

Towards Implicit Correspondence in Signed Distance Field Evolution

Miroslava Slavcheva^{1,2}
mira.slavcheva@tum.de

Maximilian Baust¹
maximilian.baust@tum.de

Slobodan Ilic^{1,2}
slobodan.ilic@siemens.com

¹ Technische Universität München

² Siemens Corporate Technology

Abstract

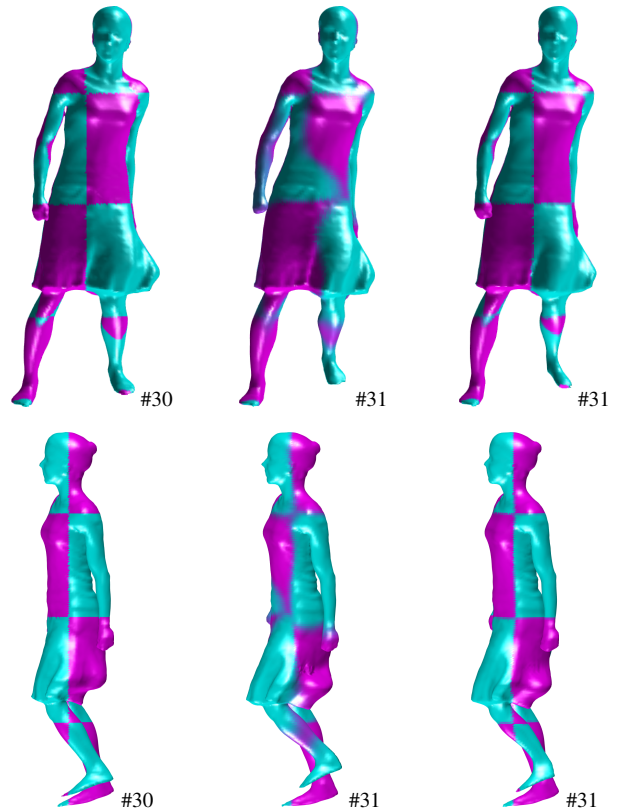
The level set framework is widely used in geometry processing due to its ability to handle topological changes and the readily accessible shape properties it provides, such as normals and curvature. However, its major drawback is the lack of correspondence preservation throughout the level set evolution. Therefore, data associated with the surface, such as colour, is lost. The objective of this paper is a variational approach for signed distance field evolution which implicitly preserves correspondences. We propose an energy functional based on a novel data term, which aligns the lowest-frequency Laplacian eigenfunction representations of the input and target shapes. As these encode information about natural deformations that the shape can undergo, our strategy manages to prevent data diffusion into the volume. We demonstrate that our system is able to preserve texture throughout articulated motion sequences, and evaluate its geometric accuracy on public data.

1. Introduction

Real-world scenes contain shapes that move and interact non-rigidly over time, *i.e.* they inhabit a 4D spatio-temporal domain. Multi-camera systems are able to recover complete, but independent 3D models of the scenes at isolated time instances [29]. However, these are not consistent over time, lacking motion information. Thus in order to enable tasks such as performance capture, primitives on a template 3D surface have to be tracked across frames of such motion sequences, following a deformation model. This challenging problem has numerous applications, among which virtual reality, 3D avatar animation and special effects.

One major difficulty is capturing non-rigid motion involving topological changes, *e.g.* when subjects interact, or when loose clothing touches or splits from other surfaces. While triangular meshes have become a common discrete surface representation for motion capture, they require tedious handling for such situations [38].

On the other hand, level set methods [20, 21] inherently



(a) Reference shape. (b) Classical level set. (c) Proposed.

Figure 1: Example of colour-preserving SDF evolution on frames 30 and 31 from the *Swing* sequence of [35]. Mainly the left leg and the body move. (a) Reference shape and texture. (b) Warped shape using standard level set evolution. Colour diffusion is clearly identifiable on the moving parts. (c) Warped shape using the proposed evolution modification with Laplacian eigenfunction term. The original texture is preserved, while the geometric error remains at the same order of magnitude as for the classical approach.

manage changing topology without need for additional processing. They are widely used in shape analysis due to the ease they provide for calculating geometric properties, such

as derivatives, normals and curvature, over a fixed Cartesian grid without parameterization. However, the underlying level set evolution involves an incremental iterative numerical scheme, in which correspondences are lost [22, 37]. This limits applications to reconstruction and modelling, but prevents tasks that require tracking data associated with the surface, such as texture mapping and identity transfer.

