

One Point Isometric Matching with the Heat Kernel

Maks Ovsjanikov¹ and Quentin Mérigot^{1,2} and Facundo Mémoli¹ and Leonidas Guibas¹

¹Stanford University, ²INRIA Sophia-Antipolis

Abstract

A common operation in many geometry processing algorithms consists of finding correspondences between pairs of shapes by finding structure-preserving maps between them. A particularly useful case of such maps is isometries, which preserve geodesic distances between points on each shape. Although several algorithms have been proposed to find approximately isometric maps between a pair of shapes, the structure of the space of isometries is not well understood. In this paper, we show that under mild genericity conditions, a single correspondence can be used to recover an isometry defined on entire shapes, and thus the space of all isometries can be parameterized by one correspondence between a pair of points. Perhaps surprisingly, this result is general, and does not depend on the dimensionality or the genus, and is valid for compact manifolds in any dimension. Moreover, we show that both the initial correspondence and the isometry can be recovered efficiently in practice. This allows us to devise an algorithm to find intrinsic symmetries of shapes, match shapes undergoing isometric deformations, as well as match partial and incomplete models efficiently.

1. Introduction

Finding structure-preserving maps between shapes is one of the most common operations in geometry processing applications, including shape matching, shape retrieval, and recognition. High quality correspondences are also essential in morphing, and shape interpolation [KMP07]. Furthermore, many measures of similarity between shapes are defined through the quality of the maps between them [MS05].

A particularly useful type of maps between shapes is an intrinsic isometry, which preserves pairwise geodesic distances between points. Since many real-world deformations, such as articulated motion, are approximately isometric, finding near isometric maps can have practical significance in non-rigid shape comparison and recognition, with applications in manufacturing and medical imaging [SGT98].

At first sight, finding an isometric map between two shapes is complicated by the fact that the space of all possible maps can be quite large. Moreover, even if the map is given, directly verifying if it is an isometry, can potentially have quadratic complexity since distance preservation should be enforced at all pairs of points. Nevertheless, isometric maps provide a very rich structure that makes isometric shape matching tractable. For example, many shape matching algorithms (e.g. [BBK06, HAWG08]) implicitly assume that the knowledge of a small number of correspondences can be extended to a full isometry. Similarly, verifying distance




Figure 1: Matching between a full and a partial model. Red is the correspondence found by our algorithm, and green are the interpolated correspondences using our method.

preservation on a small number of pairs of points can improve efficiency and robustness without sacrificing accuracy, [TBW⁺09]. However, the precise nature of these constraints, such as the minimal number of correspondences necessary to recover the full isometry, is not well understood.

In this paper, we make a step towards analyzing and exploiting the structure of isometries between a pair of shapes, by showing that under mild genericity conditions the knowledge of a single correspondence can be used to recover an isometry defined on entire shapes. We first show that the image of every point is characterized by the preservation of the heat kernel to the given correspondence points. Furthermore, we prove that generically, any map that preserves the heat kernel to a fixed point is an isometry. This allows us

to quickly check whether a potential correspondence can be extended to a full isometry, without verifying all pairwise geodesic distances. Using these observations, we devise an efficient algorithm which can be used to (1) find intrinsic symmetries of shapes, (2) match shapes undergoing isometric deformations, and (3) match partial and incomplete models efficiently. We also demonstrate that local descriptors and multiple landmark correspondences can be naturally incorporated into our method to improve robustness and accuracy.

1.1. Related Work

Finding correspondences between two possibly deformed shapes is a very well-studied problem in geometry processing with a wealth of work spanning several decades. We give a brief overview of the work most closely related to ours, in particular focusing on algorithms aimed at finding dense maps between pairs of non-rigid shapes.

Several variants of the classical ICP algorithm have been proposed to match deformable shapes when the deformation between them is small (e.g. [ACP03, WJH⁺07, LSP08]). These methods typically introduce a regularization term that allows for small elastic deformation of the shape, but cannot be easily extended to handle shapes in arbitrary poses.

Bronstein et al. [BBK06] propose an optimization-based technique to find as-isometric-as-possible mappings of a selected number of points from the source shape to the target shape. Similar ideas have been exploited by Huang et al. [HAWG08] who extend a sparse set of correspondences to a dense map using geodesic landmark coordinates and Tevs et al. [TBW⁺09] who use a robust random sampling technique to find near-isometric matches. Lasowski et al. [LTSW09] sub-sample the shape to a small set of points and formulate a probability distribution over all possible isometric matches between samples, revealing the symmetry structure of the object. Ahmed and colleagues [ATR⁺08] follow a similar paradigm, by first matching a sparse set of features extracted from a high-resolution video stream, and then obtain dense correspondences using Laplacian diffusion. Zhang et al. [ZSCO⁺08] perform a combinatorial search on the space of feature matches and use a non-rigid shape deformation technique [LSLCO05] to interpolate correspondences and measure distortion. All of the above techniques use the intuition that a small number of landmark correspondences is sufficient to obtain a dense map that preserves all pairwise distances. In this paper, we formalize this intuition by showing that, generically, the isometric map between shapes can be efficiently recovered from a single correspondence.

Spectral invariants such as the eigenfunctions of the Laplace-Beltrami operator have been used for near-isometric shape matching by Mateus et al. [MHK⁺08] and Jain et al. [JZvK07] as well as for segmentation [dGGV08, Reu10] and intrinsic symmetry detection [OSG08] among many others. In this work we employ spectral ideas, and in particular the heat kernel, to parametrize and recover isometric maps.

Note that in the case of rigid matching, the isometry group is the Euclidean group $E(3)$, where three correspondences constrain all degrees of freedom. This observation has led to extremely robust shape matching methods, e.g. [AMCO08].

Perhaps most closely related to our method is the remarkable work by Lipman and Funkhouser [LF09], who propose a shape matching method for genus zero surfaces. Their method is based on the observation that for such surfaces, the isometry group is a subgroup of the Möbius group, which can be parametrized by three distinct correspondences. In contrast, our technique is applicable to shapes of any genus and dimension, satisfying mild genericity conditions. Furthermore, unlike the Möbius map, which is conformal, we operate in the space of isometric maps, reducing the minimum number of correspondences and allowing us to provide a simple sufficient condition for a map to be an isometry.

Our work is also closely related to shape parametrization techniques that rely on Laplace-Beltrami operator and the heat kernel [BBG94, JMS10]. Unlike these techniques, our method allows to define a canonical, global embedding of the shape, which is essential for shape matching.

