Detail preserving non-rigid shape correspondences

Manika Bindal
Goa University
Taleigao, Goa
India 403206
manika.bindal@gmail.com

Venkatesh Kamat
IIT Goa
Farmagudi, Goa
India 403401
vvkamat@iitgoa.ac.in

ABSTRACT

Understanding shapes is an organic process for us (humans) as this is fundamental to our interaction with the surrounding world. However, it is daunting for the machines. Any shape analysis task, particularly non-rigid shape correspondence is challenging due to the ever-increasing resolution of datasets available. Shape Correspondence refers to finding a mapping among various shape elements. The functional map framework deals with this problem efficiently by not processing the shapes directly but rather specifying an additional structure on each shape and then performing analysis in the spectral domain of the shapes. To determine the domain, the Laplace-Beltrami operator has been utilized generally due to its capability of capturing the global geometry of the shape. However, it tends to smoothen out high-frequency features of shape, which results in failure to capture fine details and sharp features of shape for the analysis. To capture such high-frequency sharp features of the shape, this work proposes to utilize a Hamiltonian operator with gaussian curvature as an intrinsic potential function to identify the domain. Computationally it is defined at no additional cost, keeps global structural information of the shape intact and preserves sharp details of the shape in order to compute a better point-to-point correspondence map between shapes.

Keywords

shape matching, shape correspondence, functional maps

1 INTRODUCTION

Shapes in computational context refers to digital representation of any real world object such as humans, chairs, etc. These digital representations can be meshes, point clouds or voxel grids. With ever increasing technological advancements, the accuracy with which these digital representations are being captured has transformed the field of shape analysis. Particularly, shape matching is quite an interesting area enticing researchers across multiple domains from Computer Graphics, Image processing, Geometry Processing and Computer Vision. A sub-area focusses on the fundamental task of computing shape correspondences, where rather than just specifying if two shapes match, a mapping is also desired between various elements of given shapes. Major applications constitute object reconstruction, attribute transfer, Statistical modelling, Shape Interpolation and morphing, etc.

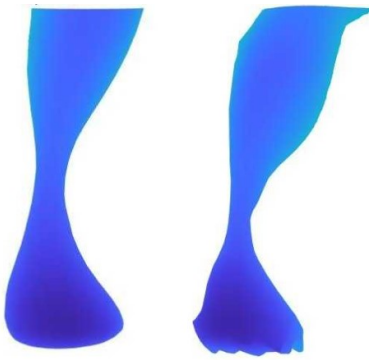

Figure 1: Reconstructed hand from human mesh via (left) 150 Laplace-Beltrami eigenfunctions (right); 150 Hamiltonian eigenfunctions

Based on how shapes can deform, varied approaches have been suggested [VKZHCO11]. Rigid shapes undergo transformations that preserve extrinsic features i.e. euclidean distances remain intact while non-rigid shapes deform anyhow [BBK07]. Rigid deformation tends to transform the shape without changing its geometry or topology via rotation or translation. Non-rigid involves changing in geometry as well as topology via stretching or bending. Finding correspondences for rigid shapes has plethora of efficient solutions. However, due to the vast space of deformations for non-rigid shapes, it is an interesting area to work. Another con-

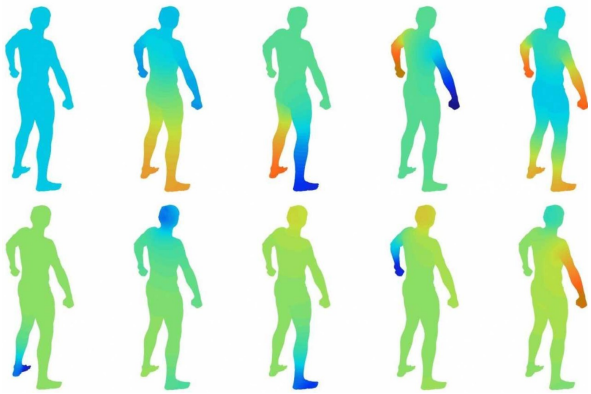

Figure 2: Laplace-Beltrami Eigenfunctions (top row); Hamiltonian Eigenfunctions (bottom row) arranged according to ascending order of eigenvalues respectively

sideration would be to establish a mapping either between all the elements of the shape or a partial subset of the elements of shape leading to determining full and partial correspondences respectively [RCB+17]. The field's vastness and complexity are evident from these aspects.