To remedy this discrepancy, researchers have investigated hybrid structures combining the advantages of both meshes and level sets. For instance, SpringLS [13, 14] provide interoperability between the two representations, allowing the user to interpret geometry in the form that is more beneficial at the current step of an algorithm.

Other authors adhere to the mesh representation and use spectral methods based on the Laplace-Beltrami operator to calculate volumetric descriptors, which are matched to identify corresponding interest points across shapes. These include the volumetric heat kernel signature [23] and its scale-invariant follow-up versions [12].

Other approaches favour the level set framework. The Particle Level Set [6] and the Marker Level Set [16, 17] methods apply the estimated motion not only to the volumetric grid, but also to a set of particles attached to the surface, and subsequently correct for their locations. Similarly, Pons *et al.* [22] maintain explicit backward correspondences to the reference shape and advect them using a system of coupled Eulerian partial differential equations.

While some of these techniques demonstrate successful results on synthetic examples or in scenarios where the level set equations are analytically defined, they all entail some overhead for representation conversion, descriptor matching, or additional equation handling. To the best of our knowledge, no method has managed to integrate correspondence tracking within the level set equation itself. This is largely due to the conflicting objectives of an evolving level set energy versus direct explicit correspondence matching.

The objective of our work is to develop a variational framework for 3D signed distance field (SDF) evolution, which allows to propagate volumetric correspondences. More specifically, we propose a novel energy term directly in the evolution equation. It is based on the intuition that shapes do not deform absolutely randomly, but follow some natural patterns, which are reflected through their Laplacian eigenfunctions [8, 15, 25, 28]. Our energy imposes similarity between the Laplacian eigenfunction representations of the reference and target shapes. As a result, our framework yields *implicit correspondences* between the SDF representations of the shapes. We demonstrate this through the application of colourful SDF evolution, in which we manage to preserve surface texture during the evolution process. Moreover, we quantitatively evaluate the geometric error on public 4D sequences, showing that it remains at the same order of magnitude as with the standard level set framework.

2. Related Work

Deformable models are commonly used for 3D reconstruction, registration, simulation, animation and motion tracking [14]. Meshes and level sets are the two representations that are most often employed for manipulating the 3D data at hand. Each has advantages in certain applications, *e.g.* meshes are more suitable for registration, where correspondences between vertices need to be estimated [2, 3, 33], while level sets are more often utilized in segmentation, where the boundary of a particular structure is determined by a propagating front [20]. However, for tasks which need to combine the two representations, there either has to be an explicit structure that can be cast into either a mesh or a level set [14], or the solution needs to implicitly provide the data associated with the other representation.

Spectral descriptors Spectral decomposition methods based on the Laplace-Beltrami operator on meshes have achieved remarkable results for non-rigid full and partial shape matching [4, 10, 11, 26, 32]. They model deformations as approximate isometries of the object boundary, *i.e.* its surface. Inspired by this success, researchers have looked into volume isometries, which are more natural to be preserved during motion. This brought about the volumetric heat kernel signatures [23], volumetric maximally stable extremal regions [12], and a variety of other signatures based on Laplace-Beltrami eigendecomposition [24, 27, 28]. They typically take an arbitrarily big subset of the operator eigenfunctions in order to build a descriptor. Subsequently, quantization and matching are required in order to determine corresponding regions for the applications of non-rigid shape retrieval and classification. However, examples are limited to mainly synthetic noise-free meshes.

Hybrid structures While the spectral descriptors are computed from a mesh representation of the shape, some authors avoid it due to the difficulty of discretizing equations on polygonal grids and the tedious calculation of projections onto the discretized surface for handling properties such as gradients [1]. Instead, they prefer to use the level set framework [21, 30, 34]. It, however, does not preserve correspondences and is therefore ill-suited for tasks such as surface registration and motion tracking. Pons *et al.* [22] were among the first to propose a way to maintain correspondences during the level set evolution. They use a system of coupled PDEs in order to track backward correspondences to the initial surface position. Their framework handles large deformations and topological changes, but is based on analytically defined motion equations of curvature-dependent speed.

The Particle Level Set method [6] provides a similar scheme, in which a set of particles are associated with the

initial surface. They are advected together with the level set evolution, and then processed for addition and deletion where topological changes occurred. Moreover, at each iteration, a correction step has to be done to ensure that the particles are still aligned with the surface. This is a complicated procedure, which might take hundreds of seconds. Therefore, speedups and modifications followed, such as the Marker Level Set [16, 17], which is still far from interactive frame rates.

More recently, SpringLS have been proposed to offer direct interoperability between meshes and level sets [13, 14]. They define a level set as a constellation of triangular surface elements, which are loosely connected via structures with the physical properties of springs, so that rigidity constraints can be applied. However, the processing speed still remains at the order of several minutes, while accuracy is only slightly better than that of Pons *et al.* [22].