Our use of the heat kernel and theoretical results are also related to the Heat Kernel Signature (HKS) of Sun et al. [SOG09], who prove that a homeomorphism between two compact manifolds that preserves the diagonal of the heat kernel at each point must be an isometry. Unfortunately, the regularization condition of continuity is rather difficult to enforce in practice. We discuss the relation of our results to [SOG09] in detail in Section 3.

Finally our isometry verification step follows the paradigm proposed in geometric hashing [WR97], and used in [LF09].

2. Heat Kernel and Heat Kernel Map

The main ingredient for our method is the heat kernel. In this Section we briefly overview its properties and define the Heat Kernel Map, used in our isometric matching algorithm.

2.1. Mathematical Background

Given a compact Riemannian manifold M without boundary, let $u(x, t) : M \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be the amount of heat at a point $x \in M$ at time t , for some initial heat distribution $f : M \rightarrow \mathbb{R}^+$ at time 0. Then, $u(x, t)$ satisfies the heat equation:

$$\frac{\partial u}{\partial t} = -\Delta_M u, \text{ and } \lim_{t \rightarrow 0} u(x, t) = f(x)$$

where Δ_M is the Laplace-Beltrami operator of M . The solution to the heat equation at time t given f , can be computed through the heat operator $H_t : L^2 \rightarrow L^2$, where L^2 is the space of all smooth, square integrable functions on M :

$$u(x, t) = (H_t f)(x) = \left(e^{-t\Delta_M} f \right)(x),$$

where the $e^{-t\Delta_M}$ is the operator exponential. An exposition of these results can be found in e.g. [Gri06].

The *heat kernel* on M is the unique function $k_t^M(x, y) : \mathbb{R}^+ \times M \times M \rightarrow \mathbb{R}^+$, such that for all $f \in L^2$, $x \in M$, $t > 0$:

$$u(x, t) = (e^{-t\Delta_M} f)(x) = \int_M k_t^M(x, y) f(y) dy.$$

Intuitively, the value of the heat kernel $k_t^M(x, y)$ measures the amount of heat transferred from point x to point y in time t .

The heat kernel has an alternative interpretation as the transition density function of Brownian motion on M (see e.g. [Hsu02] for details). If W_t^x is the Brownian motion on M starting at point x , then the probability $\mathbb{P}(W_t^x \in C)$ that at time t , W_t^x lies inside an integrable set $C \subseteq M$, is given by:

$$\mathbb{P}(W_t^x \in C) = \int_C k_t(x, y) dy.$$

This interpretation makes apparent the richness of the information contained in the heat kernel. Since the transition probability is related not only to the shortest path between two points, but is rather a weighted average over all paths possible in time t , the heat kernel intuitively contains information beyond the shortest path distance between two points. We will formalize this intuition in Section 3.

The heat kernel has many nice properties, which have made it a useful tool in several geometry processing applications [SOG09, Mém09, GBAL09]. The properties of the heat kernel that are most relevant to our work are that the heat kernel is an isometry invariant (intrinsic) and that it is multiscale.

Lemma 2.1 Let M and N be two compact, connected Riemannian manifolds without boundary. A map T from M to N is an isometry if and only if $k_t^M(x, y) = k_t^N(T(x), T(y))$ for all $x, y \in M$, $t > 0$.

The invariance of the heat kernel under isometries follows from the intrinsic nature of the Laplace-Beltrami operator. The converse follows from Varadhan's Lemma (e.g. [Gri06]) which asserts that for every pair of points x, y in M , the geodesic distance $d^2(x, y) = -4 \lim_{t \rightarrow 0} t \log k_t^M(x, y)$. Thus, any map that preserves the heat kernel must be distance-preserving. Furthermore, we show in the Appendix that any such map must be surjective and thus an isometry.

The multiscale property of the heat kernel implies that for small values of t , the heat kernel $k_t^M(x, \cdot)$ can be well approximated by the heat kernel of a small geodesic neighborhood of the point x . The authors of [SOG09] have demonstrated this property empirically, and cited theoretical results that suggest this behavior for either infinitesimal t or for large t as long as the neighborhood is large enough.

Another relation between t and the size of the neighborhood of x was obtained by Grigor'yan [Gri99], who showed that if M is a *non-compact* Riemannian geodesically complete manifold satisfying bounded volume growth conditions, and W_x^t is the Brownian motion on M starting at x , the probability that the Brownian motion W_x^t at time t belongs to the geodesic ball $B(x, \sqrt{c_n t \log t})$ goes to 1 as t grows to infinity.




Figure 2: Expected geodesic distance $\mathbb{E}(x, d(W_x^t))$ between a point x and the Brownian motion W_x^t starting at x .

This suggests (but does not prove) that the heat kernel, which is the transition density function of the Brownian motion, is determined by the neighborhood of x with geodesic radius $O(\sqrt{c_n t \log t})$. We will use this intuition when extending our isometric matching to handle incomplete shapes.

As an illustration, we show in Figure 2 the evolution of the expected value of the geodesic distance of Brownian motion from the source point x as a function of time, for three source points on a mesh. Note that this expected value is given simply as $\mathbb{E}(d(x, W_x^t)) = \int_M d(x, y) k_t(x, y) dy$. Moreover, by Markov's inequality, we can bound the amount of heat contained outside of a ball $B(x, r)$ around x by:

$$\int_{M \setminus B(x, r)} k_t(x, y) dy \leq \frac{1}{r} \int_M d(x, y) k_t(x, y) dy.$$

It is interesting to note that the graph of $\mathbb{E}(d(x, W_x^t))$ for small time values allows us to distinguish source points on a saddle from source points in a positively curved area. For the same shape as in Figure 2, and a few source points x , we computed numerically the smallest radius R such that 90% of the heat distribution $k_t^M(x, \cdot)$ is contained within the geodesic ball $B(x, R)$. If we let M' be the geodesic ball $B(x, R)$, the L^2 difference between the restrictions of the heat distributions $k_{t/2}^M(x, \cdot)$ and $k_{t/2}^{M'}(x, \cdot)$ to the ball $B(x, R)$ remains low — consistently smaller than 2%.

Note that some remarks from the point of view of homogenization of PDEs, made in [Mém09] further attempt to justify the interpretation of t as a geometric scale.