In this work, the analysis is restricted to Isometric deformations of triangulated meshes where distances along the surface are invariant, when deformed without any stretching or tearing of the surface. The proposed research aims to investigate whether utilizing the sharp features of two given triangle meshes can improve the point-to-point correspondence mapping between the respective shape by using a functional framework.

2 RELATED WORK

Shape matching can be broadly categorized into two approaches: geometric transformation-based and descriptor-based. Geometric transformation-based approaches involve finding the geometric transformation that aligns two shapes and then establishing correspondences based on proximity [GMGP05]. This approach typically involves minimizing the distance between corresponding points or features in the two shapes. Descriptor-based approaches, on the other hand, do not require the shapes to be aligned. Instead, correspondences are established based on similarities between shape descriptors or feature vectors [CCFM08]. This approach is often used when the shapes have different topologies or geometry. A hybrid approach can also be used, where the alignment and computation of correspondences are alternated and iteratively improved [Ale02]. This approach can be effective in cases where the shapes undergo non-rigid deformations, where the geometric transformation-based approach may not be suitable.

For rigid shapes, *Iterative Closest Point*, ICP algorithm [BM92] is the celebrated technique that iteratively finds shape correspondence by aligning the shapes first by finding the optimal geometric transformation consisting of rotations and translations, then by utilizing Nearest Neighbour approach to find the closest points for the computed alignment [RL01]. Other methods for rigid shapes are surveyed in [BSBW14].

For non-rigid shapes, to allow any kind of mathematical analysis is to restrict the deformation to an isometry, wherein the distances between pair of points are preserved along the surface, while the shape is undergoing isometric deformation. Computing a map for high quality meshes with large number of vertices is computationally quite intensive. To cater different metric spaces, [EK03] introduced the idea of isometrically embedding the shapes into a canonical domain to allow any kind of shape analysis task. Shapes can also be embedded spectrally [ZVKD10] by utilizing eigenmodes of linear operators defined on the shape.

The current work is inspired by functional map framework [OBCS+12] developed to efficiently compute a mapping from the function space of one shape to another and subsequently determining point-to-point correspondences between them. Shape features are projected onto functional basis to reduce computation time during analysis and capture geometry along with other properties of shape effectively. Further, a map is computed between shapes by setting up an optimization problem, which then is refined to obtain point-to-point correspondences. Refer section 3 for more details on Functional Maps.

Identification of good basis is crucial for the functional map framework to output point-to-point correspondences. [OBCS+12] proposed to use Laplace-Beltrami eigenfunctions as functional basis. To capture high-frequency information on the shape, [NVT+14] proposed *compressed manifold modes* that are sparse basis with local support upto sign flip and ordering. By explicitly controlling the region of localization, [MRCB18] introduced *localized manifold harmonics* (LMH). These properties quickly become challenging to adopt when dealing with multiple meshes together as in case of shape correspondence, pose transfer, etc. Though [KBB+13] and [EKB+15] take into account multiple meshes and obtained basis incorporating the geometric information of all the meshes, it still fails to capture high-frequency information of the shapes.

3 BACKGROUND

Functional map is a promising framework for non-rigid shape matching that finds a map between the two function spaces defined on shapes rather than a map between shapes directly. Consider two shapes $\mathscr{M}$ and $\mathscr{N}$ which are represented as triangular meshes. A point-to-point map between $\mathscr{M}$ and $\mathscr{N}$ is given as $\mathscr{T}: \mathscr{M} \to \mathscr{N}$, where for any point $p \in \mathscr{M} \implies \mathscr{T}(p) \in \mathscr{N}$. If the

number of points are same on both the shapes, then a bijection is desired where $\mathscr{T}^{-1}$ exists.

3.1 Functional Maps

Functional map framework works by defining spaces of scalar valued functions $\mathscr{F}(\mathscr{M},\mathbb{R})$ and $\mathscr{F}(\mathscr{N},\mathbb{R})$ on shapes $\mathscr{M}$ and $\mathscr{N}$ respectively. It aims at computing a linear mapping between these function spaces $\mathscr{T}_F\colon \mathscr{F}(\mathscr{M},\mathbb{R}) \to \mathscr{F}(\mathscr{N},\mathbb{R})$. Map $\mathscr{T}_F$ which associate values of $f\colon \mathscr{M} \to \mathbb{R}$ and $g\colon \mathscr{N} \to \mathbb{R}$ can be represented as a matrix $\mathbb{C} \in \mathbb{R}^{k_1 \times k_2}$ with $\phi_{\mathscr{M}}$ and $\phi_{\mathscr{N}}$ are respective basis such that $|\phi_{\mathscr{M}}| = k_1$ and $|\phi_{\mathscr{N}}| = k_2$. The framework pipeline consists of following steps:

