show that the stability of both types of techniques can be captured by the appropriate notion of a condition number. We also provide a practical way to estimate the shape condition number and show how it can be used to guide the selection of landmark correspondences between shapes. Thus we shed some light on the reasons why general shape matching remains difficult and provide a way to detect and mitigate such difficulties in practice.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms

1. Introduction

Shape matching is one of the most basic operations in digital geometry processing, with applications ranging from medical image analysis [MV98] to shape interpolation [KMP07], to reconstruction [LPC*00], and to information (e.g. texture) transfer [KS04]. In the most general setting, the goal of a shape matching method is to recover an alignment between two shapes, which can be represented as either a deformation of one of the shapes or a correspondence between points on two given shapes. Although shape matching can often be effortlessly performed by humans, practical and robust computational methods that can be used for a wide class of shapes and their deformations remain elusive.

A major challenge in designing a successful shape matching method is to define an informative measure by which potential alignments can be scored and compared. In the majority of cases, such a score is represented through an energy function which assigns a single value to each potential solution of the matching problem, and the goal is to find the alignment minimizing this score, locally or globally. The primary objective of unsupervised shape matching is to formulate an energy function entirely in terms of the shape geometry, without prior knowledge on shape functionality, part decomposition or semantics. This constraint lies at the heart of the difficulty of non-rigid shape matching, and as we will show




Figure 1: Optimal placement of a landmark to improve stability of shape matching using our method. Red to green means higher to lower improvement in stability. Note that the best landmarks are in the most featureless areas of the shape.

in this paper, prior information is often *required* to obtain a single, reliable solution.

When the matching problem is phrased in terms of finding the optimal correspondence between points on the two shapes, a commonly used strategy for defining an energy function is to measure how well a correspondence preserves distances between pairs of points on each shape. In the context of non-rigid matching, distance preservation is often interpreted in terms of geodesic distances, which translates to recovering approximate intrinsic isometries — a problem

that has recently received a lot of attention, e.g. see [MS05, BBK06, TBW*09, LF09, OMMG10] among many others.

Alternatively, when the matching algorithm seeks a lower-dimensional *deformation* of a source shape onto the target shape, and the deformed points may not lie on but only near the target, the energy function is often formulated through a trade-off between proximity to the target shape and the local stretching introduced by the deformation, e.g. [LSP08, HAWG08, ZSCO*08, WAO*09].

Note that in both formulations the solution space is associated with an energy landscape and the goal of the matching algorithm is to find the global (or often local) minimum of this energy. It is the case, however, that often there exist other solutions with very similar energy values to the optimal one, but which lie very far in the solution space. The simplest example of this scenario is matching symmetric shapes, which admit at least two equally good maps. In the presence of noise, which map will be chosen as optimal can be arbitrary, and prior information such as shape semantics is necessary to pick the single "correct" solution. Even without perfect symmetry, however, small perturbations of one of the shapes can cause the optimal solution to change wildly.

In this paper, we study the stability properties of non-rigid shape matching with respect to a given energy function. Our goal is to formalize the intuition that some shapes lead to more "difficult" energy landscapes, and this difficulty can be captured by measuring the stability of the "best" solution under small perturbations. For example, a highly detailed, non-symmetric shape may have a unique global minimum of the deformation energy when matched even against a perturbed version of itself, while a near-symmetric shape may have several equally good solutions, without a canonical way to choose between them.

To capture this intuition we introduce a stability measure, called the *shape condition number*, which measures the worst-case behavior of how much the optimal solution can change under perturbations of the input source shape or energy landscape. We then relate the condition number to a simpler notion called the *shape symmetry score* which measures how far a shape is from symmetry. We analyze the symmetry score for both correspondence-based and deformation based shape matching and show how appropriate notions can be defined and estimated in practice through an eigenanalysis of certain matrices. Finally, we use these measures to derive methods for computing landmarks on a shape that explicitly aim to improve the stability of shape matching. For example, Figure 1 shows the improvement in stability of matching after adding a landmark correspondence for different points on the cat shape. Note that the best landmarks can often lie in the most featureless parts of the shape.

Although the notion that symmetric shapes lead to ambiguities in matching and are therefore more difficult to match is intuitive and has been in the geometry processing folklore, we believe that we are the first to provide a rigorous

formulation of this idea via the concepts of the shape condition number and symmetry score and to study their properties. Furthermore, our stability-enhancing landmark selection strategy leads to landmark sequences with a distinctly different flavor and better behavior from those that have appeared earlier.

2. Related Work

Although a wealth of methods exists for non-rigid matching, the questions of stability as well as the hardness of shape matching, while observed in many settings, have not received a great deal of attention.

In the context of rigid matching, [MGPG04] studies the "funnel of convergence" of the point-to-plane Iterative Closest Point (ICP) algorithm of Chen and Medioni [CM92], which is the set of initial shape configurations that lead to convergence of ICP. The authors also propose a method that is shown to have a more stable funnel of convergence for a set of shapes. In a related work, Gelfand et al. [GIRL03] study the problem of optimal placement of landmark correspondences for ICP and show that greater stability can be achieved if the landmarks are chosen to decrease the slippage energy. Interestingly, the slippage energy is closely related to the symmetry structure of the shape, as developed and exploited in [GG04].

In non-rigid shape matching, several techniques have been proposed for finding landmark correspondences which can be used as anchors to improve stability of matching.

Memoli and Sapiro [MS05] use Farthest Point Sampling (FPS) on the geodesic distance function to find a covering of a shape by landmarks which can be used to approximate the Gromov-Hausdorff distance between a pair of shapes. Similarly, Bronstein et al. [BBK06] and Raviv et al. [RBBK07] use FPS to reduce the complexity of finding a low distortion map between a large set of points. One of the drawbacks of FPS is its reliance of the randomness to compute the first landmark point. To improve repeatability Zhang et al. [ZSCO*08] use the extremal points of the geodesic distance field as feature landmarks for correspondence. Sun et al. [SOG09] and Ovsjanikov et al. [OMMG10] use extremal points of the heat kernel function to find landmark points for correspondence. These landmarks also show a large degree of repeatability on a standard dataset [BBB*10] (see also references in [BBB*10] for a list of evaluated feature detection methods). In a similar spirit, Ruggeri et al. [RPSS10] use the extremal points of the eigenfunctions of the Laplace-Beltrami operator to find landmark correspondences. Note, however, that while landmark repeatability of several feature detectors was evaluated by Bronstein et al. [BBB*10], in most of the existing methods there is no explicit connection between the performance of landmark selection and the stability of shape matching. Note also that recently, Sun et al. [SCF10] have introduced a Gromov-Hausdorff stable

modification of geodesic distances and used it to define a notion of shape distance that they show more stable and discriminative than geodesic distortion on a benchmark set.

