# Iterative Closest Spectral Kernel Maps

Alon Shtern and Ron Kimmel
*Computer Science Department*
*Technion - Israel Institute of Technologies*
*Haifa, Israel*
*ashtern@tx.technion.ac.il, ron@cs.technion.ac.il*

***Abstract*—An important operation in geometry processing is finding the correspondences between pairs of shapes. Measures of dissimilarity between surfaces, has been found to be highly useful for nonrigid shape comparison. Here, we analyze the applicability of the spectral kernel distance, for solving the shape matching problem. To align the spectral kernels, we introduce the iterative closest spectral kernel maps (ICSKM) algorithm. The ICSKM algorithm farther extends the *iterative closest point* algorithm to the class of deformable shapes. The proposed method achieves state-of-the-art results on the Princeton isometric shape matching protocol applied, as usual, to the TOSCA and SCAPE benchmarks.**

***Keywords*-shape matching; Laplace-Beltrami operator; correspondence**

## I. Introduction

Correspondence detection between pairs of shapes lies at the heart of many operations in the field of geometry processing. The problem of acquiring correspondence between rigid shapes has been widely addressed in the literature. As for non-rigid shapes, this problem remains difficult even when the space of deformations is narrowed to nearly isometric surfaces, which approximately preserve the geodesic distances between corresponding points on each shape.

A common approach for shape matching is to define a measure of dissimilarity between shapes modeled as 2-manifolds. The well-established Gromov-Hausdorff distance measures the maximum geodesic discrepancy between pairs of corresponding points of the two given shapes [18]. The point-wise map can be inferred to as a byproduct of the evaluation of the Gromov-Hausdorff distance. This approach was embraced by the *Generalized Multi-Dimensional Scaling* (GMDS) framework [5]. Within the Gromov-Hausdorff framework, Bronstein et al. [7] suggested replacing the geodesic distance by the diffusion distance [8], exploiting the apparent stability of diffusion distances to local changes in the topology of the shape. Despite its generality and theoretical beauty, it has been a challenge to apply the Gromov-Hausdorff framework in a straightforward manner to shape matching, mainly due to its intrinsically combinatorial nature.

Kasue and Kumura [11] extended the Gromov-Hausdorff distance framework to the family of spectral methods. The *spectral kernel distance* was constructed by replacing the metric defined on the manifolds with the heat kernel. The heat kernel provides a natural notion of scale, which is useful for multi-scale shape comparison. Recently, Mémoli [17] introduced the *spectral Gromov-Wasserstein distance*, applying the theory of mass transportation. The spectral Gromov Wasserstein distance via the comparison of heat kernels satisfies all properties of a metric on the class of isometric manifolds.

The evaluation of the spectral kernel distance between two nearly isometric surfaces should be capable of discovering the mapping between them. Alas, the task is not straightforward, due to model impairments and the combinatorial nature of the problem. Therefore, to achieve highly accurate and dense correspondence, we need to make adaptations to this distance measure and design an efficient and robust optimization algorithm.

### A. Contribution

Our main observation is that the alignment of the spectral kernels and the evaluation of the spectral kernel distance between two shapes can be achieved by extending the well established Iterative Closest Point (ICP) algorithm [4], [28] to the class of nonrigid shapes. The classical ICP algorithm refines the correspondence between rigid shapes embedded in the three dimensional Euclidean space. The key idea is simple. Given an initial map between the shapes, find the best rotation and translation that aligns the shapes, apply it and calculate new correspondence by the nearest neighbor algorithm.

As for nonrigid shapes, a similar idea was presented by the iterative post-process refinement algorithm [19]. Instead of aligning the shapes in the three dimensional Euclidean domain, this method estimates the transformation that best fits the shapes in the spectral domain. Given an initial map from shape $X$ to shape $Y$, one linear constraint is generated for each point $x \in X$, and the least squares method is used to infer the transformation matrix.