1 For each shape compute invariant feature descriptors say **F** and **G** with respect to isometric deformation

2 Choose basis $\phi_{\mathscr{M}}$ and $\phi_{\mathscr{N}}$ for both the shapes

3 Create function preservation constraints by projecting feature descriptors **F** and **G**, computed in step 1 onto respective basis as **A** and **B**

3 Set up other constraints like operator commutativity or regularization constraint

4 Compute optimal functional map **C** by minimizing the following energy:

$$E(\mathbf{C}) = ||\mathbf{CA} - \mathbf{B}||^2 + \|\mathscr{S}_F^{\mathscr{N}}\mathbf{C} - \mathbf{C}\mathscr{S}_F^{\mathscr{M}}\|^2$$

5 Refine **C** further and compute point-to-point map by using ICP like algorithm

Note that $\mathscr{S}_F^{\mathscr{M}}$ and $\mathscr{S}_F^{\mathscr{N}}$ are operators mentioned in Step 3. $\mathscr{T}_F$ acts linearly between function spaces and is sufficient to compute $\mathscr{T}$. Idea is to add a structure on the shape and work on that rather than directly on the shapes. Functional maps, due to its efficiency in dealing with high-resolution shapes by reducing the dimension where shape analysis is done, works well particularly for shape matching.

3.2 Laplace-Beltrami Operator (LBO)

The self-adjoint Laplace-Beltrami Operator on manifold $\mathscr{M}$ is specified as $\Delta_{\mathscr{M}}\colon \mathscr{F}(\mathscr{M},\mathbb{R}) \to \mathscr{F}(\mathscr{M},\mathbb{R})$ which via spectral theorem, admits an eigendecomposition with non-negative eigenvalues λ and orthonormal eigenfunctions ϕ popularly known as *manifold harmonics* $\Delta_{\mathscr{M}}\phi = \lambda\phi$ [VL08]. For mesh $\mathscr{M}$ with n vertices, a popular discrete [MDSB03] cotangent Laplace-Beltrami Operator matrix $\mathbf{L}_{\mathscr{M}} \in \mathbb{R}^{n\times n}$ is defined in terms of a sparse matrix $\mathbf{W}_{\mathscr{M}} \in \mathbb{R}^{n\times n}$ containing cotangent weights and a lumped mass matrix $\mathbf{A}_{\mathscr{M}} \in \mathbb{R}^{n\times n}$ containing vertex areas as $\mathbf{L}_{\mathscr{M}} = \mathbf{A}_{\mathscr{M}}^{-1}\mathbf{W}_{\mathscr{M}}$. The eigendecomposition (refer Fig. 2) of such a Laplace-Beltrami operator can be posed as an optimization problem:

$$\min_{\Phi} \; tr(\Phi^T \mathbf{W}_{\mathscr{M}}\Phi) \qquad s.t. \;\; \Phi^T \mathbf{A}_{\mathscr{M}}\Phi = I \qquad (1)$$

where $\Phi \in \mathbb{R}^{n\times n}$ is the eigenvector matrix as $\Phi = (\phi_1, \phi_2, ..., \phi_n)$ with eigenfunction ϕ_i arranged as columns according to increasing eigenvalues. Equation (1) is also equivalent to the generalized eigenvalue problem

$$\mathbf{W}_{\mathscr{M}}\Phi = \mathbf{A}_{\mathscr{M}}\Phi\Lambda$$

where $\Lambda = (\lambda_1, \lambda_2, ..., \lambda_n)$ is the respective diagonal eigenvalue matrix. Refer [LZ10] for extensive details on spectral analysis via Laplace-Beltrami eigenfunctions.