Perhaps the most closely related to ours is recent work by Tevs and colleagues [TBW*11] who propose a guided landmark selection process for shape matching, which adds landmark points according to the amount of information contained in them, as measured by the entropy of the appropriate probability distribution. While entropy is a useful measure of information gain, it is still not clear how to connect decrease in entropy with the stability of shape matching. Moreover, while Tevs et al. provide some practical evaluation of performance under perturbation of distances, in the analysis and motivation the authors assume that the two shapes are exactly isometric, whereas our goal is precisely to study the stability of shape matching under shape perturbations.


We also note that measuring the difficulty of matching through self-similarity has been considered before in the context of string matching, in the seminal works of Knuth, Morris, Pratt [KMP77] and Boyer, Moore [BM77].

3. Motivation

Before we introduce the notion of the shape condition number, we provide some motivation as to why such a notion is useful. In numerical stability theory (see e.g. [TB97]), a problem is called *well-conditioned* if a small perturbation in the input data leads to a small perturbation in the result. For example, a linear system $Ax = b$ is well conditioned if given $A(x + \Delta x) = b + \Delta b$, small Δb leads to a small Δx for any b . In this context, the (relative) matrix condition number $\kappa(A) = \sup_{b, \Delta b} \frac{\|\Delta x\|}{\|x\|} / \frac{\|\Delta b\|}{\|b\|} = \frac{\sigma_{\max}}{\sigma_{\min}}$, where σ_{\max} and σ_{\min} are the largest and smallest singular values of A respectively, or, equivalently, if A is real and symmetric, absolute values of largest and smallest eigenvalues. A linear system is well-conditioned if $\kappa(A)$ is close to 1, and ill-conditioned if $\kappa(A)$ is large. In particular, if A is rank-deficient, $\kappa(A) = \infty$. Intuitively, $\kappa(A)$ measures the difficulty of solving a linear system with A , *regardless* of the method used.

Our main goal is to use a similar notion for non-rigid shape matching. To illustrate this problem, consider two simple shapes in Figure 2a. They represent nearly perfect star-shapes with a small modification at one of the ends (marked in red). Thus, in the noise-free case, there exists a unique map T_{opt} that minimizes distortion of Euclidean distances between pairs of points. However, when a small amount of noise is added to Y , the optimal map can change drastically, Figure 2b. Thus, even if a method obtains the “correct” map in the noiseless case, this solution is not reliable since there exist other maps with very similar distortion, but which look very different. Note that in this case, and, as we show below, in general, instability is closely related to shape symmetry.

In order to quantify and measure the stability of shape



(a) Optimal map for the noise-free matching problem. (b) Optimal map after a small amount of noise.

Figure 2: For a pair of almost perfectly symmetric shapes (deviation marked in red), although the least-distortion map T_{opt} is unique (a), it can change dramatically with a small amount of noise (b).

matching we will first need to define a way to evaluate perturbations of the input shapes, as well as the induced changes in the optimal matching solution. As we noted in the introduction, the solution to a shape matching problem between shapes X and Y is either a map $T_{\text{opt}} : X \rightarrow Y$, which optimizes some measure of distortion $D_{X,Y}(T)$, or a set of deformation parameters P_{opt} , such that when X is deformed according to P_{opt} , it lies as close as possible to Y . In this setting again, P_{opt} are commonly chosen to minimize an energy function $E_{X,Y}(P)$. Therefore, to define a condition number, we need a way to measure stability of T_{opt} and P_{opt} . In the following sections we outline several possibilities for correspondence-based and deformation-based shape matching methods.

4. Correspondence-Based Matching

In this section we define the condition number for correspondence-based matching. Given two shapes X and Y , the output of a matching is a map $T_{\text{opt}} : X \rightarrow Y$ that optimizes some distortion measure $D_{X,Y}(T)$ (see e.g. [MS05, BBK06, TBW*09, LF09] and related techniques). We would like to say that a shape X is well-conditioned if a small perturbation in Y leads to a small perturbation in T_{opt} for all Y . One of the difficulties in implementing this idea is that if Y is symmetric then there may be several optimal maps T , regardless of X . To avoid such situations we will assume that Y is a perturbation of X and will measure how far T_{opt} is from the identity.

To formalize this notion, we will treat all shapes as compact metric spaces. Thus, we will represent a shape X as a collection of point-to-point distances (e.g. geodesic or Euclidean). For simplicity, we will also assume that the shapes consist of a finite number of points, although similar statements can also be made for continuous metric spaces. Finally, we would like all shapes to be scaled so that e.g. the diameter is always 1. This is only to ensure that something like a “small distance” makes sense. Then, given two shapes X and Y , both containing N points, let:

$$T_{\text{opt}}(X, Y) = \underset{T: X \rightarrow Y}{\operatorname{argmin}} D_{X,Y}(T), \quad (1)$$

where $D_{X,Y}(T)$ is some measure of distortion of distances introduced by T , and the minimization done is over all bijections. Several measures of distortion, have been proposed in the literature, but perhaps the most common one is the L^p

norm on the difference of pairwise distances:

$$D_{X,Y}^p(T) = \left(\sum_{x,x' \in X} (d_X(x,x') - d_Y(T(x),T(x')))^p \right)^{1/p}, \quad (2)$$

where p is commonly set to be 1, 2 or ∞ . In the latter case: $D_{X,Y}^\infty(T) = \max_{x,x' \in X} |d_X(x,x') - d_Y(T(x),T(x'))|$. Note that for finite metric spaces with the same number of points, the Gromov-Hausdorff distance $d^{GH}(X,Y) = \frac{1}{2} D_{X,Y}^\infty(T_{\text{opt}})$, where T_{opt} is the bijection that minimizes $D_{X,Y}^\infty$ [Mém07]. Note also that finding a map T that minimizes D^p for any fixed p and arbitrary finite metric spaces is NP-hard, as it includes the metric Traveling Salesman Problem as a special case. This can be shown by a simple modification of the proof in [BCPP98].

Given a distortion measure $D_{X,Y}(T)$, we define $T_{\text{opt}}(X,Y)$ as the map between X and Y that minimizes $D_{X,Y}(T)$. There is an important subtlety of what happens when multiple T_{opt} exist with the same distortion. We will treat this question separately, when we define the condition number below.

In order to quantify the distortion of $T_{\text{opt}}(X,Y)$ under perturbations of Y , we will also need a way to measure distances between maps T_1 and $T_2 : X \rightarrow Y$. Here we use a similar L^p norm formulation:

$$d^p(T_1, T_2) = \left(\sum_{x \in X} (d_Y(T_1(x), T_2(x)))^p \right)^{1/p} \quad (3)$$

Now suppose \tilde{X} is a deformed version of X . For example, \tilde{X} is the information we have about X due to acquisition and representation error. Since \tilde{X} is a deformed version of X , we will assume that there is a natural bijection $\tilde{I}d : X \rightarrow \tilde{X}$. In other words, we can think of \tilde{X} as a pair $(\tilde{X}, \tilde{I}d)$, where $\tilde{I}d$ is a canonical map from X to \tilde{X} . When X is represented through a matrix of distances, we can think of \tilde{X} as a matrix of the same size with $\tilde{I}d$ given as the natural identity correspondence between the rows and columns.