The proposed *iterative closest spectral kernel maps* (ICSKM) algorithm extends this idea by finding the transformation that best matches the respective spectral kernels $K(x, x')$ and $\tilde{K}(y, y')$ of the shapes $X$ and $Y$. Now, each pair of points $x, x' \in X$ generates a linear constraint by

including its normalized kernel relation $K(x,x')/K(x,x)$. The two dimensional information, effectively improves the refinement procedure. The optimization problem is solved by the least squares method with Tikhonov regularization [10], [27]. The algorithm is shown to be robust, flexible and easy to implement. It can be used efficiently as a refinement procedure of rough or sparse correspondence detection methods. The main advantage of the ICSKM algorithm over existing methods is in the combination of the iterative post-process refinement algorithm with the two dimensional constraints of the spectral kernel, resulting in highly accurate correspondence maps.

## II. Related work

### A. spectral kernel distance

The heat kernel $K_t(x,x')$ is defined as the solution of the heat equation $\dfrac{\partial u}{\partial t} = \Delta u$, with a point heat source at $x \in X$, measured at point $x' \in X$ after time $t > 0$, where $\Delta$ denotes the *Laplace-Beltrami* (LB) operator.

Kasue and Kumura [11] defined the metric $d(X,Y)$ between the Riemannian manifolds $X$ and $Y$ by comparing their respective heat kernels

$$d(X,Y) \equiv \inf_{\substack{\varphi: X \mapsto Y \\ \psi: Y \mapsto X}} \max(\mathrm{dis}(\varphi), \mathrm{dis}(\psi)), \tag{1}$$

taking the supremum of kernel distortion for all $t > 0$

$$\begin{aligned} \mathrm{dis}(\varphi) &\equiv \sup_{x,x' \in X, t>0} u(t) d_t(x,x',\varphi(x),\varphi(x')), \\ \mathrm{dis}(\psi) &\equiv \sup_{y,y' \in Y, t>0} u(t) d_t(\psi(y),\psi(y'),y,y'), \end{aligned}$$

where $d_t(x,x',y,y')$ measures the absolute discrepancy between the heat kernels $K(x,x')$ and $\tilde{K}(y,y')$

$$|\mathrm{Vol}(X) K_t(x,x') - \mathrm{Vol}(Y) \tilde{K}_t(y,y')|.$$

Vol(X) and Vol(Y) are the volumes of $X$ and $Y$, respectively. The function $u(t) \equiv e^{-(t+1/t)}$ is used to normalize the kernels for different values of $t$, and make sure that it will not blow up as $t \to 0$. We denote $d(X,Y)$ as the *spectral kernel distance*. The spectral kernel distance is a metric between isometry classes of Riemannian manifolds, which means, in particular, that two manifolds are at zero distance if and only if they are isometric.

In practice, a more tractable $L_2$ version of Eq. (1) can be optimized by finding the map $\varphi : X \mapsto Y$ that best aligns the spectral kernels of two shapes for a fixed time $t$. In the discrete setting, the spectral kernel distortion can be formalized as

$$\min_{\varphi: X \mapsto Y} \sum_{x,x' \in X} |K_t(x,x') - \tilde{K}_t(\varphi(x),\varphi(x'))|^2. \tag{2}$$

### B. Post-process iterative refinement algorithm

The post-process iterative refinement algorithm [19] takes as input an initial map, iteratively finds the transformation matrix between the spectral bases of the two compared shapes, and outputs a dense correspondence between the shapes. Here, we use the first $n$ Laplace-Beltrami eigenfunctions as the spectral basis [3], [14], [25], [26]. The eigendecomposition of the LB operator consists of non-negative eigenvalues $0 = \lambda_0 < \lambda_1 < \cdots < \lambda_i < \cdots$, with corresponding eigenfunctions $\Phi \equiv \{\phi_0, \phi_1, \cdots, \phi_i, \cdots\}$ that forms an orthonormal basis, which is well suited for representing near isometric shapes [1], [19]. In this case, the post-process iterative refinement algorithm is similar to the well known *Iterative Closest Point* (ICP) [4], [28] in $n$ dimensions, except that it is performed in the natural spectral domain, rather than the standard Euclidean space.