3.3 Hamiltonian Operator

Hamiltonian operator H is a classical operator from quantum mechanics, that appears in Schrodinger's equation describing the wave motion of a particle. On a manifold mesh $\mathscr{M}$, Hamiltonian operator is described as the extension of Laplace-Beltrami operator $\mathbf{L}_{\mathscr{M}}$:

$$\mathscr{H}_{\mathscr{M}}(f) = \mathbf{L}_{\mathscr{M}}(f) + \mu\ \mathbf{V}_{\mathscr{S}}(f)$$

with parameter $\mu \in \mathbb{R}$ and $\mathbf{V}_{\mathscr{M}} : \mathscr{M} \to \mathbb{R}_+$ a potential function on $\mathscr{M}$ [CSBK18]. Since Hamiltonian operator is the sum of two self-adjoint operators, it is also self-adjoint and hence, admits an eigendecomposition with real eigenvalues ζ and orthonormal eigenfunction ψ as $H\psi = \zeta\psi$ where ζ denotes particle energy at stationary eigenstate ψ. Refer Fig. 2 for first few Hamiltonian Eigenfunctions on the shape sorted according to increasing eigenvalues. Note that $\psi(x)$ i.e. eigenstate at point x on manifold represents the wave function of a particle where $|\psi(x)|^2$ specifies the probability of finding the particle at x. The generalized eigenvalue problem for Hamiltonian operator is specified as

$$(\mathbf{W}_{\mathscr{M}} + \mu\ \mathbf{A}_{\mathscr{M}} diag(v))\ \Psi = \mathbf{A}_{\mathscr{M}}\Psi\Theta$$

, where $\Psi \in \mathbb{R}^{n\times n}$ is the orthonormal eigenvector matrix as $\Psi = (\psi_1, \psi_2, ..., \psi_n)$ with eigenfunctions ψ_i arranged as columns, v is an n-dimensional potential vector and $\Theta = (\zeta_1, \zeta_2, ..., \zeta_n)$ is the respective diagonal eigenvalue matrix. Parameter μ controls the trade-off between global and local support of eigenbasis. [CSBK18] introduced Hamiltonian operator to shape analysis domain.

Figure 3: Gaussian Curvature visualized on wolf shape

4 MOTIVATION

Since functional map is a flexible framework, there is an opportunity to make improvements at any step of the discussed pipeline (ref. sec.3.1). Though various attempts have been made to update functional maps by modifying it at various stages, basis selection remains a crucial step in the framework since it characterizes the domain in which the analysis is going to take place. Moreover, the basis should reduce representation complexity, be stable and compact. Laplace-Beltrami Eigenfunctions [Lev06] have mainly been utilized as basis to compute the desired mapping due to multi-scale property and invariance to isometric deformations of shapes. Though Laplacian eigenfunctions are compact and stable, it tends to smoothen out sharp features on the shape, which hampers the analysis. Also global nature of these eigenfunctions make it sensitive to topological changes. However, when dealing with challenging datasets, information regarding localized and detailed features of shape become significant, which Laplace-Beltrami eigenfunctions fail to capture. This work is motivated by proposing a basis that better captures the shape geometry by picking up sharp features of the shape and be computationally viable.

5 CONTRIBUTION

For Hamiltonian operator defined on shapes, the potential function is responsible for localizing the region so that Hamiltonian eigenfunctions capture high frequency information along with preserving the properties captured by Laplace-Beltrami eigenfunctions. From computing perspective, since discrete potential function is described as a diagonal matrix it amounts to no additional computation for basis over Laplace-Beltrami basis (ref sec. 3.3).

Since Gaussian curvature (K) fully characterizes the geometry of the shape [Ale02] and is given as the product of the principal curvatures $K = k_1 * k_2$, it picks up the regions with negative and positive curvatures i.e. regions where sharp features of the shape appear. Refer Fig.3 for visualizing gaussian curvature on the shape where positive curvature regions are marked with yellow such as paws, while negative curvature regions are marked with red such as inside of the ear. Hence justifies the selection as a potential function.

In this work, the contribution is to suggest to use gaussian curvature as potential function to determine Hamiltonian basis, as it better captures the shape geometry which leads to better point-to-point correspondences, without any additional cost. Refer fig. 1 where shape signal was first projected onto each Laplace-Beltrami and Hamiltonian basis respectively and subsequently reconstructed via 150 of each set of bases. Note that Hamiltonian basis capture sharp features where fingers are also identified, while Laplace-Beltrami basis has smoothen out these details via utilizing same number of respective basis.

6 IMPLEMENTATION

6.1 Dataset

TOSCA dataset [BBK08] consisting of hi-resolution non-rigid shapes in a variety of poses have been utilized. The database contains a total of 80 objects, including 11 cats, 9 dogs, 3 wolves, 8 horses, 6 centaurs, 4 gorillas, 12 female figures, and two different male figures, containing 7 and 20 poses. Each object is a triangulated mesh with vertices, edges and triangular faces. Ground truth vertex-to-vertex correspondences are also provided, which are utilized to evaluate the performance of proposed technique.