One way to measure the distortion between \tilde{X} and X is via the Gromov-Hausdorff distance. However, our task is made easier since \tilde{X} is a deformation of X and we have the canonical map $\tilde{I}d$. In this case, we can simply use:

$$d(X, \tilde{X}) = D_{X, \tilde{X}}(\tilde{I}d),$$

where $D_{X, \tilde{X}}(\tilde{I}d)$ is one of the distortion measures above.

Given these measures, we define the condition number of X :

$$\kappa(X) = \sup_{\tilde{X}} \frac{d(\tilde{I}d, T_{\text{opt}}(X, \tilde{X}))}{d(X, \tilde{X})}. \quad (4)$$

Note that if multiple T_{opt} exist with the same distortion, we take the one that leads to the *highest* $\kappa(X)$.

We call $\kappa(X)$ the condition number of X . Intuitively, the condition number will be large if a small change in X can trigger a large change in the optimal map. In particular, the condition number measures the difficulty of finding the canonical

identity map under perturbations of X . If the identity map remains optimal after a perturbation of X , then $\kappa(X)$ will be small, and conversely, if a small perturbation of X alters the optimal map $T_{\text{opt}}(X, \tilde{X})$ and takes it very far away from the identity, then $\kappa(X)$ will be large.

Discussion Note that the condition number $\kappa(X)$ captures the worst-case behavior of matching a shape to its deformed copy. Given X and \tilde{X} , any algorithm that tries to optimize $D_{X, \tilde{X}}(T)$ is susceptible to uncertainty in the knowledge of \tilde{X} . Specifically, if all that is known about X and \tilde{X} is pairwise distances (or quantities derived from them such as the Laplace-Beltrami operator), then any uncertainty in the distances will result in uncertainty in the optimal map found by the algorithm. For a shape X , the condition number measures the worst case behavior in terms of \tilde{X} .

Note that in many cases some information, such as landmark correspondences is known, or could potentially be specified by the user. Such external information could improve the stability of isometric shape matching. We can incorporate this information into the definition of $\kappa(X)$ by modifying the definition of the optimal map T_{opt} . Rather than optimizing over all maps in Equation 1, we can restrict the choice to only those maps that satisfy the known constraints. We exploit this fact when we use our definition for computing landmark correspondences in Section 6.

4.1. Condition number and Symmetry

One disadvantage of the notion of condition number introduced above is the optimization over all possible shapes \tilde{X} , making it hard to estimate in practice. Intuitively, however, the difficulty of matching a shape must be already encoded in the shape itself. For example, if $\kappa(X) = \infty$, this means that the shape X admits a non-identity map with 0 distortion. In other words, X is symmetric. Thus $\kappa(X) = \infty$ if and only if X has a perfect symmetry.

To capture this intuition, we introduce another quantity, which is closely related to the shape condition number, that we call the *global symmetry score*, and which will allow us to relate the shape condition number to its symmetry structure. Given a compact shape X , we define its *global symmetry score* $\sigma(X)$ as:

$$\sigma(X) = \sup_{T: X \rightarrow X} \frac{d(T, Id)}{D_{X,X}(T)}, \quad (5)$$

where T again varies over all bijections $X \rightarrow X$. Intuitively, $\sigma(X)$ is attained for the self-map T that is as far as possible from the identity while introducing the least possible distortion. Note that if X is perfectly symmetric and admits a non-identity map with zero distortion, then $\sigma(X) = \infty$.

The main result of this section is the following Theorem, which relates $\sigma(X)$ with $\kappa(X)$:

Theorem 4.1 For any finite shape X and the L^p distortion measures introduced above, $\sigma(X) \leq \kappa(X) \leq 2\sigma(X) + 1$.

Proof: See Appendix.

Moreover, the following Lemma shows that when X and \tilde{X} are isometric, then $\sigma(X) = \kappa(X)$.

Lemma 4.2 A necessary and sufficient condition for $\kappa(X) = \sigma(X)$, is that the condition number $\kappa(X)$ is achieved by a shape \tilde{X} that is isometric to X , i.e.: $D_{X,\tilde{X}}(T_{\text{opt}}(X,\tilde{X})) = 0$.

Proof: See Appendix.

The two results above show that the difficulty of shape matching, when measured by the stability of the optimal solution is intimately related to the symmetry structure of the shape. This is also intuitively clear since a highly symmetric shape will have several equally good maps when mapped to another shape, and conversely a highly unstable shape must have near-isometric self-maps that are far from identity.

4.2. Computing The Global Symmetry Score

Given a shape X , represented as a symmetric matrix of pairwise distances between points, our goal is to estimate its global symmetry score $\sigma(X)$. Moreover, we will assume that both $D_{X,X}(T)$ and $d_X(T,Id)$ are given by the L^2 norm as defined in Equations 2 and 3 above.

To estimate $\sigma(X)$, we use a construction similar to the one by Leordeanu and Hebert [LH05]. Namely, for a shape X represented by a finite collection of N points, let A be a matrix of size $N^2 \times N^2$, where rows and columns represent pairs of points, and $A((x,x'),(y,y')) = (d_X(x,x') - d_X(y,y'))^2$. Moreover, let B be a diagonal matrix of the same size as A , and $B((x,x')) = d_X(x,x')^2$. Then, if λ_0 and λ_{\max} are the smallest and largest eigenvalues of the generalized eigenvalue problem $Ax = \lambda Bx$, we have: $\frac{1}{\sqrt{\lambda_{\max}}} \leq \sigma(X) \leq \frac{1}{\sqrt{\lambda_0}}$.

To see this, note that every N^2 -size vector x can be interpreted as a set of correspondences, by representing it in matrix form, taking blocks of N entries from x for each row. In particular, if x is a binary vector and in its matrix representation every row and every column contain exactly one 1, then x represents a bijection. Thus, using the Rayleigh quotient:

$$\begin{aligned} \frac{1}{\lambda_0} &= \sup_x \frac{x^T Bx}{x^T Ax} \geq \sup_{x \in \text{bijections}} \frac{x^T Bx}{x^T Ax} \\ &= \sup_{T: X \rightarrow X} \frac{\sum d(x, T(x))^2}{\sum (d(x, x') - d(T(x), T(x')))^2} \\ &= \sup_{T: X \rightarrow X} \frac{d^2(T, Id)}{D_{X,X}^2(T)} = \sigma^2(X). \end{aligned}$$

The lower bound is established similarly. One consequence of this result is that using Theorem 4.1 we can bound the condition number of X simply as $\kappa(X) \leq 2 \frac{1}{\sqrt{\lambda_0}} + 1$. Such a bound is useful when e.g. selecting a set of landmark correspondences that would reduce the condition number. Note that a similar algorithm for symmetry detection was recently

proposed by Chertok and Keller [CK10], although importantly the authors did not use the normalizing matrix B .