3. Overview

Having experienced issues with selecting appropriate step sizes for particle-based level set approaches (which is an issue addressed by the authors themselves [14]), we set out to develop a framework which is entirely based on level set evolution, and maintains correspondences in an implicit way. We take inspiration in part by the Laplace-Beltrami operator, whose spectrum is an isometry invariant of the shape, independent of its spatial position or parameterization, and is even dubbed to "understand" geometry [8, 25]. In analogy to physical vibration models, it is indicative of the trajectories in which a surface is able to deform [8].

The Laplace-Beltrami is an operator associated with the surface, *i.e.* the volume boundary of an object, and therefore convenient methods for calculating it from a mesh representation exist. While it is invariant to isometric deformations [28], it is more natural for the volume to be preserved during articulated motion. However, the volumetric Laplacian shares similar invariance properties only if a very fine grid with appropriate boundary conditions is used, which might be prohibitively expensive for practical 3D scenarios [28]. Nevertheless, it has been shown that the Laplacian of a voxel representation of a shape is able to handle its articulations [15]. Thus it is sufficient for our main goal of human performance capture.

Therefore, we propose to stay within the level set framework, where objects are represented via voxel grids. We utilize the eigenfunctions of the Laplacian, which capture information about natural non-rigid motion patterns, and include them directly as an energy term in a variational framework. Thus, without explicitly tracking correspondences, we are able to implicitly infer them. This is demonstrated through texture transfer during level set evolution, as shown in Figure 1 and the results that follow.

4. Proposed Approach

In this section we describe our signed distance field evolution method. More specifically, we focus on the novel Laplacian eigenfunction energy term, which permits the propagation of implicit correspondences. In addition, we highlight the modifications to the standard level set framework, which are required to make that propagation possible.

4.1. Signed Distance Field Evolution

The original level set framework [20, 21] is based on evolution with curvature-dependent speed. This is appropriate for applications such as segmentation, in which a randomly initialized shape is evolved towards a target boundary. Since we are interested in deforming a 3D shape to a given target, we instead prefer the variational formulation more commonly used for reconstruction [9, 31]. Our choice of level set is the signed distance field, which provides a convenient way to distinguish between inside and outside of objects, and has proven extremely beneficial for shape representation and reconstruction in recent years [5, 19].

Let ϕ_{input} be the SDF which is being evolved towards the target SDF ϕ_{target} . Note that they occupy the same volume and are discretized using the same voxel size.

We aim to estimate a 3D flow field Ψ which warps the input shape towards the target, *i.e.* it seeks to align $\phi_{input}(\Psi)$ with ϕ_{target} . The deformation field has the same resolution as the SDFs and assigns a displacement vector (u, v, w) to each voxel center (x, y, z) in the world coordinate frame. Therefore, our approach resembles scene flow [7, 36], but the warp field is per voxel rather than per 3D point in a stereo or RGB-D image.

Thus our SDF evolution **data term**, similarly to any level set approach, is a sum over voxels:

$$E_{data}(\Psi) = \frac{1}{2} \sum_{x,y,z} (\phi_{input}(x+u, y+v, z+w) - \phi_{target}(x, y, z))^2. \quad (1)$$

Note that the vector (u, v, w) is different for each voxel, as the flow field Ψ consists of all of these displacements, organized in a regular grid.

The level set framework is iterative and typically requires re-initialization after every couple of iterations in order to ensure that the geometric properties of the level set are correct. This can be done by explicitly applying a PDE that enforces the gradient magnitude to equal unity everywhere in the volume [20]. Alternatively, it can be achieved directly through the variational framework, whereby a **level set regularizer term** is added to softly impose this constraint [9]. As a variational framework is our goal too, we use the term

introduced by Li *et al.* [9]:

$$E_{reg}(\Psi) = \frac{1}{2} \sum_{x,y,z} (|\nabla \phi_{input}(x+u, y+v, z+w)| - 1)^2, \quad (2)$$

where the ∇ operator takes the gradient of the SDF. Numerically, we implement it as a central difference.

4.2. Preventing Diffusion

The terms in Eq. 1 and 2 form a variational energy, which is solved iteratively via a numerical scheme, such as gradient descent, based on the respective Euler-Lagrange equations. At each iteration step of the evolution process, the warp field is updated and applied to the current shape. As this results in displacements to non-integer voxel locations, the new state of the SDF has to be obtained via tri-linear interpolation [31, 37]. This is the main reason why correspondences are not preserved in the level set framework [22], which is manifested as colour diffusion when attempting to propagate texture, as shown in Figure 1b. To reduce this diffusion, instead of incrementally calculating the deformation field, we always estimate it relative to the initial shape. We thus obtain the warped SDF via a single interpolation step from ϕ_{input} , rather than a sequence of interpolations applied at each iteration. As we will show in the experimental section, this results in a slightly inferior geometric accuracy, but is a crucial modification for preserving implicit correspondence.