Let $\varphi : X \mapsto Y$ be a bijective mapping between shapes $X$ and $Y$. If we are given a scalar function $f : X \mapsto \mathbb{R}$, then, we can obtain a corresponding function $g : Y \mapsto \mathbb{R}$ by the composition $g = f \circ \varphi^{-1}$. Given the bases $\Phi$ and $\tilde{\Phi}$ on the shapes $X$ and $Y$, respectively, we can represent $f$ as a row vector $\boldsymbol{a}$ with coefficients $a_i$, and equivalently, $g$ as a row vector $\boldsymbol{b}$ with coefficients $b_i$. It is easy to show that we can write a linear transformation $\boldsymbol{a} = \boldsymbol{b}C$, where the transformation matrix $C$ is independent of $f$ and is completely determined by the bases $\Phi$, $\tilde{\Phi}$ and the map $\varphi$.

Now, suppose we have point-to-point correspondences, such that each point $x \in X$ corresponds to some point $y \in Y$ by the mapping $y = \varphi(x)$. In this case, the delta function $\delta_x$ at point $x \in X$ corresponds to the delta function $\tilde{\delta}_y$ at point $y = \varphi(x)$. We can represent the delta function $\delta_x$ in the basis $\Phi$ by

$$\boldsymbol{a}_x = \Phi(x) = (\phi_1(x), \phi_2(x), \ldots, \phi_i(x), \ldots).$$

Equivalently, the function $\tilde{\delta}_y$ can be represented in the basis $\tilde{\Phi}$ as

$$\boldsymbol{b}_y = \tilde{\Phi}(y) = (\tilde{\phi}_1(y), \tilde{\phi}_2(y), \ldots, \tilde{\phi}_i(y), \ldots).$$

Then, we can construct the function preservation constraints $A = BC$, where the corresponding matrices $A$ and $B$ are built by stacking the row vectors $\boldsymbol{a}_x$ and $\boldsymbol{b}_y$, respectively. Therefore, at every iteration of the refinement procedure, we can infer the transformation matrix $C$ from previous correspondences by solving $A = BC$ with the least square method. Then, a new map can be found by searching for the point $y \in Y$, such that the row vector $\tilde{\Phi}(y)C$ is the closest to $\Phi(x)$.

## III. Iterative closest spectral kernel maps

Motivated by the definition of the spectral kernel distortion of Eq. (2), we wish to find the map $\varphi : X \mapsto Y$ that aligns the compatible spectral kernels, $K(x,x')$ and $\tilde{K}(y,y')$. We adopt a similar approach to the post-process

iterative refinement algorithm, by constructing corresponding functions over the two shapes. The trivial functions that represent point-to-point correspondence are the delta functions. The key idea is to impose the spectral kernel constraints on these delta functions. Accordingly, if the point $x \in X$ maps to $y = \varphi(x)$ and the point $x' \in X$ maps to $y' = \varphi(x')$, then, the function

$$f_{x,x'} \quad = \quad (K(x,x') \,/\, |K(x,x)|)\delta_x,$$

should correspond to

$$g_{y,y'} \quad = \quad (\tilde{K}(y,y') \,/\, |\tilde{K}(y,y)|)\tilde{\delta}_y.$$

We point out that $x$, $x'$ are constant parameters that define the function $f_{x,x'}$. As seen in Section II-B, the LB basis representation of the delta function $\delta_x$ at a point $x \in X$ is simply $\Phi(x)$. Therefore, the function $f_{x,x'}$ in the basis $\Phi$, and eqivalently the function $g_{y,y'}$ in the basis $\tilde{\Phi}$, can be represented by

$$\begin{aligned} \boldsymbol{a}_{x,x'} &= (K(x,x') \,/\, |K(x,x)|)\Phi(x), \\ \boldsymbol{b}_{y,y'} &= (\tilde{K}(y,y') \,/\, |\tilde{K}(y,y)|)\tilde{\Phi}(y). \end{aligned}$$

In this case, we can construct the corresponding matrices $A$, $B$ by stacking the row vectors $\boldsymbol{a}_{x,x'}$ and $\boldsymbol{b}_{y,y'}$, respectively. Notice that we normalize the kernels, so that $K(x,x') \,/\, |K(x,x)| = 1$, $\forall x = x'$.