6.2 Methodology

In this work, wolf meshes are considered to illustrate implementation details - wolf0 and wolf1. Functional framework to compute vertex-to-vertex correspondences between two triangular meshes is utilized, each mesh depicting a different pose of wolf shape from TOSCA dataset. Refer section 3.1 for the steps involved to determine correspondences via functional map.

First step is to compute feature descriptors for both the meshes, hence heat kernel signatures (HKS) [SOG09] and wave kernel signatures (WKS) [ASC11] are utilized which are defined via Laplace-Beltrami eigenvalues. Next step is to identify the bases for the shapes. For this Hamiltonian operator is selected as via potential function, bases can be enhanced to incorporate better shape geometry. Various potential functions with intrinsic features like gaussian curvature, gaussian curvature with absolute values and with extrinsic features like mean curvature and others were tried out with varied values of parameter μ (ref. Sect.3.3). Finally μ is

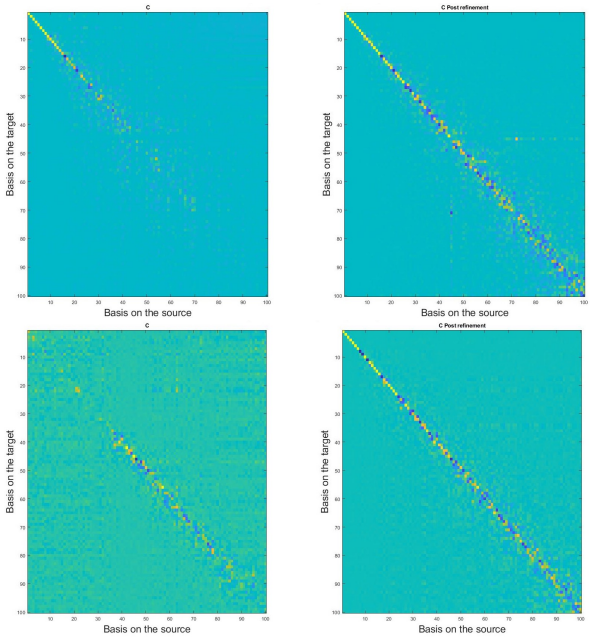

Figure 4: Functional map (pre and post refinement) via Laplace-Beltrami basis (top row); via Hamiltonian basis (bottom row) for wolf Mesh

takes as 5000 empirically and gaussian curvature was selected as the potential function to determine Hamiltonian operator. First 100 Hamiltonian eigenfunctions were considered as basis for the function spaces of shapes, arranged according to increasing eigenvalues for the next step.

Then function descriptors computed earlier were projected onto these chosen basis, so as to reduce the complexity of further shape processing. Along with Laplacian commutativity constraint the energy functional as discussed in section 3.1 is minimised via gradient-descent approach to get an optimal functional map. Procedure depicted in Section 6.2 of [OBCS$^+$12] is utilized to refine the obtained functional map and also compute point-to-point correspondences simultaneously.

Refer Fig. 4 top row, that specifies functional map with 100 basis functions; previous to and post the refinement step via Laplace-Beltrami basis and Fig 4 bottom row for Hamiltonian basis. Note that Laplace-Beltrami basis picks up the low frequency features while Hamiltonian picks up high frequency features in the initial optimal map before refining the map.

To verify if the utilized approach performs better, accuracy of obtained point-to-point correspondence map needs to be established. For that, geodesic error is computed by summing up all geodesic distances from computed mapping of points to ground-truth mapping. For a vertex p in source mesh, let the obtained corresponding vertex in target mesh is q and the ground-truth establishes vertex r in target mesh as the corresponding vertex for vertex p from source mesh, then the geodesic error at vertex p is the geodesic distance between vertices q and r on target mesh. Summing up geodesic error for each vertex on source mesh is represented as geodesic error for the obtained mapping.

Obtained results are compared with the existing approaches of Laplace-Beltrami basis [OBCS$^+$12] and compressed manifold modes [NVT$^+$14] as it claims to pick sharp details from the shapes.