2.2. Heat Kernel Map

Given a compact manifold M and a fixed source point p , we define the Heat Kernel Map $\Phi_p^M(x)$ of any point $x \in M$:

$$\Phi_p^M : M \rightarrow F, \quad \Phi_p^M(x) = k_t^M(p, x),$$

where F is the space of functions from \mathbb{R}^+ to \mathbb{R}^+ . Thus, Φ_p^M associates with every point $x \in M$ a real-valued function of one parameter t given by $k_t^M(p, x)$. Note that the Heat Kernel Map depends on the source point p , which can be any point on the manifold M .

One of the main results of this paper is that under mild genericity conditions, the Heat Kernel Map is *injective*. In other words, $\Phi_p^M(x) = \Phi_p^M(y)$ if and only if $x = y$. We will use this observation to guide our isometric matching method.




Figure 3: Overview of our method: Given a shape M and a point p on M , we construct the Heat Kernel Map, which associates for each point $x \in M$, a function from $t \in \mathbb{R}$ to \mathbb{R} , given by $t \mapsto k_t^M(p, x)$. Under mild genericity conditions, we prove that each point x is associated with a unique function (Theorem 3.1). Given another shape N , and a point q corresponding to p , we match any other point $x \in M$, by finding the unique point $y \in N$ such that their Heat Kernel Maps coincide. We show that the global map constructed this way is an isometry (Theorem 3.5). Moreover, each step can be done efficiently in practice (Sections 4, 5).

Our general approach is as follows (see Figure 3 for a summary): given two compact manifolds M and N , and a known correspondence between a point $p \in M$ and $f(p) \in N$, any isometry $f : M \rightarrow N$ must preserve the heat kernel, and thus for any point $x \in M$, $\Phi_p^M(x) = \Phi_{f(p)}^N(f(x))$. But since $\Phi_{f(p)}^N$ is injective, there cannot be another point $z \in N$ such that $\Phi_{f(p)}^N(f(x)) = \Phi_{f(p)}^N(z)$. Therefore, every point $x \in M$ will have a *unique* corresponding point $f(x) \in N$, characterized by the preservation of the heat kernel: $\Phi_p^M(x) = \Phi_{f(p)}^N(f(x))$. To prove this, we need two genericity assumptions.

Generic Manifolds: A compact Riemannian manifold M is called *generic* if its Laplace-Beltrami operator does not have repeated eigenvalues. Note that this is the same genericity condition considered by Sun et al. [SOG09]. It is well known (see [BU83], Theorem 3.1) that this condition holds for a generic Riemannian metric on any smooth manifold of dimension at least 2. Therefore, given a manifold M that does not satisfy the genericity assumption, an infinitesimal perturbation of the metric makes M generic.

Generic Points: A point p in M is called *generic* if $\phi_i^M(p) \neq 0$ for every eigenfunction ϕ_i^M of the Laplace-Beltrami operator of M . Note that the set of generic points of M has full measure. Indeed, the nodal domain of every eigenfunction ϕ_i^M is $(d-1)$ -dimensional, and as such has zero d -volume [Che76]. Thus, the countable union of all nodal domains has measure zero, and its complement has full measure.

3. Properties of the Heat Kernel Map

Our first result is the following:

Theorem 3.1 Let M be a generic connected compact manifold without boundary and p a generic point on M . Then the Heat Kernel Map is injective: $\Phi_p^M(x) = \Phi_p^M(y)$ only if $x = y$.

To prove Theorem 3.1, we start with a technical Lemma that gives some intuition about the heat kernel and the information that it provides. This Lemma is well known (note that a similar one was proved in [SOG09]), but we include its proof in the Appendix for completeness.

Lemma 3.2 Let (λ_i) be an increasing sequence of non-negative numbers that tend to infinity, and suppose that for

every positive time t , $a(t) = \sum_{i=0}^{\infty} a_i e^{-\lambda_i t} = 0$ where the convergence is absolute. Then all of the coefficients a_i are zero.

Remark 3.3 Note that by this Lemma, if $a(t) = \sum_{i=0}^{\infty} a_i e^{-\lambda_i t}$, and $b(t) = \sum_{i=0}^{\infty} b_i e^{-\mu_i t}$ are two absolutely convergent series with $a_i \neq 0$ and $b_i \neq 0$ then $a(t) = b(t)$ for all t if and only if $a_i = b_i$ and $\lambda_i = \mu_i$ for all i . Thus, the knowledge of the function $a(t)$ for all t is equivalent to the knowledge λ_i, a_i .

Proof of Theorem 3.1 First, recall that the heat kernel on any compact manifold M has the following spectral expansion (see e.g. [Gri06]) :

$$k_t^M(p, x) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i^M(p) \phi_i^M(x), \quad (1)$$

where λ_i^M, ϕ_i^M are the eigenvalues and eigenfunctions of the Laplace-Beltrami operator on M respectively. The eigenvalues λ_i are all non-negative and only accumulate at infinity. By Mercer's theorem, the spectral expansion converges absolutely for all t , which allows us to use Lemma 3.2.

Suppose two points $x, y \in M$ are such that $\Phi_p^M(x) = \Phi_p^M(y)$. The spectral expansion of this equality shows that:

$$\sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i^M(p) \phi_i^M(x) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i^M(p) \phi_i^M(y) \quad \forall t$$

Applying Lemma 3.2 on the difference of the two sides, we get that for every i :

$$\phi_i^M(p) (\phi_i^M(x) - \phi_i^M(y)) = 0$$

and, since p is generic, $\phi_i^M(p)$ should agree with $\phi_i^M(y)$ for every i . But this implies that for any L^2 function $h : M \rightarrow \mathbb{R}$, $h(x) = h(y)$, which can only hold if $x = y$. \square

Intuitively, Theorem 3.1 establishes that for a generic manifold M , a generic point p can be used to provide a unique name for every point of the manifold, given by $\Phi_p^M(x)$. Moreover, since the Heat Kernel Map is intrinsic, this name is invariant under isometric deformations of the shape. A direct consequence of this result is the following:

Corollary 3.4 If M and N be two generic connected compact manifolds without boundary, and p a generic point on M . Then, if $f, g : M \rightarrow N$, are two isometries such that $f(p) = g(p)$, then $f(x) = g(x)$ for all $x \in M$.