By recalling that for nearly isometric shapes, the correspondence we are looking for should be represented by a nearly-diagonal $C$ [13], we can submit an element-wise off-diagonal penalty $W$ and formulate the following problem

$$\underset{C}{\operatorname{argmin}} \left\| A - BC \right\|_F^2 + \beta \left\| W \odot C \right\|_F^2, \tag{3}$$

where $\beta$ is a tuning parameter. The symbol $\odot$ represents the matrix element-wise multiplication operation. The matrix $W$ is chosen, such that as $(i,j)$ is located farther from the diagonal of the matrix $W$, the element-wise penalty $W_{i,j}$ increases. The minimization of Eq. (3) can be obtained separately for each column of $C$ by the least squares method, with Tikhonov regularization [10], [27].

The iterative closest spectral kernel maps algorithm is summarized in Algorithm 1. For a given initial correspondence $\hat{\varphi}_0(x)$, $\forall\, x \in X$, the algorithm provides the transformation matrix $C$ computed in Step 3, and the point-wise map $\hat{\varphi}(x)$ found in Step 4, which can be used to approximate the spectral kernel distortion of Eq. (2).

As an additional option, one can discard correspondences before estimating the transformation matrix in Step 3. For example, corresponding triangles with flipped orientation are expected to be outliers. In that case, it is advisable to filter out correspondences that belong to such inversely oriented coupled triangles.

The ICSKM algorithm can be viewed as a generalization of the post-process iterative refinement algorithm. This is noticed by setting the kernel $K(x,x')$ to be the heat kernel $K_t(x,x')$. In that case, as $t \to 0$ the normalized kernel $K_t(x,x') \,/\, |K_t(x,x)| \to 0$ for $x \neq x'$, and the only constraints that remain are $\boldsymbol{a}_{x,x} \to \Phi(x)$ and $\boldsymbol{b}_{y,y} \to \tilde{\Phi}(y)$.

**Algorithm 1 :** ITERATIVE CLOSEST SPECTRAL KERNEL MAPS

**for** $\ell = 1$ to $L$ **do**

1) Calculate the spectral coefficients

$$\begin{aligned} \boldsymbol{a}_{x,x'} &= (K(x,x') \,/\, |K(x,x)|)\Phi(x), \\ \boldsymbol{b}_{y,y'} &= (\tilde{K}(y,y') \,/\, |\tilde{K}(y,y)|)\tilde{\Phi}(y), \end{aligned}$$

for $x, x' \in X$, $y = \hat{\varphi}_{\ell-1}(x)$, $y' = \hat{\varphi}_{\ell-1}(y')$,

using the correspondence $\hat{\varphi}_{\ell-1}$ provided by the previous iteration.

2) Compose the constraint matrices $A_\ell$, $B_\ell$ by stacking the row vectors $\boldsymbol{a}_{x,x'}$, $\boldsymbol{b}_{y,y'}$ respectively.
3) Find the optimal transformation matrix $C_\ell$ that minimizes

$$\left\| A_\ell - B_\ell C_\ell \right\|_F^2 + \beta \left\| W \odot C_\ell \right\|_F^2.$$

4) For each point $x \in X$, find new map $\hat{\varphi}_\ell(x)$ by searching for the point $y \in Y$, that minimizes the Euclidean distance between the row vectors $\tilde{\Phi}(y)C_\ell$ and $\Phi(x)$, applying

$$\hat{\varphi}_\ell(x) = \underset{y \in Y}{\operatorname{argmin}} \left\| \Phi(x) - \tilde{\Phi}(y)C_\ell \right\|_2.$$

**end for**

*Implementation*

In all our experiments we used the same choice of parameters. In general, we chose our parameters for achieving the most accurate results in a reasonable time. To that end, we used $n = 200$ eigenfunctions of the LB operator.