With allowed normalized geodesic error threshold of 0.1, results are provided in Table 1 to compare results for wolf and Human meshes in different poses from TOSCA dataset. It shows that proposed basis performs better empirically over Laplace-Beltrami basis and wins over from compressed manifold modes by a very good margin. Refer Fig. 5 for obtained accuracy of point-to-point correspondences in terms of percentage of correct correspondences computed with respect to total number of vertices (or points) on the shapes, via Laplace-Beltarmi basis, proposed Hamiltonian basis and via compressed manifold modes. Note that proposed basis give similar *exact error* i.e. percentage of point-to-point correspondence map with zero error to that of Laplace-Beltrami basis. However, with minimal error allowed proposed basis perform much better than other two. To enhance the empirical validity of our proposed work, accuracy of point-to-point correspondences for Human meshes from the same dataset are presented in the table, which justifies the use of proposed basis over existing ones.

Basis Type	**Accuracy with less than 0.1 geodesic error**	
	Wolf	Human
LBO Eigenfunctions	84.7%	64.09%
CMM	14.73%	12.6%
Proposed basis	98.6%	72.23%

Table 1: Different basis were utilized to determine functional space to compute point-to-point correspondences

For visualization purpose refer Fig.7 for wolf meshes, to see geodesic error plot as heat map on the shape itself in case of proposed and Laplace-Beltrami basis respectively. In case of Hamiltonian basis, at the very end of tail the error is high and rest of the shape has minimal error. However, for Laplace-Beltrami basis the error is scattered over the shape and particularly present at all regions with sharp features like paws, ears, etc. Compressed manifold modes performed quite poorly, hence excluded from this visualization. Similar to the wolf meshes, human meshes (refer Fig 8) also show similar result, wherein the geodesic error induced via Laplace-Beltrami bases is scattered while for Hamiltonian bases similarly show the localized error as depicted in wolf case.

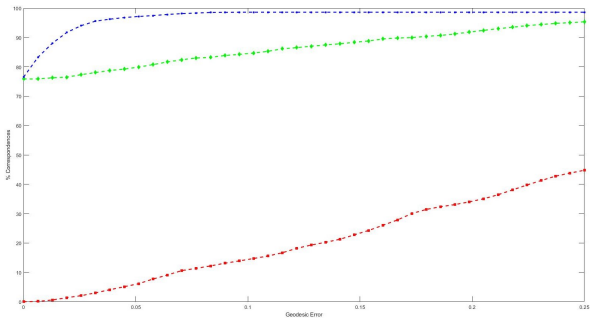

Figure 5: Accuracy of point-to-point correspondences; via Laplace-Beltrami basis(green line-diamond markers); via Hamiltonian basis (blue line-circular markers); via compressed manifold modes (red line-square markers) in the functional map framework for wolf mesh

Application

Texture mapping is considered as application for the computed correspondences via proposed basis. Refer Fig.6 for texture is transferred from wolf0 shape to wolf1, where particularly lower body of wolf is shown to highlight problematic areas.

7 DISCUSSION

In this work, the effect of selecting an intrinsic potential function, particularly Gaussian curvature, to describe Hamiltonian operator with respect to functional map framework has been studied. Due to modulations in manifold harmonics via an intrinsic potential function, Hamiltonian basis picked sharper features as compared to Laplace-Beltrami basis. These sharp features helped improve overall accuracy of point-to-point correspondences computed via functional framework considering Hamiltonian eigenfunctions as basis. Compressed manifold modes are localized basis which also aims at picking up the high frequency details, but fails to work for non-rigid shape matching because of the induced sparseness along with absence of global information determining the shape geometry. Proposed basis overcomes both these issues by being able to capture global information along with sharper features of the shape, hence fared well.

Based on empirical evidence gathered from testing multiple shape pairs in our study, it has been observed that introducing sharp feature information into the basis leads to an improved representation of shape geometry. This, in turn, enhances the ability of the basis to capture the geometric features of the shape. However, error localization is a phenomenon that is seen across multiple meshes, which might be due to Gaussian curvature picking up high curvature regions while not picking up no curvature zones or flat regions of the shape.

8 FUTURE SCOPE

The potential function proposed to compute Hamiltonian eigenfunctions is specified as an intrinsic feature of the shape which is fixed for each shape before the analysis initiates, be it some intrinsic invariant or any step function specifying a particular region on the shape. Along with its fixed nature, it shows error localization possibly due to Gaussian curvature not picking up no curvature zones. An interesting future prospect would be to compute a potential function that adapts itself such that it automatically picks those regions that are critical to both shapes considered together and in-turn enhances point-to-point correspondences too.