Proof: Since M and N are isometric by assumption, $\lambda_i^M = \lambda_i^N$ for all i . Thus, by Lemma 3.2, and using the genericity of p , $\Phi_i^N(f(p))^2 = \Phi_i^M(p)^2 \neq 0$ for every i , and $f(p)$ is generic. Since both f and g are isometries, they must preserve the Heat Kernel Map: for any $x \in M$, $\Phi_p^M(x) = \Phi_{f(p)}^N(f(x)) = \Phi_{g(p)}^N(g(x))$. But $f(p) = g(p)$ by assumption, and the $\Phi_{f(p)}^N$ is injective, so $f(x) = g(x) \forall x$. \square

Therefore, given a generic point p of M , there is at most one isometry f from M to N that sends p to a given point q in N . If such an isometry exists, the image $f(x)$ of a point x is characterized by the preservation of the heat kernel to p : $f(x)$ is the only point s.t. $k_t^M(p, x) = k_t^N(q, f(x))$ for all t .

This immediately suggests an algorithm for finding correspondences between points on M and N . Suppose we know that a point $p \in M$ must map to a point $q \in N$ under some unknown isometry $f: M \rightarrow N$. Then, we can compute the Heat Kernel Map Φ_p^M of points on M and the Heat Kernel Map Φ_q^N of points on N . By the above Theorem, for any point $x \in M$, the corresponding point $f(x)$ is the unique point such that $\Phi_p^M(x) = \Phi_q^N(f(x))$ (see Figure 3).

Unfortunately, we often do not know even a single correspondence between shapes. One remedy suggested by Theorem 3.1 is to fix a point $p \in M$ and try all possible correspondences $q \in N$. For each potential correspondence (p, q) , we can try to extend it using the procedure suggested above. Namely, for every point $x \in M$ try to locate a point $f(x) \in N$ such that $\Phi_p^M(x) = \Phi_q^N(f(x))$. Note however, that even if we succeed to find a corresponding point for each x , we still need to verify that the global map we get with this procedure is an isometry. However, as shown in the following Theorem, preservation of the Heat Kernel Map turns out to be sufficient to ensure pairwise isometric consistency in the exact case. This means, in particular that if we succeed in finding a correspondence for each point x using this simple procedure, then we will have constructed an isometry.

Theorem 3.5 Let M and N be two generic connected compact manifolds and p a generic point on M . Then, any map f such that $f(p)$ is generic and $k_t^M(p, x) = k_t^N(f(p), f(x))$, for every time $t > 0$ and every point x in M , is an isometry.

Note that if the heat kernel was preserved for every *pair* of points x, y in M , this Theorem would be nothing but Lemma 2.1. The non-trivial part consists in showing that the conclusion remains true even when one of the points in M is fixed.

Proof. Setting $x = p$, we get $k_t^M(p, p) = k_t^N(f(p), f(p))$ for all $t > 0$. Thus,

$$\sum_{i=0}^{\infty} e^{-t\lambda_i^M} (\Phi_i^M(p))^2 = \sum_{i=0}^{\infty} e^{-t\lambda_i^N} (\Phi_i^N(f(p)))^2.$$

By the genericity assumption, the coefficients $(\Phi_i^M(p))^2$ and $(\Phi_i^N(f(p)))^2$ are all non-zero. Thus, by Remark 3.3, for every i , $\lambda_i^M = \lambda_i^N$ and $\Phi_i^N(f(p)) = \alpha_i \Phi_i^M(p)$, where $\alpha_i \in \{\pm 1\}$.

Moreover, since $k_t^M(p, x) = k_t^N(f(p), f(x))$, we have:

$$\sum_{i=0}^{\infty} e^{-t\lambda_i^M} \Phi_i^M(p) \Phi_i^M(x) = \sum_{i=0}^{\infty} e^{-t\lambda_i^N} \Phi_i^N(f(p)) \Phi_i^N(f(x))$$

Lemma 3.2 again shows that for every i , $\Phi_i^N(f(x)) = \alpha_i \Phi_i^M(x)$, where $\alpha_i \in \{\pm 1\}$ is independent of x . Therefore for any pair of points (x, y) in M and any positive t ,

$$\begin{aligned} k_t^M(x, y) &= \sum_i e^{-t\lambda_i^M} \Phi_i^M(x) \Phi_i^M(y) \\ &= \sum_i \alpha_i^2 e^{-t\lambda_i^N} \Phi_i^N(f(x)) \Phi_i^N(f(y)) = k_t^N(f(x), f(y)). \end{aligned}$$

Thus, by Lemma 2.1, the function f is an isometry. \square

Note that Theorem 3.5 is rather surprising since it shows that any map that preserves relations of all points to a single point p for all t , also preserves pairwise relations between the points. This, of course, would not be true if the map only preserved distances to p , since it would not disambiguate between points lying on the same geodesic circle around p .

Genericity Conditions Both Theorem 3.1 and Theorem 3.5 may not hold if the manifolds are not generic. The simplest case of a non-generic manifold is the sphere, where for any point p , $k_t(p, x) = k_t(p, y)$ for all t , whenever $d(p, x) = d(p, y)$. Since the geodesic circle, in general, contains more than a single point, both theorems will not hold. Note that a consequence of Theorem 3.1 is that for any generic manifold M , the set of points $\{x \in M, \text{ s.t. } x = T(x) \text{ for some non-trivial self-isometry } T: M \rightarrow M\}$ has zero measure. Similarly, if M is generic, then every self-isometry $T: M \rightarrow M$ is idempotent ($T \circ T = Id$) [OSG08] and the isometry group must be discrete [BU83, Proposition 3.4.].

Nevertheless, we note that Theorem 3.1 can be generalized to show that if the Laplace-Beltrami operator of a manifold M has eigenvalues with bounded multiplicity k , then k generic points can be used to define an injective map. The proof follows the steps of 3.1, and we omit it due to space constraints. As we will see in Section 6 adding extra correspondences can also be useful in practice to increase robustness and the quality of the recovered dense maps.

Manifolds with Boundary Note that although Theorem 3.1 and Theorem 3.5 were stated for compact manifolds without boundary, similar statements also hold for manifolds with boundary with a choice of boundary conditions on the Laplace-Beltrami operator. We omit the treatment of the boundary conditions for simplicity, but see for example Section 4 of [Hsu02] for a discussion of the Dirichlet heat kernel for manifolds with boundary.