The lower bound, on the other hand, albeit much looser, is useful since it bounds the stability of any shape matching algorithm that tries to match X by optimizing a distortion measure D , while using only intrinsic quantities (i.e. distances and their derived quantities, such as the Laplace-Beltrami operator) without prior knowledge on the alignment.

Let us note that while the eigenvalue upper bound $\sigma(X) \leq 1/\sqrt{\lambda_0}$ is conceptually useful, in practice we observed that it is often too loose because it substitutes vectors representing bijections with arbitrary N^2 -vectors. Thus, we use a quasi-convex optimization method to estimate $\sigma(X)$ as follows: given matrices A and B as above, we use binary search to iteratively check whether $x^T Bx - \lambda x^T Ax \leq 0$ for a fixed λ is feasible or not. To check the feasibility, we use the linear programming relaxation technique which is frequently used in solving Markov Random Fields [WJW05]. The basic idea is to let $Y = xx^T$ such that the constraint $x^T Bx - \lambda x^T Ax \leq 0$ is linear in Y . Then we relax $Y = xx^T$ by the marginalization constraints. We refer the reader to [WJW05] for exact details of this approach.

5. Condition number for deformation-based matching

As mentioned in the introduction, another class of methods for matching two shapes X and Y aims at deforming X to align it as well as possible with Y under some deformation penalty e.g. [LSP08, HAWG08, ZSCO*08, WAO*09]. In this formulation, the output of a shape matching method is a set of deformation parameters P that minimize some energy $E_{X,Y}(P)$. For example, P may consist of $3n$ variables, each corresponding to a displacement of a point of X along one of the 3 coordinate axes. The goal of the matching algorithm is then to find a displacement P that would minimize the energy $E_{X,Y}(P)$. Note that in this scenario the majority of methods use numerical optimization techniques to find a local minimum of $E_{X,Y}(\cdot)$.

To define the condition number and the symmetry score in this setting, we will need a way to measure distances in the deformation parameter space as well as distances between shapes. Here we will assume that P can be represented as a vector of deformation parameters, and that when a shape X is matched against itself $E_{X,X}(\mathbf{0}) = 0$ is a global minimum of the energy function. In other words, we assume that for a non-deformed version of X , a vector of zeros leads to zero distortion and $E_{X,X}(P) \geq 0$ for any P . We define the condition number, simply as:

$$\kappa(X) = \sup_{\tilde{X}} \frac{d(P_{\text{opt}}, \mathbf{0})}{d(X, \tilde{X})},$$

where $P_{\text{opt}} = \text{argmin} E_{X,\tilde{X}}(P)$, and if several deformations exist that lead to the same error, we pick the one that maximizes $\kappa(X)$. The exact nature of the distance $d(X, \tilde{X})$ is not

crucial in our setting. However, we will assume that it has the following form: if $\tilde{X} = X(P)$, i.e. if \tilde{X} is obtained by deforming X according to P , then $d(X, \tilde{X}) \leq \sqrt{E_{X,X}(P)}$. This is motivated by the fact that the deformation energy is usually the sum of the squares of both pairwise distances between points and some additional terms related to deviation from rigidity. Moreover, we will assume that if $\tilde{X} = X(P)$ then P is among the global minima of $E_{X,\tilde{X}}(\cdot)$.

To define the symmetry score, we simply use:

$$\sigma(X) = \sup_P \frac{d(P, \mathbf{0})}{\sqrt{E_{X,X}(P)}}. \quad (6)$$

Similarly to above, we have:

Lemma 5.1 For any shape X , $\sigma(X) \leq \kappa(X)$.

Proof Under the above assumptions, the proof is straightforward: given deformation parameters P , construct a shape $\tilde{X} = X(P)$. Then note that P must be among the global minima of $E_{X,\tilde{X}}$ and moreover $d(X, \tilde{X}) \geq \sqrt{E_{X,X}(P)}$ by assumption. Therefore, $\frac{d(P, \mathbf{0})}{\sqrt{E(P)}} \leq \frac{d(P, \mathbf{0})}{d(X, \tilde{X})} \leq \kappa(X)$. Since this holds for any P , the result follows. \square

Note that the symmetry score $\sigma(X)$ measures the worst-case behavior for any set of deformation parameters P . However, in many settings, methods that use a deformation penalty aim at finding local minima of the deformation energy. In this case, to measure the stability of a *local* minimum, we need to restrict our attention to only infinitesimal deformations:

$$\sigma_{\text{inf}} = \lim_{\epsilon \rightarrow 0} \sup_{d(P, \mathbf{0}) < \epsilon} \frac{d(P, \mathbf{0})}{\sqrt{E_{X,X}(P)}}.$$

Note that unlike the global version of the symmetry score, its local counterpart can be thought as quantifying the certainty that a particular solution is indeed a *local* minimum, and is closely related to the slippage analysis of Gelfand et al. [GG04]. The two measures are complimentary and can be used together to measure the difficulty of shape matching.

5.1. Computation of Local Symmetry Score

Interestingly, if $d(P, \mathbf{0})$ is measured as the L^2 norm of the vector P , i.e. $d(P, \mathbf{0}) = \|P\|_2 = (\sum P_i^2)^{1/2}$ just like the computation of the global symmetry score for correspondence-based matching in Section 4.2, the computation of σ_{inf} reduces to an eigenvalue problem. Namely, $\sigma_{\text{inf}} = 2/\sqrt{\lambda_{\min}(H)}$, where H is the Hessian matrix of the deformation energy $E_{X,X}$ evaluated at 0.