Our empirical evidence suggests that the GPS kernel [21], [23], that is,

$$K(x,x') \quad = \quad \sum_i \frac{1}{\lambda_i} \phi_i(x)\phi_i(x'),$$

provides superior qualities for correspondence detection, compared to other kernels we tested. The number of iteration has been set to $L = 45$. At each iteration, a subset of 2000 points are matched. The off-diagonal penalty $W_{i,j}^2 = \dfrac{|\tilde{\lambda}_i - \lambda_j|}{\lambda_j} U_i$ was set to be proportional to the difference of the eigenvalues $\tilde{\lambda}_i$ and $\lambda_i$ that correspond to $\tilde{\phi}_i$ and $\phi_i$, and scaled by the $i_{\text{th}}$ entry of $U = \text{diag}(B^T B)$. The tuning parameter $\beta$ was set to 0.1.

The system was implemented in MATLAB. All the experiments were executed on a 3.00 GHz Intel Core i7 machine with 32GB RAM. Run-times for pairs of shapes of various sizes from the TOSCA dataset are shown in Table I.

Table I
RUN-TIMES (IN SECONDS) OF THE PROPOSED METHOD, EVALUATED ON SHAPES FROM THE TOSCA DATASET.

| # of vertices | 4344 | 19248 | 25290 | 45659 | 52565 |
|---|---|---|---|---|---|
| $n = 50$ eigenfunctions | | | | | |
| runtime | 32 | 39 | 41 | 51 | 56 |
| $n = 100$ eigenfunctions | | | | | |
| runtime | 50 | 59 | 62 | 74 | 80 |
| $n = 200$ eigenfunctions | | | | | |
| runtime | 115 | 135 | 140 | 153 | 165 |

## IV. RESULTS

We tested the proposed method on pairs of shapes represented by triangulated meshes from both the TOSCA database [6] and the SCAPE database [2]. The TOSCA dataset contains densely sampled synthetic human and animal surfaces, divided into several classes with given ground-truth point-to-point correspondences between the shapes within each class. The SCAPE dataset contains scans of real human bodies in different poses. We compare our results to several correspondence detection methods.

- Iterative Closest Spectral Kernel Maps - the method proposed in this paper. Initial coarse map is found by comparing the first few matched eigenfunctions of the LB operator [24]. Another option is to use a small number of landmark points.
- Functional Maps + Blended (TOSCA only) - the functional maps based post-process iterative refinement algorithm. We use the results shown in [19]. There, the post-process procedure refines the correspondence provided by the Blended method [12].
- Blended - the method proposed by Kim et al. that uses a weighted combination of isometric maps [12].
- Möbius Voting - the method proposed by Lipman et al. counts *votes* on the conformal Möbius transformations [15].
- Permuted Sparse Coding + MSER (SCAPE only) - the approach proposed by Pokrass et al. finds correspondence by using methods from the field of sparse modeling [20] . We note that this method depends on the ability to detect repeatable regions between shapes. There, maximally stable extremal regions (MSER) are used as a preprocessing step [16].

Fig. 1 compares the ICSKM algorithm with existing methods on the TOSCA benchmark, using the evaluation protocol proposed in [12]. The distortion curves describe the percentage of surface points falling within a relative geodesic distance from what is assumed to be their true locations. For each shape, the geodesic distance is normalized by the square root of the shape's area. It is evident from the benchmark that the proposed method significantly outperforms existing ones.

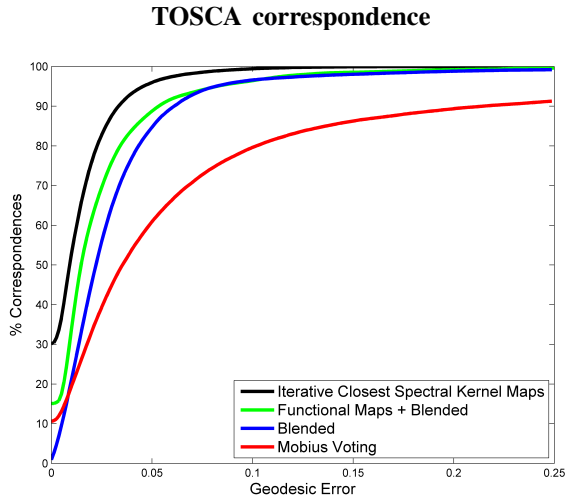


Figure 1. Evaluation of the iterative spectral kernel maps algorithm applied to shapes from the TOSCA database, using the protocol of [12].