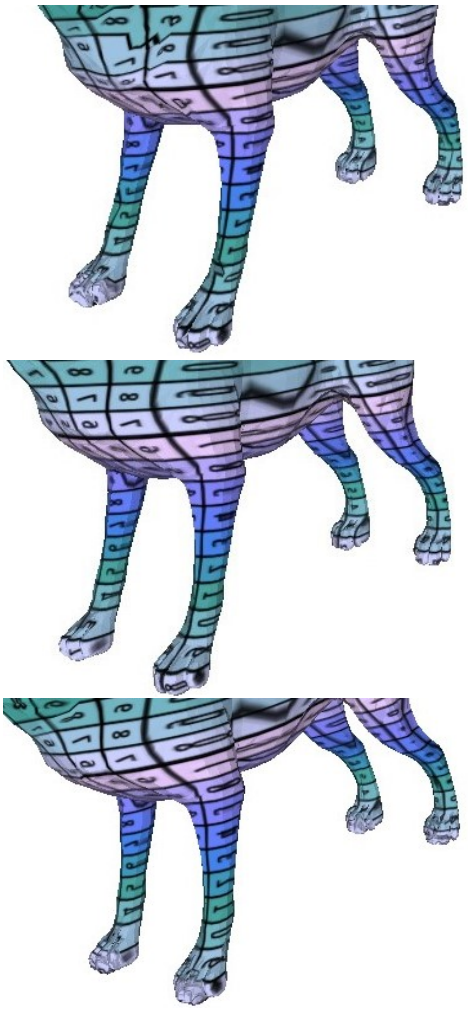

Figure 6: Texture mapping specified from wolf0 to wolf1 shape via computed correspondences; via Hamiltonian basis (top); via ground-truth (middle); via Laplace-Beltrami basis (bottom)

A potential avenue for future research is to investigate shape correspondences in non-rigid shapes that permit stretching and other deformations beyond isometric ones. This would present a greater challenge in the analysis, thereby providing an opportunity to explore new methods and approaches in this area.

Figure 7: Geodesic Error Plot visualized as heat map on the shape itself; via Hamiltonian basis (top); via Laplace-Beltrami basis (bottom) for wolf mesh

9 ACKNOWLEDGEMENTS

This work was financially supported by Visvesvaraya PhD Scheme, MeitY, Government of India under Grant MEITY-PHD-1090.

10 REFERENCES

[Ale02] Marc Alexa. Recent advances in mesh morphing. In *Computer graphics forum*, volume 21, pages 173–198. Wiley Online Library, 2002.

[ASC11] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)*, pages 1626–1633. IEEE, 2011.

[BBK07] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Rock, paper, and scissors: extrinsic vs. intrinsic similarity of non-rigid shapes. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–6. IEEE, 2007.

[BBK08] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008.

[BM92] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, volume 1611, pages 586–607. International Society for Optics and Photonics, 1992.

[BSBW14] Ben Bellekens, Vincent Spruyt, Rafael Berkvens, and Maarten Weyn. A survey of rigid 3d pointcloud registration algorithms. In *AMBIENT 2014: the Fourth International Conference on Ambient Computing, Applications, Services and Technologies, August 24-28, 2014, Rome, Italy*, pages 8–13, 2014.

[CCFM08] Umberto Castellani, Marco Cristani, Simone Fantoni, and Vittorio Murino. Sparse points matching by combining 3d mesh saliency with statistical descriptors. In *Computer Graphics Forum*, volume 27, pages 643–652. Wiley Online Library, 2008.

[CSBK18] Yoni Choukroun, Alon Shtern, Alex M Bronstein, and Ron Kimmel. Hamiltonian operator for spectral shape analysis. *IEEE transactions on visualization and computer graphics*, 2018.

[EK03] Asi Elad and Ron Kimmel. On bending invariant signatures for surfaces. *IEEE Transactions on pattern analysis and machine intelligence*, 25(10):1285–1295, 2003.

[EKB+15] Davide Eynard, Artiom Kovnatsky, Michael M Bronstein, Klaus Glashoff, and Alexander M Bronstein. Multimodal manifold analysis by simultaneous diagonalization of laplacians. *IEEE transactions on pattern analysis and machine intelligence*, 37(12):2505–2517, 2015.