To see this simply note that locally: $E_{X,X}(P) = E_{X,X}(\mathbf{0}) + \nabla E_{X,X}(\mathbf{0}) \cdot P + \frac{1}{2} P^T H P + o(\|P\|^2)$. Since $E(\mathbf{0}) = 0$, and $E(P) \geq 0$ by assumption, meaning $\nabla E_{X,X}(\mathbf{0}) = 0$, we get, up to third order: $E_{X,X}(P) = \frac{1}{2} P^T H P + o(\|P\|^2)$. Therefore:

$$\sigma_{\text{inf}} = \sup_P \frac{2\|P\|}{\sqrt{P^T H P}} = \frac{2}{\sqrt{\lambda_{\min}(H)}}.$$

Relation to prior work: Note that Gelfand et al. [GIRL03] used a similar notion as a measure of stability of the ICP algorithm. In their work, the deformation parameters $P = (c \tilde{c})$

was a 6-dimensional vector corresponding to a linearization of translation and rotation, and the deformation energy was the point-to-plane ICP metric [CM92]. Gelfand et al. then used the condition number of the Hessian of the deformation energy as a measure of stability of ICP. This number is simply the ratio of the largest and smallest eigenvalues. The difference of their definition from ours simply comes from the fact that our notion of stability is absolute rather than relative, which is reasonable since we are only considering the solutions around identity (a relative measure would also be unclear for correspondence-based matching since it would require us to define a norm of a map). Note, however, that qualitatively the two notions behave similarly, and e.g. tend to ∞ for rank-deficient matrices, which correspond to self-symmetric (in the rigid case, slippable [GG04]) shapes. Note also that spectral analysis of a deformation energy's Hessian has been used for shape analysis in the works of of Hildebrandt [HSvTP10] and Huang [HWAG09] among others.

6. Landmark selection

One of the main applications for the notions of stability introduced above is finding ways in which stability can be improved in practice. Note that the condition number gives a lower bound on the stability of *any* method used to match a given shape using only its metric structure. However, in many cases additional information may exist that can be used during matching. The most common strategy of introducing information for shape matching is by using landmark correspondences, which can be either specified by the user or precomputed separately. Thus, when performing shape matching the set of allowable shape deformations or maps is restricted to only those that preserve known landmark correspondences. Note that even a few landmark correspondence can often greatly constraint the space of possible solutions and potentially improve stability significantly [LF09, OMMG10]. However, with the notable exception of the work of Tevs et al. [TBW*11] the optimal placement of landmark correspondences for shape matching has received little attention.

Our stability measures and especially the symmetry score provide a general way to compute landmark points on a shape in order to improve stability of shape matching. In the definition of the symmetry score Equations 5 and 6, adding landmark correspondences would mean restricting the search to only those maps T and those deformation parameters P that *fix* the known landmark points. Note that the symmetry score can only become smaller with addition of constraints on the space of solutions.

In practice when computing the symmetry score using the methods in Section 4.2 or Section 5.1 (in the case where P represents displacements of points) adding landmark correspondences simply translates into removing the corresponding rows and columns from the matrices before performing the eigen-decomposition.




Figure 3: Symmetry score $\sigma(X)$ computed on four near-symmetric shapes. Note that the fish has an approximate reflectional symmetry (in the undeformed state), as does the hand.

To find the optimal landmark points, we must therefore find those points such that after removing the corresponding rows and columns the symmetry score decreases the most, or, equivalently, the smallest eigenvalue increases the most.

To implement this idea, we use a greedy sampling strategy that evaluates for each point on the shape the reduction of the symmetry score if this point is fixed. We then sample the points that lead to the greatest reduction in the symmetry score. In the following section we illustrate this general method on two types of techniques for shape matching, which are most useful for global and local search respectively. In all of the examples below we compute landmark samples using the isometric correspondence-based matching using the L^2 norm as described in Section 4.2. We call this method **Condition Number Based Sampling (CNBS)**. We emphasize that CNBS is not a feature detection method and its goal is not to find the most repeatable landmark points that are easy to identify and match on copies of the shape. Rather, our goal is to find those points such that if their correspondences are known then the shape matching problem becomes easiest. This is especially useful when the correspondences are provided by a user, and our goal is to minimize manual effort, before the matching problem is sufficiently easy to be solved automatically. Note also that by introducing landmark points we get a more refined measure of shape matching difficulty. Indeed, by concatenating the scores after adding landmark points, we obtain a *vector* of symmetry scores for each shape, which could be used to e.g. distinguish between symmetric shapes such as a human and a sphere which initially have infinite symmetry scores.

7. Results

We first illustrate the relevance of the notion of the symmetry score $\sigma(X)$ for near-symmetric shapes. Figure 3 shows the computed symmetry scores on four shapes with progressively more pronounced symmetries, ending with the horse shape which is perfectly symmetric. Here and in all of the experiments below we first sub-sampled the shapes using Farthest Point Sampling (FPS) to only include 400 points, and computed the symmetry score using the L^2 norm of the geodesic distortion as described in Section 4.2. Note that although in theory our relaxation provides only an upper




Figure 4: Heat maps showing the improvement in the symmetry score after fixing landmarks at different points on the shape. Blue to red corresponds to little to high improvement in the symmetry score for the David (a) and Dog (b) models in the TOSCA dataset.

bound on the symmetry score, in practice it follows very closely the near-symmetry structure of the shape. For a perfectly symmetric horse shape the symmetry score $\sigma(X) = \infty$ as expected, while for near symmetric hand (which has an approximate front-back as well as reflectional symmetries), fish (which represents a bending of a perfectly symmetric shape) and foot shapes, the symmetry score decreases.

In the rest of this section we illustrate the performance of our Condition Number Based Sampling (CNBS) for landmark selection. As described in Section 6 our method is greedy and selects landmarks by considering the decrease in the symmetry score after the landmark point is fixed during the computation of $\sigma(X)$. Figures 1 and 4 show this improvement in the symmetry score for different points on the Cat, David and Dog shapes taken from the TOSCA dataset [BBK08]. As expected, landmark points close to the symmetry axis of the shape introduce little if any improvement in the symmetry score and by extension the condition number. Figure 5 also shows the first landmark as picked by the FPS, the Planned Landmark Selection by Tevs et al. [TBW*11], and our Condition Number Based Sampling. In an attempt to reduce randomness in FPS, here and in all of the evaluation below, we always selected the farthest point from a random point as the initial landmark in FPS.

Evaluation and comparison with prior work To evaluate our landmark selection method, we compare it with the standard Farthest Point Sampling (FPS) and the Planned Landmark Sampling (PLS) method of Tevs et al. [TBW*11]. Our goals are two-fold: to compare the stability properties of the landmark selection methods as measured by the symmetry score and to demonstrate that the symmetry score provides a useful measure for the “hardness” of matching the shapes.

For this, we evaluate landmark selection methods using two criteria: first we compute the symmetry score of the shapes after introducing the landmarks using the method described in Section 4.2. Then, we use these sample points as landmark correspondences between pairs of shapes within the same class of deformed models, and use the spectral matching approach of Leordeanu and Hebert [LH05] to discover the remaining correspondences. To evaluate the matchings,

we measure the average distance of the computed correspondences from the ground truth across all shapes in the class.

In this experiment, we used 8 classes of shapes from the TOSCA dataset of Bronstein et al. [BBK08]. Within each class, we used the undeformed (null) shape to compute the landmark points and evaluate the symmetry score. We then matched the null shape against each remaining shape in the class using the selected points as landmark correspondences and spectral matching of Leordeanu and Hebert [LH05]. To evaluate the matching, we measured the average (across all matchings of null shape to other shapes in the same class) geodesic distance of the computed correspondences from the ground-truth ones (note that shapes within a class in the TOSCA dataset contain the same number of points with corresponding vertices given by vertex id). The landmarks on the null shapes were computed by the authors of [TBW*11].