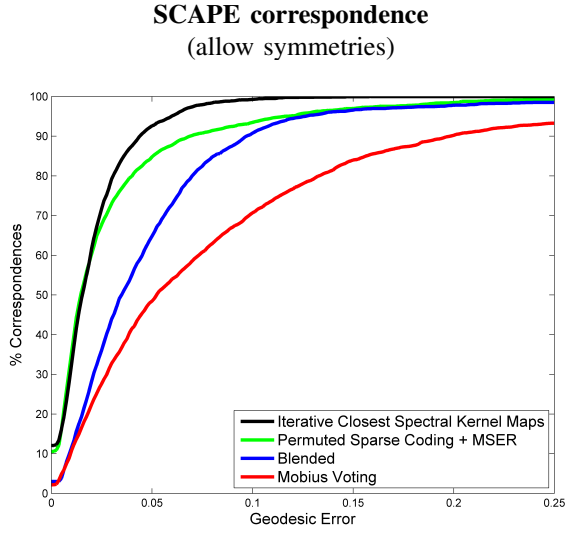


Figure 2. Evaluation of the iterative closest spectral kernel maps algorithm applied to shapes from the SCAPE database, using the protocol of [12] with allowed symmetries.

Fig. 2 compares the proposed correspondence algorithm with existing methods on the SCAPE database, again using the evaluation protocol proposed in [12], allowing symmetric flip for a selected number of feature points. Remark: In the evaluation, the correct symmetry is automatically chosen for the shape as a whole.

Table II displays the percentage of correspondences that fall within different values of relative geodesic distances. It is interesting to focus on large geodesic errors. Unlike other methods, in the proposed approach only one of 200 points has a geodesic error larger than 0.1.

Table II
PERCENTAGE OF SURFACE POINTS FALLING WITHIN A RELATIVE GEODESIC ERROR FOR DIFFERENT METHODS (TOSCA).

| Geodesic circle | 0.025 | 0.050 | 0.100 | 0.150 |
|---|---|---|---|---|
| **ICSKM** | **82.5** | **95.9** | **99.5** | **99.9** |
| F. Maps + Blended | 69.5 | 88.7 | 96.4 | 98.5 |
| Blended | 55.9 | 84.7 | 96.6 | 98.0 |
| Möbius Voting | 39.3 | 60.9 | 79.6 | 86.2 |

**ICSKM landmark points initialization**
(TOSCA+SCAPE)

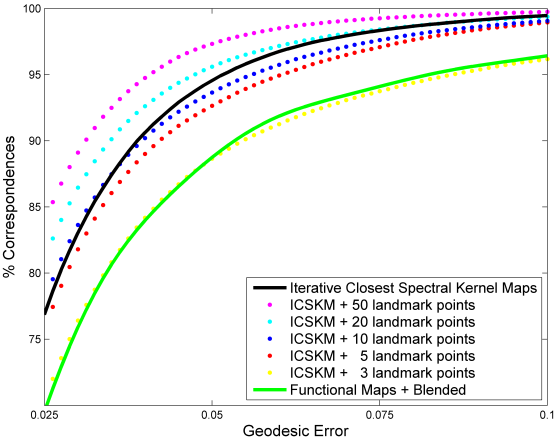


Figure 3. Evaluation of the iterative closest spectral kernel maps algorithm initialized by landmark points.

We continue investigating the refinement capabilities of the ICSKM algorithm. For that goal, we provide the algorithm with 3 to 50 landmark points, that were randomly selected from the ground-truth mapping. Fig. 3 compares the dense maps produced by the algorithm with these initial constraints. Observe that with just five landmark points, the algorithm outperforms previous state-of-the-art methods. We have also applied the ICSKM algorithm to non-isometric shapes taken from the TOSCA database. Fig. 4, displays the distortion curves for different pair of classes. For each class, we used the manually selected landmark points specified in [12]. Half of these points were used to provide the algorithm with initial correspondence. The rest of the points were used to evaluate the geodesic error. Fig. 5 demonstrates how the mapping produced by the ICSKM algorithm initialized with 7 landmark points, transfers the texture from a wolf to a cat and from a dog to a horse

Finally, we illustrate how the proposed method is able to find the intrinsic reflective symmetry axis of nonrigid shapes. Intrinsic symmetry detection can be viewed as finding correspondence from a shape to itself [22]. Following this approach, we search for a self-map with flipped orientation. In Fig. 6 we visualize the distance between a point and its image for several shapes from the TOSCA database.