Relation to the HKS As mentioned in the Introduction, Theorem 3.5 is related to the Informative Theorem for the Heat Kernel Signature (HKS) in [SOG09], which establishes that any *homeomorphism* T between two generic manifolds M and N , s.t. $k_t^M(x, x) = k_t^N(T(x), T(x)) \forall x, t$ must be an isometry. Note that this does not imply that the map $x \mapsto k_t(x, x)$ is injective, or that if two points have the same

HKS, there exists an isometry mapping one to the other. Moreover, even if such an isometry exists, to construct it, one has to operate in the space of continuous maps, which is difficult to enforce for discretely sampled shapes. Consider, for example, two copies, M and N , of a shape with a reflectional symmetry (e.g. a human shape), so that its only self-isometry is idempotent, and it can satisfy the genericity conditions. The necessity for continuity is clear since for every point on M there are *at least* two points on N with the same HKS. In contrast, Theorem 3.1 proves injectivity of the Heat Kernel Map, while Theorem 3.5 does not assume continuity, which is essential for constructing isometric maps in practice. Note that, although not done here, we leave explicit construction of examples of generic manifolds with non-injective HKS or non-trivial isometry groups as future work.

4. Algorithm

In this Section, we convert the observations of Section 3 into an algorithm for finding approximate isometries in practice. Our algorithm consists of two stages: obtaining a few landmark correspondences and extending them to a dense map on the whole shape. Note that Theorem 3.1 ensures that generically a single correspondence is sufficient to recover the isometry. However, several correspondences can improve stability or be used if the shapes are not generic.

I. Feature Detection Given two shapes M and N , represented as triangle meshes, we first detect a sparse set of feature points $P \subseteq M$ and $Q \subseteq N$. In this work, we use the local maxima of the Heat Kernel Signature [SOG09] for a large time t to identify feature points. Thus, we consider a point p to be a feature point on M if $k_t(p, p) > k_t(x, x)$ for all x in the two ring neighborhood of p . These feature points have been shown to be stable and repeatable in practice [BBB⁺10], but other feature detection methods (such as extrema of Gauss curvature) can also be used to reduce complexity.

II. Single Feature Matching Given the sets of feature points $P \subseteq M$ and $Q \subseteq N$, we fix a point $p \in P$ (either randomly or by user selection) and iterate over points $q_i \in Q$ to find possible candidate correspondences of p . To evaluate each potential match (p, q_i) , we compute the Heat Kernel Maps Φ_p^M and $\Phi_{q_i}^N$, by discretizing for every point $x \in M$ the function $k_t^M(p, x)$ into a set of time samples $\{t_j\}$. Thus, for every point $x \in M$ we compute the vector of size $J = |\{t_j\}|$, whose component j is $k_{t_j}^M(p, x)$, for a set of pre-determined time values t_j . Similarly, for every point $y \in N$, we compute the vector $k_{t_j}^N(q_i, y)$. We then evaluate the quality of the match (p, q_i) by computing the sum of squared distances between heat kernel maps:

$$E(p, q_i) = \sum_{x \in M} \min_{y \in N} \left\| \Phi_p^M(x) - \Phi_{q_i}^N(y) \right\|^2.$$

One choice of distance that works well in practice is:

$$\left\| \Phi_p^M(x) - \Phi_{q_i}^N(y) \right\| = \sup_j t_j \left| k_{t_j}^M(p, x) - k_{t_j}^N(q_i, y) \right|. \quad (2)$$

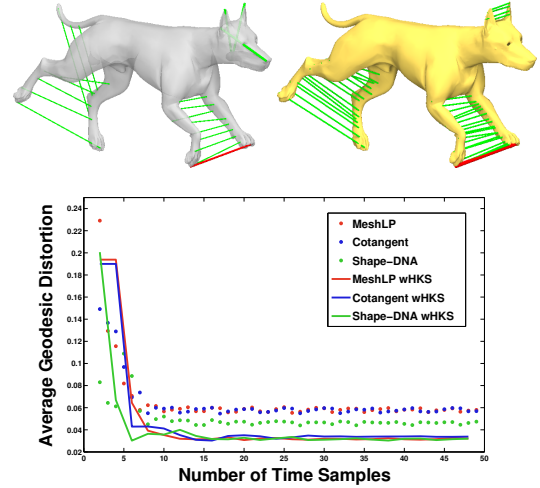


Figure 4: Top Left: Intrinsic symmetry detection without HKS [SOG09]. The extrapolated map (green) is continuous but incorrectly matches the tail to the leg. Top Right: Incorporating HKS fixes these correspondences. Bottom: Quality of the dense map for different discretizations of the Laplace-Beltrami operator. Note that the average geodesic distortion levels out at 10-15 time samples.

Note that we scale each coordinate j of $\Phi_p^M(x)$ by t_j . This ensures the heat kernel does not blow up as t goes to 0, for 2-dimensional surfaces [Gri06, Theorem 3.9]. Also note that by Theorem 3.5, if $E(p, q) = 0$ then there must exist an isometry mapping p to q . In practice, we order the potential matches (p, q_i) in increasing order of $E(p, q_i)$ and only consider those with small error.

III. Correspondence Propagation Given a landmark correspondence (p, q) , assuming the genericity conditions above, for any $x \in M$, Theorem 3.1 ensures that there will exist a unique point $y \in N$ such that $\left\| \Phi_p^M(x) - \Phi_q^N(y) \right\| = 0$. In practice, we simply do the best we can:

$$f(x) = \operatorname{argmin}_{y \in N} \left\| \Phi_p^M(x) - \Phi_q^N(y) \right\|.$$

Thus to every point $x \in M$ we associate its nearest neighbor from N in the space of Heat Kernel Maps, breaking ties arbitrarily. In other words, we find the point from N whose Heat Kernel Map is the closest to that of x . In practice, this can be easily done using a nearest neighbor data structure.

IV. Local Descriptors & Correspondence Augmentation

In the ideal case of exact isometries between generic shapes the procedure outlined above will result in a dense isometric map between the two shapes. However, most shapes in practice are not exactly isometric and can fail to be generic. Here, we introduce two heuristics that greatly improve the quality of the dense maps in practice.

First, we note that any isometric map between two shapes must preserve local descriptors at each point. Therefore, a

map $f : M \rightarrow N$ that preserves the Heat Kernel Map at each point must also preserve intrinsic quantities such as Gauss curvature. The Heat Kernel Signature (HKS) proposed in [SOG09] and independently in [GBAL09] is a powerful local descriptor that characterizes each point in a multiscale fashion. It also fits naturally within our framework, since it is defined through preservation of the heat kernel. To employ the HKS, we augment the Heat Kernel Map by introducing extra $J = |\{t_j\}|$ coordinates for each point x , defined as $k_{t_j}^M(x, x)$. Thus, each point is now associated with a vector of size $2J$, where the first J coordinates are the Heat Kernel Map: $k_{t_j}^M(p, x)$ and the next J are the HKS. Figure 4 demonstrates the effect of introducing the Heat Kernel Signature when using our method for intrinsic symmetry detection.