Figure 6 shows the landmark points selected by Farthest Point Sampling (FPS), Planned Landmark Sampling (PLS) and our (CNBS) method on 6 null shapes. The red points correspond to the first three landmarks, while the green ones represent the remaining landmarks. We used the same number of landmark points as provided by Tevs et al. and this number is reported in the leftmost column of Figure 7a.

The table in Figure 7a shows the estimate of the symmetry score after adding one, three and all landmark points on each null shape, the maximum number being defined by PLS. Note that the first landmark selected by PLS is computed by only comparing descriptor values. Thus it may not reduce the symmetry score significantly. Indeed, Figure 5 shows the first landmark point on the cat shape for the three sampling strategies. The first point picked by PLS lies close to the symmetry axis thus not reducing the symmetry score significantly. However, with 3 landmark points, both PLS and our method significantly outperform FPS, while for a larger number of landmarks all methods seem to have comparable performance.

The table in Figure 7b, shows the performance of spectral matching [LH05] after fixing 3 and all landmark points. In particular, we compute dense correspondences and measure the average geodesic distance of the computed correspondences to the ground-truth in percent of the half-diameter of the shape. Note that we averaged this distance over matchings from the null shape to each shape in a given class, which roughly simulate a set of noisy deformations of the null shape. Note that, as predicted, the symmetry score shown in 7a correlates very well with the average error in 7b. In other words, for shapes and landmarks with high symmetry score, the average error also increases, which suggests that the symmetry score and by extension the condition number are meaningful measures for the hardness of matching a shape. Therefore, based on these preliminary results, we conclude that 1) the condition number is consistent with the matching error, 2) our method is comparable to or slightly better than PLS and both methods are better than FPS, and 3) with more samples, these methods tend to be identical.




Figure 5: First landmark selected by FPS (yellow), PLS [TBW*11] (red) and our method (blue). Note the proximity of the red and yellow points to the symmetry axis.

8. Conclusion & Future Work

In this paper, we have analyzed the stability of shape matching for a fixed distortion measure and showed that it can be captured by a notion called the shape condition number. We also showed that the condition number is closely related to a simpler measure which we called the shape symmetry score. We demonstrated ways in which the condition number can be estimated in practice, and described how the computation can be used to guide to the selection of landmark points to improve stability of shape matching.

A major limitation of our analysis is that it does not address partial shape matching, which can result in energy landscapes with comparable local minima in the absence of shape symmetries. Analyzing stability of partial shape matching is an important and challenging problem for future work. Similarly, we are interested in analyzing different matching *methods* for their stability with respect to the lower bound provided by the condition number, and extending stability analysis of other problems beyond shape matching.

Acknowledgments This work was supported by NSF grants CCF 1011228, a KAUST-Stanford AEA grant, and a Stanford Graduate Fellowship. We also thank the anonymous reviewers for the valuable comments and suggestions.

References

- [BBB*10] BRONSTEIN A. M., BRONSTEIN M. M., BUSTOS B., CASTELLANI U., CRISTANI M., FALCIDIENO B., GUIBAS L. J., ET AL.: SHREC'10 Track: Feature detection and description. In *Proc. 3DOR* (2010), pp. 79–86. 2
- [BBK06] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103, 5 (2006), 1168–1172. 2, 3
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 7, 8
- [BCPP98] BURKARD R., CELA E., PARDALOS P., PITSOULIS L.: The quadratic assignment problem. *Handbook of Combinatorial Optimization* (1998), 241–338. 4
- [BM77] BOYER R. S., MOORE J. S.: A fast string searching algorithm. *Commun. ACM* 20 (October 1977). 3
- [CK10] CHERTOK M., KELLER Y.: Spectral symmetry analysis. *IEEE Trans. PAMI* 32, 7 (2010), 1227–1238. 5
- [CM92] CHEN Y., MEDIONI G.: Object modelling by registration of multiple range images. *Image Vision Comput.* 10 (1992), 145–155. 2, 6

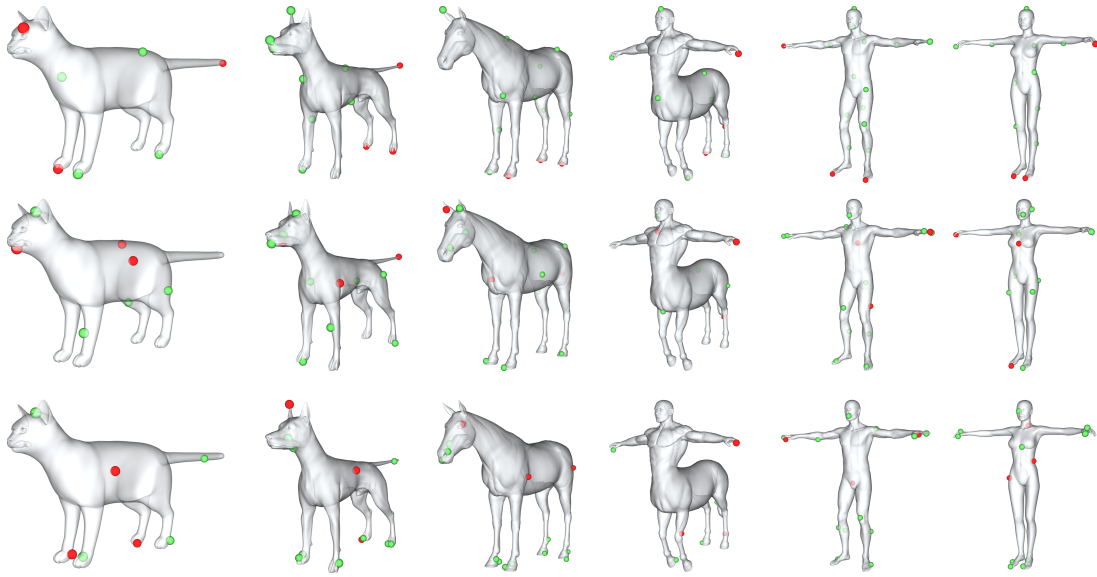


Figure 6: Landmark points selected by Farthest Point Sampling (top), method of Tevs et al. [TBW*11] (middle) and our method (bottom). The points marked in red are the first 3 landmarks, while the points in green are the remaining landmarks.