**Non-isometric shapes correspondence**

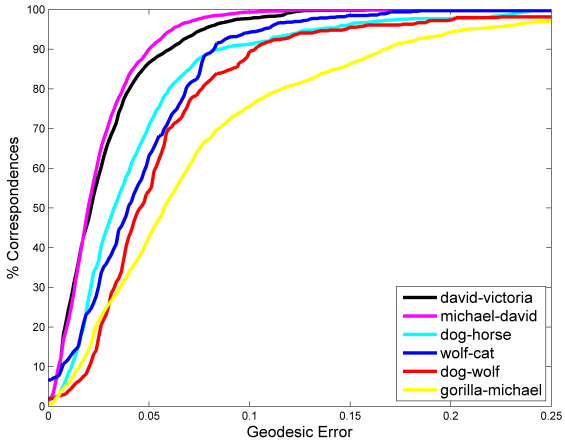


Figure 4. Evaluation of the ICSKM algorithm applied to non-isometric shapes from the TOSCA database. For the primates and animals categories, the algorithm is provided with pairs of 18 and 11 landmark points, respectively. The distortion curves are evaluated by calculating the geodesic error of 18 corresponding points for the primates category, and 10 corresponding points for the animals category.

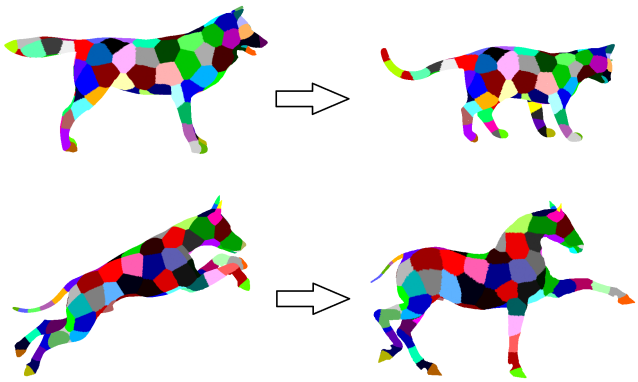

Figure 5. Texture mapping of non-isometric shapes. The textures of wolf and dog shapes were transfered to cat and horse shapes, respectively. The ICSKM algorithm was initialized by 7 landmark points selected using the farthest point strategy [9].

## V. CONCLUSIONS

A new method for correspondence detection between nonrigid shapes was introduced. The method is based on the evaluation of the spectral kernel distance, optimized by an ICP based approach in the spectral domain. We have demonstrated the effectiveness of the ICSKM algorithm by achieving state-of-the-art results on shape matching benchmarks. In the future, we intend to apply the ICSKM algorithm for other purposes, such as registration of rigid shapes, matching stereo images, and comparing deformable shapes with texture, and to study the potential and the limitations of the proposed approach for refining correspondences between shapes with topological noise or partially missing data.

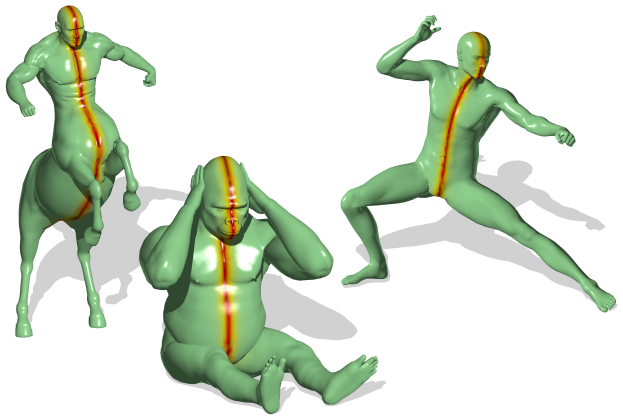

Figure 6. Symmetry axis of several shapes from the TOSCA database.