Similarly, note that we can augment the Heat Kernel Map by introducing extra landmark correspondences. In other words, if we know that a point $p \in M$ is mapped to a point $q \in N$, and $p_2 \in M$ is mapped to $q_2 \in N$ we can associate with each point $x \in M$ a combined Heat Kernel Map, where the extra J coordinates are $k_{t_j}^M(p_2, x)$. In practice, we greedily add a feature correspondence (p_2, q_2) that results in the smallest combined error E .

Note that although we may introduce extra coordinates to the Heat Kernel Map, all of the quantities remain commensurable, and as a result we can use the same distance metric to compare them. Namely, we always use the weighted L^∞ distance, defined in equation (2).

5. Implementation

The two main ingredients to implement the method described above are the computation of the Heat Kernel Map, and nearest neighbor search in the space of Heat Kernel Maps. Below we outline our choices in practice.

Computing the heat kernel on a triangle mesh has been considered in e.g. [SOG09, GBAL09], where the authors use the spectral expansion (1) to compute the heat kernel and its restriction to the diagonal. Namely, given a triangle mesh discretizing the shape, and a matrix L discretizing the Laplace-Beltrami operator with eigenvalues λ_i and eigenvectors ϕ_i , the heat kernel is given as:

$$k_t^M(x, y) = \sum_{i=0}^N e^{-t\lambda_i} \phi_i^M(x) \phi_i^M(y),$$

This choice is also possible in our setting. When the mesh is large, however, computing many eigenvalues and eigenvectors of the matrix L , even if it is sparse, can be demanding. In this case, we can still compute the Heat Kernel Map via sparse matrix exponentiation. Namely, as noted in [SOG09] if $L = A^{-1}W$, where A is the diagonal matrix of area weights, and W is a positive definite weight matrix, the heat kernel matrix for a given time value t is given as: $K_t = \exp(-tL)A^{-1}$, where $\exp(-tL)$ is the matrix exponential. Recall that the Heat Kernel Map is defined via the heat kernel to a fixed point p . Thus, we only need




Figure 5: Intrinsic symmetry detection using our method.

to compute one column of the matrix K_t , corresponding to the point p . This can be done by computing the product $K_t e_p = \exp(-tL)A^{-1}e_p$, where e_p is the indicator vector with a one in the coordinate corresponding to p .

Several techniques have been proposed to evaluate the product $\exp(-tL)v$ for a given vector v without computing the matrix exponential itself. We use the highly stable Krylov subspace projection method implemented in the Expokit software [Sid98]. This method only requires the implementation of the matrix vector multiplication Lv . Note that sparse matrix vector multiplication can be easily parallelized. In practice, we split the matrix L into contiguous sets of rows and issue separate threads to compute the product Lv independently for each set. Note that sparse matrix multiplication has also been implemented in graphics hardware (e.g. [BG08]), which can improve efficiency. Moreover, an efficient multi-resolution approach to computing the heat kernel has been recently proposed by Vaxman et al. [VBCGar].

To compute nearest neighbors in the space of Heat Kernel Maps, we use the ANN [MA] library that supports efficient nearest neighbor search in high dimensions with L^∞ norm.

6. Results

6.1. Symmetry Detection & Choice of Parameters

We first note that the algorithm described in Section 4 can be used to discover intrinsic symmetries of shapes (as in e.g. [OSG08]). Indeed, intrinsic symmetry detection can be regarded as a special case of isometric shape matching, where the two shapes are identical. To demonstrate the performance of our algorithm for different choices of the parameters, we first compare its performance for different discretizations of the Laplace-Beltrami operator, and different number of time samples when computing the Heat Kernel Map. The bottom of Figure 4 shows the average geodesic distortion of the interpolated dense maps for one fixed landmark correspondence using the Laplace-Beltrami operator discretization of Reuter et al. [RWP06] called Shape-DNA (using linear FEM), the method of Belkin et al. [BSW08] called MeshLP and the cotangent weight scheme of Meyer et al. [MDSB02]. To compute the average geodesic distortion, we sampled 100 points $\{s_i\}$ on the mesh, and then computed $\sum_i \sum_j |d(s_i, s_j) - d(f(s_i), f(s_j))|$, where $d(\cdot, \cdot)$ is the geodesic distance on the mesh, and f is the interpolated




Figure 6: Matching deformable shapes under small (a) and medium (b) geodesic distortion. Note that in both cases we identify the symmetry of the objects by finding two possible maps. (a) left: using a single landmark correspondence is sufficient to get a high quality interpolated map, (a) middle: a single landmark correspondence provides a good map, but the ears are switched, which is fixed by computing the second landmark correspondence (a) right. Even under considerable geodesic distortion, 2 landmark correspondences are enough to find high quality direct and symmetric maps, (b) left and right respectively.

map given a landmark correspondence (Section 4). In this example, and elsewhere, we have observed that Shape-DNA performs the best without using the Heat Kernel Signature. However, with the addition of the HKS, all three methods perform similarly. Moreover, the distortion levels out beyond 10-15 time samples. Therefore, the dimensionality of the Heat Kernel Map is relatively low (we used 10 time samples in all the examples below). Finally, the top of Figure 4, shows the effect of adding HKS to the Heat Kernel Map. On the left, the dense map is computed without HKS, and it erroneously matches the tail of the model to its leg. Using HKS fixes this effect and allows us to find a high quality dense map using only one landmark correspondence. Note also the decrease of the average geodesic distortion when using HKS for all three discretizations. Figure 5 shows the result of applying our method for intrinsic symmetry detection on a few sample shapes from the Nonrigid World 3D Database [BB07]. In all of the remaining examples we used the cotangent weight scheme [MDSB02] due to its simplicity and comparable performance. To compute the time discretization of the Heat Kernel Map, we normalized each shape to have unit total area by scaling the area matrix A and computed 10 time samples between $t = 0.001$ and $t = 2$.

In all of the Figures, the computed landmark correspondences are shown in red, whereas the interpolated correspondences are shown in green. We stress that both the landmark correspondences and the interpolated ones are computed automatically, and thus there is no need for human intervention, except possibly when exploring symmetric maps, where the second-best match is used for the initial feature correspondence.