	One Landmark			Three Landmarks			All Landmarks				Three Landmarks			All Landmarks		
	FPS	PLS	CNBS	FPS	PLS	CNBS	FPS	PLS	CNBS	# pts	FPS	PLS	CNBS	FPS	PLS	CNBS
Horse	40.6	Inf	17.3	21.1	15.9	13	8.1	7.5	7.8	14	1.8	1.36	1.06	0.66	0.64	0.61
Victoria	58.2	58.2	11.8	9.4	7.4	6.3	4.7	3.8	3.7	13	0.76	0.57	0.48	0.39	0.34	0.33
Michael	80.1	56.4	19.1	9	7.6	6.1	5	4.7	4.8	16	0.71	0.62	0.54	0.42	0.4	0.39
David	57.5	43.9	12	11.02	10.15	5.1	6.49	5.6	4.2	15	0.87	0.77	0.4	0.58	0.44	0.31
Dog	Inf	53.8	12.3	20.46	6.31	7.2	5.28	5.44	5.57	12	1.55	0.64	0.61	0.47	0.46	0.43
Centaur	Inf	28.2	14.5	12.74	8.51	7.77	7.9	6.38	5.43	7	1.01	0.64	0.71	0.59	0.51	0.52
CAT	Inf	46.2	23.5	15.21	8.87	10.22	6.56	6.51	6.42	7	1.27	0.81	0.84	0.56	0.54	0.54
Wolf	266.1	25.2	11.2	5.82	4.51	4.64	2.74	2.71	2.56	5	0.48	0.37	0.37	0.23	0.22	0.21

(a) Estimates of the symmetry score after adding landmark correspondences

(b) Dense correspondence accuracy

Figure 7: (a) Estimates of the symmetry score after fixing one, three and all landmarks using Farthest Point Sampling (FPS), Planned Landmark Sampling (PLS) [TBW*11], and our CNBS method. The total number of landmarks fixed is shown in the last column. Note that the first landmark in PLS is chosen using only point descriptor comparison, (b) Average distance of the dense correspondences to the ground-truth (in percent of half-diameter of the shape), after applying spectral matching [LH05] to obtain dense correspondences from sparse ones.

[GG04] GELFAND N., GUIBAS L.: Shape segmentation using local slippage analysis. In *Proc. SGP* (2004), pp. 219–228. 2, 6

[GIRL03] GELFAND N., IKEMOTO L., RUSINKIEWICZ S., LEVOY M.: Geometrically stable sampling for the ICP algorithm. In *3DIM* (2003). 2, 6

[HAWG08] HUANG Q.-X., ADAMS B., WICKE M., GUIBAS L. J.: Non-rigid registration under isometric deformations. *CGF (Proc. SGP)* 27, 5 (2008), 1449–1457. 2, 5

[HSVTP10] HILDEBRANDT K., SCHULZ C., VON TYCOWICZ C., POLTHIER K.: Eigenmodes of surface energies for shape analysis. In *Advances in Geometric Modeling and Processing* (2010), vol. 6130 of *Lecture Notes in Computer Science*, Springer, pp. 296–314. 6

[HWAG09] HUANG Q., WICKE M., ADAMS B., GUIBAS L.: Shape decomposition using modal analysis. *Computer Graphics Forum* 28, 2 (2009), 407–416. 6

[KMP77] KNUTH D. E., MORRIS J., PRATT V. R.: Fast pattern matching in strings. *SIAM Journal of Computing* 6, 2 (1977), 323–350. 3

[KMP07] KILIAN M., MITRA N. J., POTTMANN H.: Geometric modeling in shape space. In *Proc. SIGGRAPH* (2007), pp. 64:1–64:8. 1

[KS04] KRAEVOY V., SHEFFER A.: Cross-parameterization and compatible remeshing of 3d models. In *Proc. of SIGGRAPH* (2004), pp. 861–869. 1

[LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. In *Proc. of SIGGRAPH* (2009), vol. 28:3, pp. 72:1–72:12. 2, 3, 6

[LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *Proc. of ICCV* (2005), pp. 1482–1489. 5, 7, 8, 9

[LPC*00] LEVOY M., PULLI K., CURLESS B., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINZTON M., ANDERSON S., DAVIS J., GINSBERG J., SHADE J., FULK D.: The digital Michelangelo project: 3D scanning of large statues. In *Proc. SIGGRAPH* (2000), pp. 131–144. 1

[LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. *CGF (Proc. SGP'08)* 27, 5 (2008). 2, 5

[Mém07] MÉMOLI F.: On the use of Gromov-Hausdorff Distances for Shape Comparison. In *Symposium on Point Based Graphics* (2007), pp. 81–90. 4

[MGPG04] MITRA N. J., GELFAND N., POTTMANN H., GUIBAS L.: Registration of point cloud data from a geometric optimization perspective. In *Proc. SGP* (2004), pp. 23–31. 2

- [MS05] MÉMOLI F., SAPIRO G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.* 5, 3 (2005). 2, 3
- [MV98] MAINTZ J., VIERGEVER M. A.: A survey of medical image registration. *Medical Image Analysis* 2 (1998), 1 – 36. 1
- [OMMG10] OVSJANIKOV M., MERIGOT Q., MEMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *CGF* 29, 5 (2010), 1555–1564. 2, 6
- [RBBK07] RAVIV D., BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Symmetries of non-rigid shapes. In *Proc. NRTL (ICCV)* (2007), pp. 1–7. 2
- [RPSS10] RUGGERI M. R., PATANÈ G., SPAGNUOLO M., SAUPE D.: Spectral-driven isometry-invariant matching of 3d shapes. *IJCV* (2010), 248–265. 2
- [SCF10] SUN J., CHEN X., FUNKHOUSER T.: Fuzzy geodesics and consistent sparse correspondences for deformable shapes. *CGF (Proc. SGP)* 29, 5 (2010). 2
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *CGF (Proc. SGP)* 28, 5 (2009). 2
- [TB97] TREFETHEN L. N., BAU D.: *Numerical Linear Algebra*. SIAM: Society for Industrial and Applied Mathematics, 1997. 3
- [TBW*09] TEVS A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.-P.: Isometric registration of ambiguous and partial data. In *Proc. CVPR* (2009), pp. 1185–1192. 2, 3
- [TBW*11] TEVS A., BERNER A., WAND M., IHRKE I., SEIDEL H.-P.: Intrinsic shape matching by planned landmark sampling. In *Proc. Eurographics* (2011), p. to appear. 3, 6, 7, 8, 9
- [WAO*09] WAND M., ADAMS B., OVSJANIKOV M., BERNER A., BOKELOH M., JENKE P., GUIBAS L., SEIDEL H.-P., SCHILLING A.: Efficient reconstruction of nonrigid shape and motion from real-time 3d scanner data. *ACM Trans. Graph.* 28 (May 2009), 15:1–15:15. 2, 5
- [WJW05] WAINWRIGHT M. J., JAAKKOLA T., WILLSKY A. S.: Map estimation via agreement on trees: message-passing and linear programming. *IEEE Transactions on Information Theory* 51, 11 (2005), 3697–3717. 5
- [ZSCO*08] ZHANG H., SHEFFER A., COHEN-OR, ZHOU Q., VAN KAICK O., TAGLIASACCHI A.: Deformation-driven shape correspondence. In *Proc. SGP* (2008), pp. 1431–1439. 2, 5

9. Appendix

Proof of Theorem 4.1 Throughout this proof we will assume that $D_{X,Y}$ and $d_{X,Y}$ are given as L^p norms, according to Equations 2 and 3 with the same p .