4.3. Laplacian Eigencolour Term

As the eigenfunction representation results in a colouring of the voxels, which describe the natural deformation modes of the shape, we also call it *eigencolour*. To build it, we start with a binary voxel representation, which can be obtained by thresholding an SDF: signed distances with value less or equal to zero belong to the object. Let the total number of these occupied voxels be k . Then the Laplacian matrix $\mathbf{L} \in \mathbb{N}^{k \times k}$ is built as follows:

$$\mathbf{L}_{i,j} = \begin{cases} 1 & \text{if } i \neq j \text{ and } i \Leftrightarrow j, \\ -6 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Above we use the symbol \Leftrightarrow to denote that the voxels with linearized indices i and j share a common side, *i.e.* they are 6-neighbours in the 3D grid.

Next, we solve the Laplacian eigenvalue problem, also known as the Helmholtz equation [25], obtaining the eigenfunctions of \mathbf{L} . The full spectrum of the Laplacian (or rather, the Laplace-Beltrami) reflects all possible ways in which the shape can deform isometrically. However, since real-world data may contain noise, we discard high-frequency eigenfunctions. Instead, we want to capture only

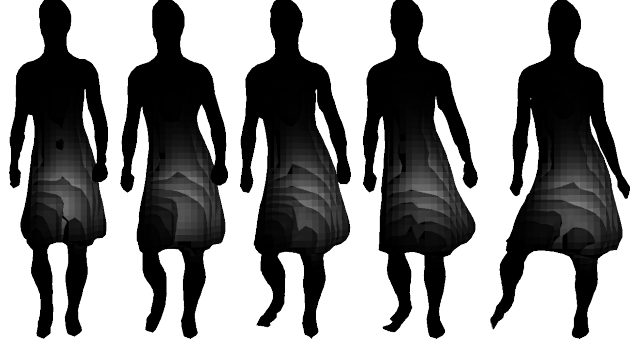


Figure 2: Visualization of the Laplacian eigencolour representation of several poses of the same subject. The patterns around skirt folds confirm that the chosen lowest-frequency eigenfunction is suitable for capturing important characteristics of the 3D shape, while discarding sporadic details.

the most significant characteristics of the shape. Therefore, we decide to retain only the eigenfunction with smallest non-zero eigenvalue. Figure 3 compares several eigenfunction colourings of the same shape using different eigenvalues. It clearly demonstrates that the contours of the higher-frequency eigenfunctions are too unstable.

Similarly, Figure 2 depicts the lowest-frequency Laplacian eigencolourings of the same shape in different poses. The skirt is the most motile part of the object and this is where the contours saturate. Moreover, they form similar patterns across the different poses, indicating that the eigenfunctions indeed encode the most natural deformations of the captured shape, regardless of its articulation.

As the eigenfunction of \mathbf{L} is a k -element vector, we pad it to the size of the volume and de-linearize its indices to obtain l , which is the *eigencolouring* of the volume. It is a scalar field of the same resolution as the original SDF grid. Padding is effectively done in the outside area of the shape, and we use the minimum value entry of the eigenfunction to pad because this does not contradict the gradient of l . Furthermore, we normalize the values to the interval $[-1, 1]$, as different shapes will have slightly different entries in the eigenfunctions.

Since ϕ_{input} and ϕ_{target} represent articulations of the same shape, we expect their lowest-frequency Laplacian eigenfunction representations l_{input} and l_{target} to be similar. We integrate this intuition into our variational framework through the following **Laplacian eigencolour term**:

$$E_{eig}(\Psi) = \frac{1}{2} \sum_{x,y,z} (l_{input}(x+u, y+v, z+w) - l_{target}(x, y, z))^2. \quad (4)$$

Note that E_{eig} is very similar to the data term E_{data} . In fact, it is possible to use only E_{eig} as a data term, but since