6.2. Isometric Matching

To demonstrate the performance of our algorithm for isometric shape matching, we first consider shapes undergoing deformation with moderate geodesic distortion. Figure 6 shows the result of our method for deformed shapes of the armadillo model. In Figure 6(a) (left), we show that a single landmark correspondence produces a high quality dense map on the entire shape, when the maximum geodesic distortion of the ground-truth correspondence is 5.3% of the shape




Figure 7: Direct and symmetric maps found on two pairs of shapes from the dataset by Vlasic et al. [VBMP08]

diameter. Moreover, we are able to find the intrinsic symmetry of the armadillo, by considering the second best match for the initial feature correspondence (Figure 6(a) middle). This map, however, flips the two ears of the armadillo, which can be fixed by computing a second landmark correspondence (Figure 6(a) right). Similarly, even under considerable geodesic distortion (13.2% of the diameter), our method can find both the direct and the symmetric map (Figure 6(b)).

Note that in all of the examples, the symmetric maps are computed on the same shapes. To avoid clutter, however, we flip one of the shapes for visualization.

Figure 8(a) (left) further demonstrates the effect of adding the second landmark correspondence on the SCAPE dataset [ASK⁺05]. In this case, although the first landmark correspondence is found correctly (Figure 8(a) left), the interpolated correspondences map both arms of one of the models to one arm on the other. Note, that despite this, the interpolated map is consistent in that it associates the correct parts of the arms. Having a second landmark correspondence allows us to find the correct direct and symmetric maps on a variety of poses 8(a) (middle and right) and (Figure 8(b)).

Figure 7 shows the result of our matching algorithm on a few sample shapes from the dataset by Vlasic et al. [VBMP08]. Again, two landmark correspondences allow us to find both direct and symmetric isometric maps between the shapes.

Partial and Approximate Matching

As mentioned in the introduction, our method can be naturally used for matching incomplete shapes. Recall that, as we illustrated in Section 2.1, the heat kernel $k_t(p, \cdot)$ for small




Figure 8: Isometric matches in the SCAPE dataset [ASK⁺05]. (a) left, the landmark correspondence is computed correctly, but the map misses one of the arms. The second landmark correspondence allows to obtain good dense maps: middle, right (a), (b).

values of t reflects the properties of a local neighborhood of p . Thus, to match partial and incomplete shapes, we can construct the Heat Kernel Map by only considering small values of t . Figure 9 shows the result of applying our method for matching an arm to a human body and a leg to a horse model. Note that the arm model is significantly denser (containing 9k points, whereas the human model is 10k points), and the two maps shown in the figure are computed for the same arm model. To compute the heat kernel for the partial models, we used Dirichlet boundary conditions, and computed the Heat Kernel Embedding for time values such that the heat kernel at points close to the boundary remains negligible.

Finally, Figure 10 shows matching obtained on deformable models of the human head. Note that the pair of shapes on the left has different topology (the mouth of the leftmost shape is open), yet our method allows to find approximate maps that respect the local and global structure of the shapes. Note that although we have obtained noisy matches in other similar examples, our method is capable to recover maps even when no near isometry exists. Exploring the stability of our method under topological changes is an interesting problem for the future, both theoretically and in practice.

7. Conclusion and Future Work

In this paper we have described a method for finding dense isometric maps between a pair of generic shapes. We show that the isometry can be recovered from a single correspondence by proving that it is the unique map that preserves the heat kernel to a fixed point. We apply our technique to find intrinsic symmetries of shapes, match shapes undergoing isometric deformations, as well as match partial and incomplete models efficiently.

Perhaps the most interesting and challenging problem for future work is to provide quantitative statements about the behavior of the Heat Kernel Map under near-isometries. Such statements would, in principle, include the inverse Lipschitz constant of the Heat Kernel Map, as well as an approximate version of Theorem 3.5. Finally, it is interesting to




Figure 9: Matching of incomplete models.




Figure 10: Matching shapes with different genus.

consider whether the space of near-isometries can also be parametrized by a small set of correspondences.

Acknowledgments. This work was supported by NSF grants 0808515 and 0808515, NIH grant GM-072970, ONR grant N00014-09-1-0783, a grant from the King Abdullah University of Science and Technology, and a gift from Google, Inc. The authors would also like to thank the anonymous reviewers for the helpful suggestions and comments.

References

- [ACP03] B. Allen, B. Curless, and Z. Popović. The space of human body shapes: reconstruction and parameterization from range scans. *ACM TOG*, 22(3):587–594, 2003.
- [AMCO08] D. Aiger, N. J. Mitra, and D. Cohen-Or. 4-points congruent sets for robust surface registration. *ACM Trans. Graph. (Proc. SIGGRAPH 2008)*, 27(3):1–10, 2008.
- [ASK⁺05] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. SCAPE: shape completion and animation of people. *ACM TOG*, 24(3), 2005.
- [ATR⁺08] N. Ahmed, C. Theobalt, C. Ross, S. Thrun, and H.P. Seidel. Dense correspondence finding for parametrization-free animation reconstruction from video. In *Proc. CVPR*, 2008.
- [BB07] A. Bronstein and R. Bronstein, M. and Kimmel. *Numerical geometry of non-rigid shapes*. Springer-Verlag, 2007.
- [BBB⁺10] AM Bronstein, MM Bronstein, B. Bustos, U. Castellani, M. Crisani, B. Falcidieno, LJ Guibas, I. Kokkinos,