1. $\sigma(X) \leq \kappa(X)$

Proof Our goal is to show that for any bijection $T : X \rightarrow X$, we can construct \tilde{X} such that:

$$\frac{d(T, Id)}{D_{X,X}(T)} \leq \frac{d(\tilde{Id}, T_{\text{opt}}(X, \tilde{X}))}{d(X, \tilde{X})}.$$

For this, given $T : X \rightarrow X$, we let \tilde{X} be such that $d_{\tilde{X}}(\tilde{Id}(x), \tilde{Id}(x')) = d_X(T(x), T(x'))$. This implies:

$$d(X, \tilde{X}) = D_{X, \tilde{X}}(\tilde{Id}) = D_{X,X}(T)$$

Moreover, note that the map $S : X \rightarrow \tilde{X}$ defined as $S = \tilde{Id} \circ T^{-1}$ has zero distortion since for all x, x' :

$$d_X(x, x') = d_{\tilde{X}}(\tilde{Id}(T^{-1}(x)), \tilde{Id}(T^{-1}(x'))) = d_{\tilde{X}}(S(x), S(x')).$$

Therefore, $S \in T_{\text{opt}}(X, \tilde{X})$. Moreover, $d(T, Id) = d(\tilde{Id}, S)$, since for any x :

$$d_X(T(x), x) = d_{\tilde{X}}(\tilde{Id}(x), \tilde{Id}(T^{-1}(x))) = d_{\tilde{X}}(\tilde{Id}(x), S(x)).$$

Therefore,

$$\kappa(X) \geq \frac{d(\tilde{Id}, T_{\text{opt}}(X, \tilde{X}))}{d(X, \tilde{X})} \geq \frac{d(\tilde{Id}, S)}{d(X, \tilde{X})} = \frac{d(T, Id)}{D_{X,X}(T)}.$$

Since this holds for any T , $\kappa(X) \geq \sigma(X)$ as claimed. \square

2. $\kappa(X) \leq 2\sigma(X) + 1$

Proof Given a shape X and its deformed version (\tilde{X}, \tilde{Id}) with a map $T : X \rightarrow \tilde{X}$, s.t. $D_{X, \tilde{X}}(T) \leq D_{X, \tilde{X}}(\tilde{Id})$ our goal is to produce a map $S : X \rightarrow X$ such that:

$$\frac{d(S, Id)}{D_{X,X}(S)} \geq \frac{1}{2} \frac{d(\tilde{Id}, T)}{d(X, \tilde{X})} - \frac{1}{2}.$$

For this, simply let $S = T^{-1} \circ \tilde{Id}$. Then, note that:

$$D_{X,X}(S) \leq D_{\tilde{X}, X}(T^{-1}) + D_{X, \tilde{X}}(\tilde{Id}) = D_{X, \tilde{X}}(T) + D_{X, \tilde{X}}(\tilde{Id}).$$

Moreover,

$$d(S, Id) \geq d(\tilde{Id}, T) - D_{X, \tilde{X}}(T),$$

since $d(S, Id)$ and $D_{X, \tilde{X}}(T)$, are given by L^p norms for the same p . Thus, assuming $d(\tilde{Id}, T) \geq D_{X, \tilde{X}}(T)$ (otherwise $\kappa(X) \leq 1$ and the bound holds trivially), and using the triangle inequality on $D_{X,X}(S)$:

$$\begin{aligned} \frac{d(S, Id)}{D_{X,X}(S)} &\geq \frac{d(\tilde{Id}, T) - D_{X, \tilde{X}}(T)}{D_{X, \tilde{X}}(T) + D_{X, \tilde{X}}(\tilde{Id})} \geq \frac{d(\tilde{Id}, T)}{2D_{X, \tilde{X}}(\tilde{Id})} - \frac{D_{X, \tilde{X}}(T)}{2D_{X, \tilde{X}}(\tilde{Id})} \\ &\geq \frac{d(\tilde{Id}, T)}{2D_{X, \tilde{X}}(\tilde{Id})} - \frac{D_{X, \tilde{X}}(\tilde{Id})}{2D_{X, \tilde{X}}(\tilde{Id})} = \frac{1}{2} \frac{d(\tilde{Id}, T)}{D_{X, \tilde{X}}(\tilde{Id})} - \frac{1}{2}. \end{aligned}$$

Proof of Lemma 4.2

Proof We first prove necessity: suppose $\kappa(X) = \sigma(X)$. Let S_{opt} be the self-map that achieves $\sigma(X)$, and \tilde{X} be such that: $d_{\tilde{X}}(\tilde{Id}(x), \tilde{Id}(x')) = d_X(S_{\text{opt}}(x), S_{\text{opt}}(x'))$, and the map $T_2 : X \rightarrow \tilde{X}$, $T_2 = \tilde{Id} \circ S^{-1}$. Then, T_2 achieves zero distortion and moreover, since

$$\kappa(X) \geq \frac{d(T_{\text{opt}}(X, \tilde{X}), \tilde{Id})}{d(\tilde{X}, X)} \geq \frac{d(T_2, \tilde{Id})}{d(\tilde{X}, X)} = \sigma(X),$$

and we are given that $\kappa(X) = \sigma(X)$, the inequalities are equalities and $\kappa(X)$ is achieved by \tilde{X} and the isometry T_2 .

To prove sufficiency, suppose that $\kappa(X)$ is achieved by a shape \tilde{X} and a map $T : X \rightarrow \tilde{X}$ such that $D_{X, \tilde{X}}(T) = 0$. Then, consider the map $S : X \rightarrow X$, $S = T^{-1} \circ \tilde{Id}$. Since T has zero distortion: $d_X(S(x), S(x')) = d_{\tilde{X}}(T(S(x)), T(S(x'))) = d_{\tilde{X}}(\tilde{Id}(x), \tilde{Id}(x')) \forall x, x' \in X$. Thus: $D_{X,X}(S) = D_{X, \tilde{X}}(\tilde{Id})$. Similarly: $d_X(S, Id) = d(\tilde{Id}, T)$. Therefore, S is such that, $\kappa(X) = \frac{d(S, Id)}{D_{X,X}(S)}$, which implies, using Lemma 4.1, that $\kappa(X) = \sigma(X)$. \square