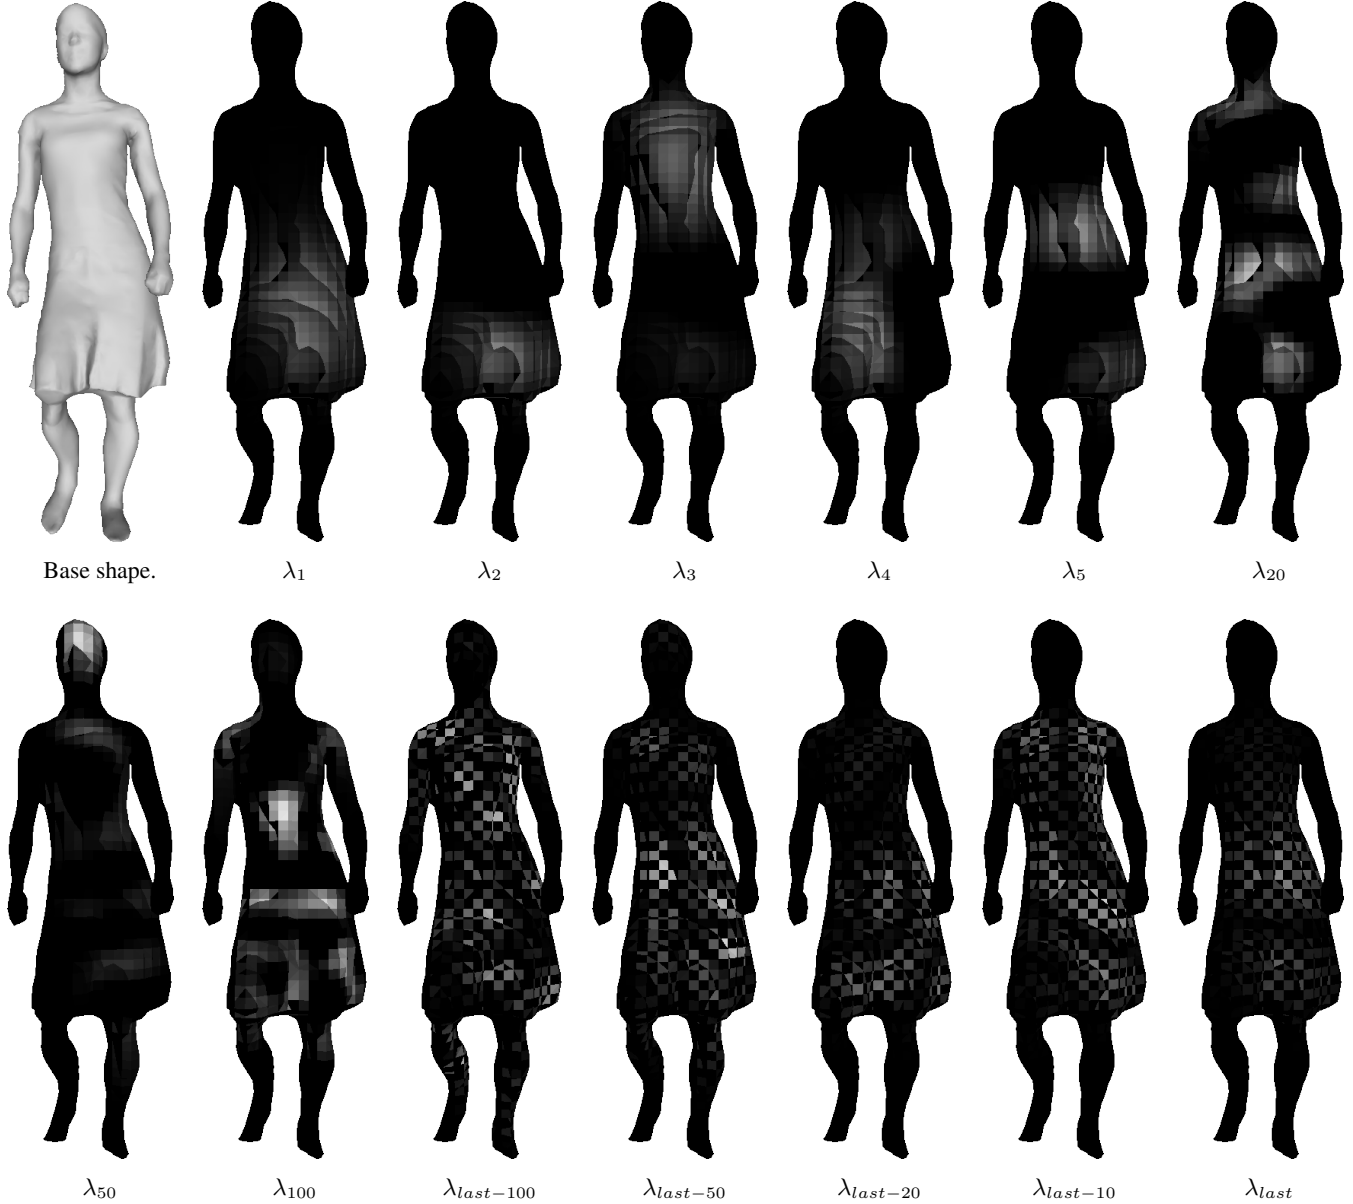


Figure 3: Laplacian eigenfunction visualization. λ_1 corresponds to the smallest eigenvalue, λ_2 to the second-smallest, *etc.* λ_{last} denotes the largest eigenvalue, out of a total of 3151 in this case, while $\lambda_{last-10}$ is the 10th largest and so on. The contours indicate that the smaller eigenvalues, corresponding to the lower-frequency eigenfunctions, capture more general and significant characteristics of the shape. On the other hand, the larger eigenvalues are associated with eigenfunctions containing a lot of high-frequency noise. Therefore, we choose the smallest eigenvalue for our framework.

it stems from a binary voxel representation rather than an SDF with smooth transitions between values, it would result in a blocky output. Exploring this possibility is, however, an interesting direction for future work.

Finally, the **complete energy** then becomes:

$$E(\Psi) = E_{data}(\Psi) + \omega_{reg}E_{reg}(\Psi) + \omega_{eig}E_{eig}(\Psi). \quad (5)$$

We solve it via a gradient descent scheme, which terminates

when the sum of squared distances update per voxel becomes smaller than 10^{-7} . The weights of the other terms in the presented examples are $\omega_{reg} = 0.05$ and $\omega_{eig} = 1$.

4.4. Implementation Details

We implemented the presented framework in MATLAB. Processing a volume of 128^3 voxels requires 0.25 seconds per iteration on a laptop equipped with an Intel i7-4900MQ