## ACKNOWLEDGMENT

The authors would like to thank Yonathan Aflalo, Anastasia Dubrovina and Matan Sela for stimulating discussions throughout this research. This work has been supported by grant agreement no. 267414 of the European Community's FP7-ERC program.

## REFERENCES

[1] Y. Aflalo and R. Kimmel. Spectral multidimensional scaling. *Proceedings of the National Academy of Sciences*, 110(45):18052–18057, 2013.

[2] D. Anguelov, P. Srinivasan, H.-C. Pang, D. Koller, S. Thrun, and J. Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. *Advances in neural information processing systems*, 17:33–40, 2005.

[3] E. Beltrami. Ricerche di analisi applicata alla geometria. *Giornale di Matematiche*, 2:355–375, 1864.

[4] P. J. Besl and N. D. McKay. Method for registration of 3-D shapes. In *Robotics-DL tentative*, pages 586–606. International Society for Optics and Photonics, 1992.

[5] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences of the United States of America*, 103(5):1168–1172, 2006.

[6] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer, 2008.

[7] A. M. Bronstein, M. M. Bronstein, M. Mahmoudi, R. Kimmel, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision*, 89(2):266–286, 2010.

[8] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the National Academy of Sciences of the United States of America*, 102(21):7426–7431, 2005.

[9] D. S. Hochbaum and D. B. Shmoys. A best possible heuristic for the k-center problem. *Mathematics of operations research*, 10(2):180–184, 1985.

[10] A. E. Hoerl and R. W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.

[11] A. Kasue and H. Kumura. Spectral convergence of Riemannian manifolds. *Tohoku mathematical journal. Second series*, 46(2):147–179, 1994.

[12] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. In *ACM Transactions on Graphics (TOG)*, volume 30, page 79. ACM, 2011.

[13] A. Kovnatsky, A. M. Bronstein, M. M. Bronstein, K. Glashoff, and R. Kimmel. Coupled quasi-harmonic bases. *Computer Graphics Forum*, 2013.

[14] B. Lévy. Laplace-Beltrami eigenfunctions towards an algorithm that understands geometry. In *Shape Modeling and Applications, 2006. SMI 2006. IEEE International Conference on*, pages 13–13. IEEE, 2006.

[15] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. In *ACM Transactions on Graphics (TOG)*, volume 28, page 72. ACM, 2009.

[16] R. Litman, A. M. Bronstein, and M. M. Bronstein. Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics*, 35(3):549–560, 2011.

[17] F. Mémoli. Spectral Gromov-Wasserstein distances for shape matching. In *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, pages 256–263. IEEE, 2009.

[18] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5(3):313–347, 2005.

[19] M. Ovsjanikov, M. Ben Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30, 2012.

[20] J. Pokrass, A. M. Bronstein, M. M. Bronstein, P. Sprechmann, and G. Sapiro. Sparse modeling of intrinsic correspondences. *Eurographics Computer Graphics Forum*, 2013.

[21] H. Qiu and E. R. Hancock. Clustering and embedding using commute times. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(11):1873–1890, 2007.

[22] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Full and partial symmetries of non-rigid shapes. *International journal of computer vision*, 89(1):18–39, 2010.

[23] R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 225–233. Eurographics Association, 2007.

[24] A. Shtern and R. Kimmel. Matching the LBO eigenspace of non-rigid shapes via high order statistics. *Axioms*, 3(3):300–319, 2014.

[25] N. Sochen, R. Kimmel, and R. Malladi. A general framework for low level vision. *Image Processing, IEEE Transactions on*, 7(3):310–318, 1998.

[26] G. Taubin. A signal processing approach to fair surface design. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 351–358. ACM, 1995.

[27] A. Tikhonov. Solution of incorrectly formulated problems and the regularization method. In *Soviet Math. Dokl.*, volume 5, pages 1035–1038, 1963.

[28] C. Yang and G. Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992.