[GMGP05] Natasha Gelfand, Niloy J Mitra, Leonidas J Guibas, and Helmut Pottmann. Robust global registration. In *Symposium on geometry processing*, volume 2, page 5. Vienna, Austria, 2005.

[KBB+13] Artiom Kovnatsky, Michael M Bronstein, Alexander M Bronstein, Klaus Glashoff, and Ron Kimmel. Coupled quasi-harmonic bases. In *Computer Graphics Forum*, volume 32, pages 439–448. Wiley Online Library, 2013.

[Lev06] Bruno Levy. Laplace-beltrami eigenfunctions towards an algorithm that" understands" geometry. In *IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*, pages 13–13. IEEE, 2006.

[LZ10] Bruno Levy and Richard Hao Zhang. Spectral geometry processing. In *ACM SIGGRAPH Course Notes*, 2010.

[MDSB03] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H Barr. Discrete differential-geometry operators for triangulated 2-manifolds.

In *Visualization and mathematics III*, pages 35–57. Springer, 2003.

[MRCB18] Simone Melzi, Emanuele Rodolà, Umberto Castellani, and Michael M Bronstein. Localized manifold harmonics for spectral shape analysis. In *Computer Graphics Forum*, volume 37, pages 20–34. Wiley Online Library, 2018.

[NVT+14] Thomas Neumann, Kiran Varanasi, Christian Theobalt, Marcus Magnor, and Markus Wacker. Compressed manifold modes for mesh processing. In *Computer Graphics Forum*, volume 33, pages 35–44. Wiley Online Library, 2014.

[OBCS+12] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):1–11, 2012.

[RCB+17] Emanuele Rodolà, Luca Cosmo, Michael M Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. In *Computer graphics forum*, volume 36, pages 222–236. Wiley Online Library, 2017.

[RL01] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. In *3-D Digital Imaging and Modeling, 2001. Proceedings. Third International Conference on*, pages 145–152. IEEE, 2001.

[SOG09] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum*, volume 28, pages 1383–1392. Wiley Online Library, 2009.

[VKZHCO11] Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. In *Computer Graphics Forum*, volume 30, pages 1681–1707. Wiley Online Library, 2011.

[VL08] Bruno Vallet and Bruno Lévy. Spectral geometry processing with manifold harmonics. In *Computer Graphics Forum*, volume 27, pages 251–260. Wiley Online Library, 2008.

[ZVKD10] Hao Zhang, Oliver Van Kaick, and Ramsay Dyer. Spectral mesh processing. In *Computer graphics forum*, volume 29, pages 1865–1894. Wiley Online Library, 2010.

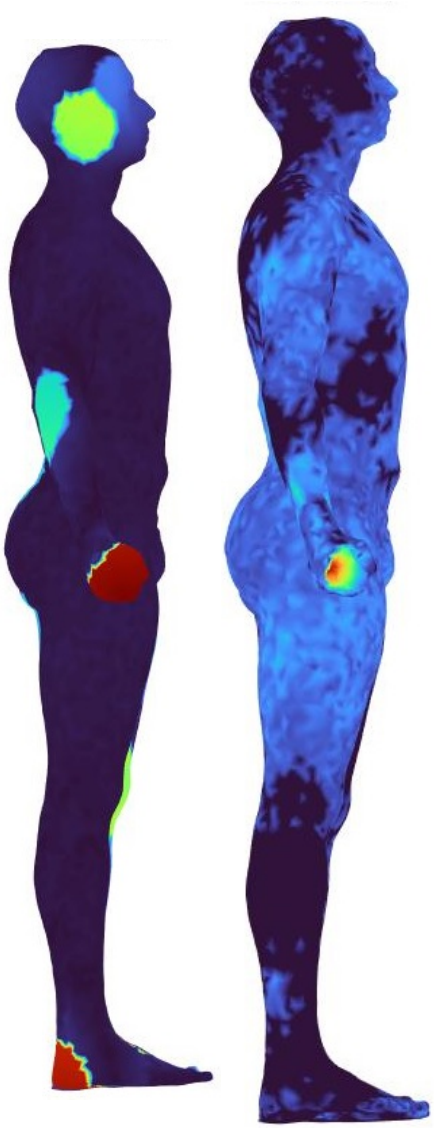

Figure 8: Geodesic Error Plot visualized as heat map on the shape itself; via Hamiltonian basis (left); via Laplace-Beltrami basis (right) for Human mesh