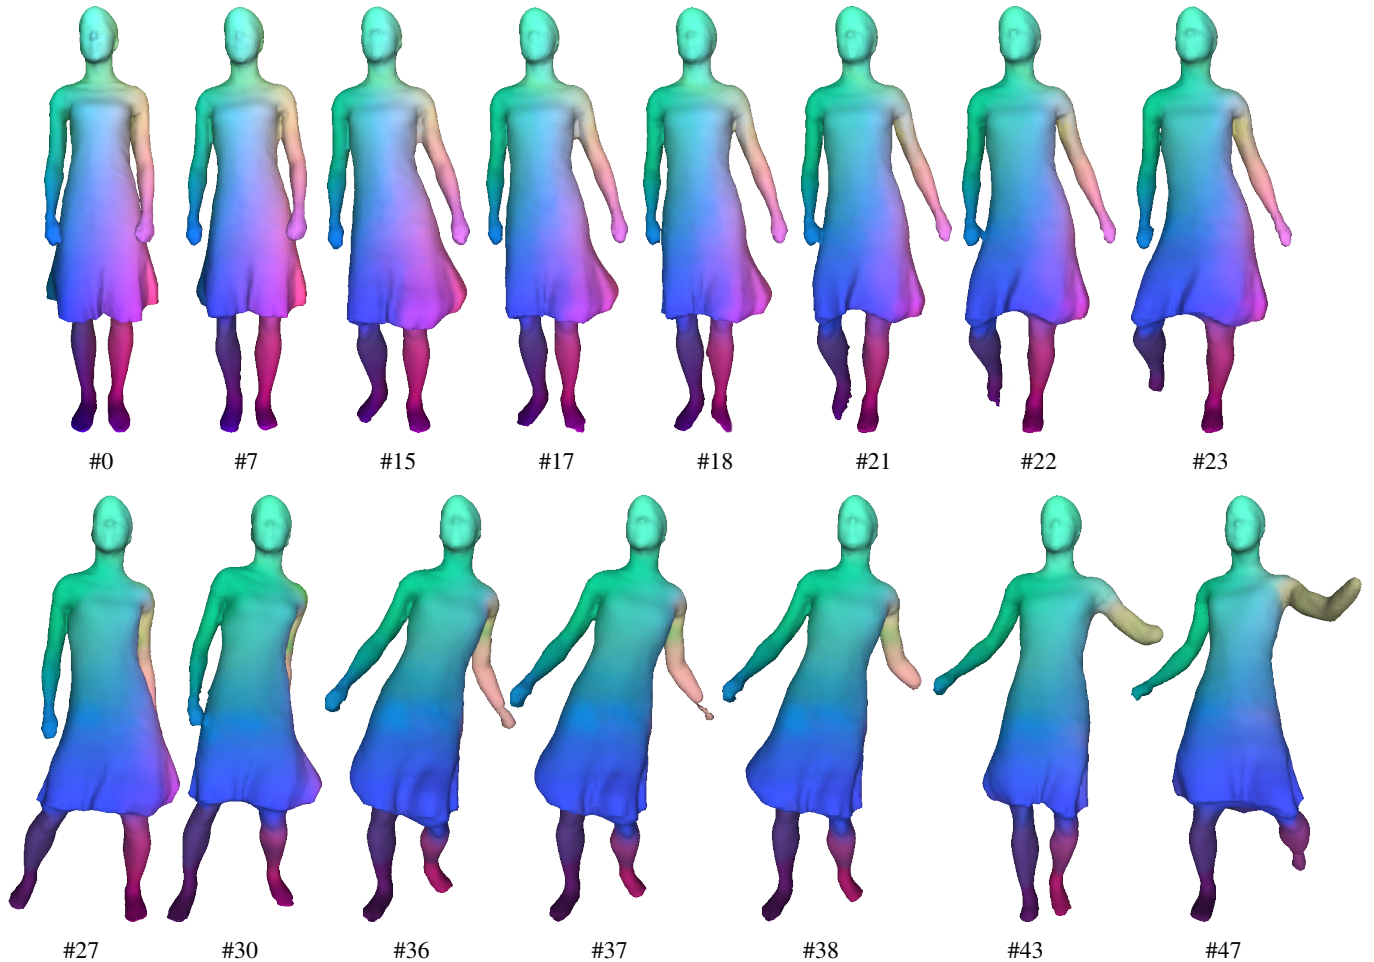


Figure 4: Results of texture transfer on the MIT dataset [35] T_{swing} sequence. The frame index is shown below each image. The top row demonstrates that our strategy is able to implicitly preserve correspondences, and consequently keep a stable texture, when the inter-frame motion is smooth. The bottom row shows the limitations of the technique under abrupt motion, namely texture diffusion (blue replacing purple on the skirt) and inability of the non-incremental scheme to recover the arm.

CPU at 2.80 GHz. This compares favourably with the times reported by SpringLS [13] and the Particle Level Set [6], both of which require several minutes for their respective applications. As most operations run independently over voxel grids, they can be parallelized. We thus expect an optimized C++/ CUDA implementation to bring further significant performance gains.

Our framework always converges in less than 120 iterations, even when non-consecutive frames are taken. This is, however, slower than the classical incremental variational level set version, which typically requires less than 80 steps. Increasing the weight of the eigencolour term can drastically reduce the number of iterations by more than 2-3 times, but this might have a slightly adverse effect on the geometric error. Therefore, the relative weight selection in the energy functional brings a trade-off between accuracy and speed.

5. Results

Now we demonstrate the capabilities of our proposed framework. The application we tackle is texture transfer in 4D motion capture, where consecutive frames exhibit smooth motions [35]. Note that since we rely on the volume Laplacian, we do not expect to be able to handle larger pose differences as used in non-rigid shape matching, *e.g.* as in the popular FAUST dataset [2].

For comparison, we use a baseline which corresponds to a standard variational signed distance field evolution, *i.e.* based on the energy functional

$$E_{baseline} = E_{data} + \omega_{reg} E_{reg}. \quad (6)$$

As explained in the method section, it can be applied in an incremental way, relative to the state in the previous iteration, or in a non-incremental way, always relative to ϕ_{input} .