- V. Murino, M. Ovsjanikov, et al. SHREC 2010: robust feature detection and description benchmark. In *Proc. 3DOR*, 2010.
- [BBG94] P. Bérard, G. Besson, and S. Gallot. Embedding riemannian manifolds by their heat kernel. *GAF*, 4:373–398, 1994.
- [BBK06] A.M. Bronstein, M.M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. Natl. Acad. Sci.*, 103(5):1168–1172, 2006.
- [BG08] N. Bell and M. Garland. Efficient sparse matrix-vector multiplication on CUDA. NVIDIA Technical Report, 2008.
- [BSW08] M. Belkin, J. Sun, and Y. Wang. Discrete Laplace operator on meshed surfaces. In *Proc. SOCG*, pages 278–287, 2008.
- [BU83] S. Bando and H. Urakawa. Generic properties of the eigenvalue of the Laplacian for compact Riemannian manifolds. *Tôhoku Math. J. (2)*, 35(2):155–172, 1983.
- [Che76] S.-Y. Cheng. Eigenfunctions and nodal sets. *Comment. Math. Helv.*, 51(1):43–55, 1976.
- [dGGV08] F. de Goes, S. Goldenstein, and L. Velho. A hierarchical segmentation of articulated bodies. *Comp. Graph. Forum*, 27:5:1349–1356, 2008.
- [GBAL09] K. Gebal, J.A. Bærentzen, H. Aanæs, and R. Larsen. Shape Analysis Using the Auto Diffusion Function. *Comp. Graph. Forum*, 28(5):1405–1413, 2009.
- [Gri99] A. Grigor'yan. Escape rate of brownian motion on riemannian manifolds. *Applicable Analysis*, 71(1):63–89, 1999.
- [Gri06] A. Grigor'yan. Heat kernels on weighted manifolds and applications. *Contemporary Mathematics*, 398:93–193, 2006.
- [HAWG08] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Comput. Graph. Forum (Proc. SGP 2008)*, 27(5):1449–1457, 2008.
- [Hsu02] E. Hsu. *Stochastic Analysis on Manifolds*. AMS, 2002.
- [JMS10] P.W. Jones, M. Maggioni, and R. Schul. Universal local parametrizations via heat kernels and eigenfunctions of the laplacian. *Ann. Acad. Scient. Fen.*, 35:1–44, 2010.
- [JZvK07] V. Jain, H. Zhang, and O. van Kaick. Non-rigid spectral correspondence of triangle meshes. *International Journal on Shape Modeling*, 13(1):101–124, 2007.
- [KMP07] M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. *ACM TOG*, 26(3 #64), 2007.
- [LF09] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH 2009)*, 28(3):72, 2009.
- [LSLCO05] Y. Lipman, O. Sorkine, D. Levin, and D. Cohen-Or. Linear rotation-invariant coordinates for meshes. *ACM Trans. Graph.*, 24(3):479–487, 2005.
- [LSP08] H. Li, R. W. Sumner, and M. Pauly. Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum*, 27(5), 2008.
- [LTSW09] R. Lasowski, A. Tevs, H.-P. Seidel, and M. Wand. A probabilistic framework for partial intrinsic symmetries in geometric data. In *Proc. ICCV*, 2009.
- [MA] D.M. Mount and S. Arya. ANN: A library for approximate nearest neighbor searching. <http://www.cs.umd.edu/~mount/ANN/>.
- [MDSB02] M. Meyer, M. Desbrun, P. Schröder, and A.H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and mathematics*, 3, 2002.
- [Mém09] F. Mémoi. Spectral Gromov-Wasserstein distances for shape matching. In *Proc. NORDIA*, 2009.
- [MHK⁺08] D. Mateus, R. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In *Proc. CVPR*, 2008.
- [MS05] F. Mémoi and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5(3):313–347, 2005.
- [OSG08] M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. *Comp. Graph. Forum*, 27(5):1341–1348, 2008.
- [Reu10] M. Reuter. Hierarchical shape segmentation and registration via topological features of laplace-beltrami eigenfunctions. *Proc. IJCV*, 89:2:287–308, 2010.
- [RWP06] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-Beltrami spectra as "Shape-DNA" of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006.
- [SGT98] A. Singh, D. Goldgof, and D. Terzopoulos, editors. *Deformable Models in Medical Image Analysis*. IEEE Press, 1998.
- [Sid98] R. B. Sidje. Expokit: a software package for computing matrix exponentials. *ACM Trans. Math. Softw.*, 24(1), 1998.
- [SOG09] J. Sun, M. Ovsjanikov, and L. Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Comp. Graph. Forum*, 28(5), 2009.
- [TBW⁺09] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. Isometric registration of ambiguous and partial data. In *Proc. CVPR*, 2009.
- [VBCGar] A. Vaxman, M. Ben-Chen, and C. Gotsman. A multi-resolution approach to heat kernels on discrete surfaces. *Proc. SIGGRAPH*, 2010 (to appear).
- [VBMP08] D. Vlastic, I. Baran, W. Matusik, and J. Popović. Articulated mesh animation from multi-view silhouettes. *ACM Trans. Graph.*, 27(3), 2008.
- [WJH⁺07] M. Wand, P. Jenke, Q. Huang, M. Bokeloh, L. Guibas, and A. Schilling. Reconstruction of deforming geometry from time-varying point clouds. In *Proc. SGP*, pages 49–58, 2007.
- [WR97] Haim J. Wolfson and Isidore Rigoutsos. Geometric hashing: An overview. *IEEE Comput. Sci. Eng.*, 4(4):10–21, 1997.
- [ZSCO⁺08] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. van Kaick, and A. Tagliasacchi. Deformation-driven shape correspondence. *Comp. Graph. Forum*, 27(5), 2008.

8. Appendix

Proof of Lemma 2.1 Consider a heat-kernel preserving map T . By Varadhan's Lemma, T must be distance preserving. To show that the map T is an isometry, we only need to prove that it is surjective. First note that the preservation of the eigenvalues, and Weyl's local law imply that the two manifolds are of the same dimension. The distance-preserving property further implies that for almost every y in M , the Jacobian $J_y T$ is well-defined and orthogonal. This allows us to use the change-of-variable formula (which holds for maps that are differentiable almost everywhere) for the integral of the function $y \in N \mapsto k_t^N(x, y)$ for a fixed x and $t > 0$, under the map T :

$$\int_M |J_y T| k_t^N(T(x), T(y)) dy = \int_{T(M)} k_t^N(x, y) dy.$$

The orthogonality of $J_y T$ and the heat-kernel invariance of T imply that the left-hand side of this equation is 1. Since the heat kernel has full support for any $t > 0$, $T(M)$ has full measure. By continuity of T and compactness of M , this implies that $T(M) = N$. \square

Proof of Lemma 3.2: Suppose $j = \min_j$ s.t. $a_j \neq 0$. Then:

$$|a_j| = |e^{\lambda_j t} a(t) - a_j| \leq c(t) := \sum_{i>j} |a_i| e^{(\lambda_i - \lambda_j)t}$$

By absolute convergence, $c(t)$ is well defined and decreasing. Set $\mu_i = \lambda_i - \lambda_j$, and $t = (h+1)\mu_{j+1}^{-1}$, for any $h > 0$. Then,

$$0 \leq c(t) = \sum_{i>j} |a_i| e^{-(h+1)\frac{\mu_i}{\mu_{j+1}}} \leq e^{-h} c(\mu_{j+1}^{-1})$$

Thus, the limit of c is zero at infinity, and $a_j = 0$. \square