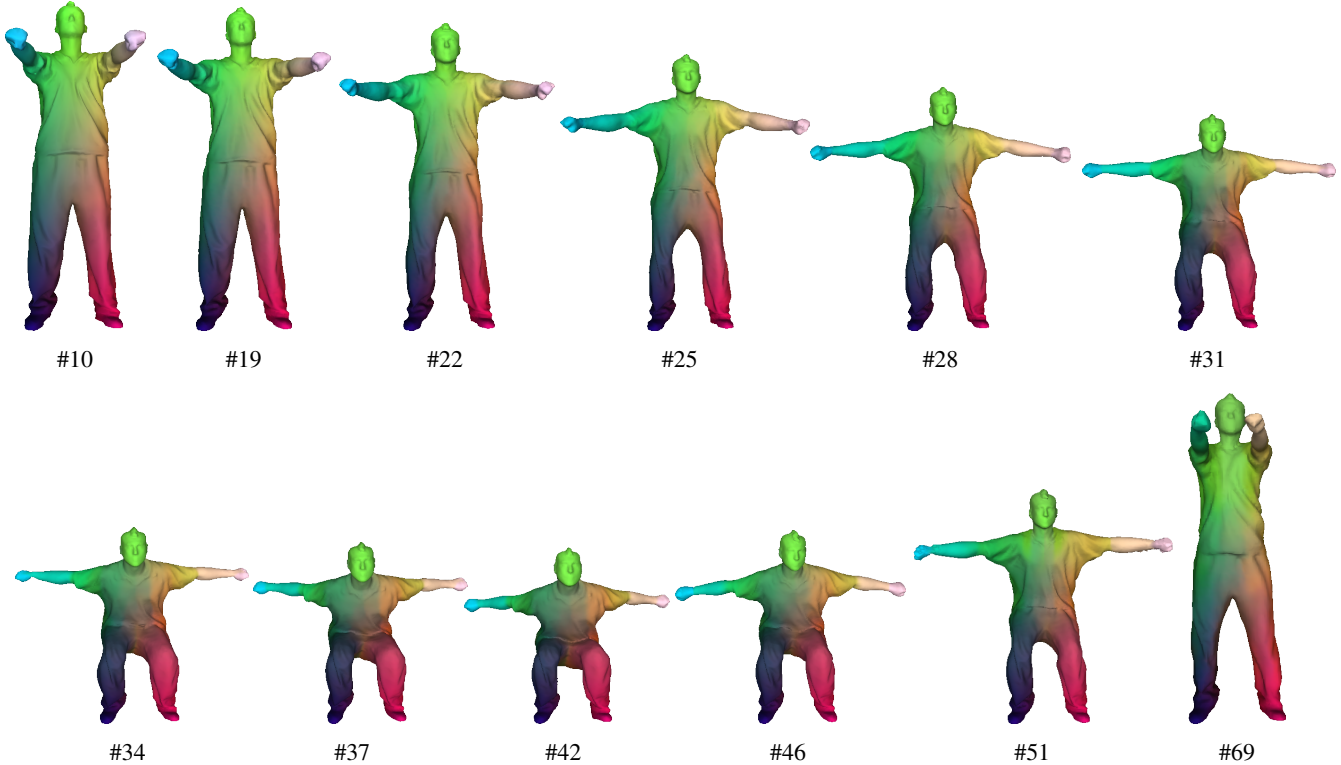


Figure 5: Results of texture transfer on the MIT dataset [35] D_squat sequence. The frame index is shown below each image. As there is no extremely abrupt motion, our framework manages to preserve a stable texture.

A visual comparison with the incremental baseline was already shown in Figure 1, where it is clearly visible that the classical variational level set framework leads to colour diffusion inside the volume.

Next, we use data from the MIT dataset [35], which contains multi-view mesh reconstructions of dynamic single-subject scenes. However, we do not make use of the meshes, but convert them to their signed distance representation and then feed them into our framework.

In Figures 4 and 5 we colour the voxels of an initial shape according to their 3D locations and aim to propagate this texture over time. Note that we do not deform a template separately to each frame, but start with one coloured grid and continuously update it to match every incoming frame. This strategy is more prone to drift, since once the colour diffuses, it cannot be remedied, but we view it as more appropriate for evaluation of 4D capture.

The resulting shapes are geometrically plausible. In fact, the average geometric error is 1.2 cm, which is well below the 1.6 cm voxel size used for discretization, and not much higher than the 0.7 cm error of the incremental baseline.

The texture is preserved well throughout the first 25 frames of the T_swing sequence, *i.e.* when the motion is not too abrupt, as shown on the top row of Figure 4. Larger motions, however, present a challenge, *e.g.* when the girl makes

a big step and her skirt executes a larger motion. Then the blue colour on the skirt diffuses from the right part towards the left one. Afterwards, once the motion is smaller again, texture stays stable in the next slower motion part of the sequence. However, the left arm motion is so large in frames 36-45 that the non-incremental strategy cannot recover the geometry. Consequently, the colour information there is terminally lost. Even though the movement becomes smaller afterwards and the wrist can be reconstructed again, it can only extrapolate colour from the parts that were left. As shown on the same sequence in KillingFusion [31], the classical incremental strategy does not face such problems in recovering geometry - but would fail in transferring texture, as indicated in Figure 1.

Figure 5 shows the first squat of the D_squat sequence. As the motion is smoother, the colours are well preserved. Therefore, we conclude that our framework is able to keep implicit correspondences for smooth motions - improving over the classical level set approach, but still has issues under arbitrarily large movements.

Finally, to demonstrate the contribution of the proposed *Laplacian eigencolour term* over the baseline in its non-incremental version, we apply a simple texture to the torso and legs of the man in a frame of D_squat , and compare its preservation over the several following frames in Figure 6.

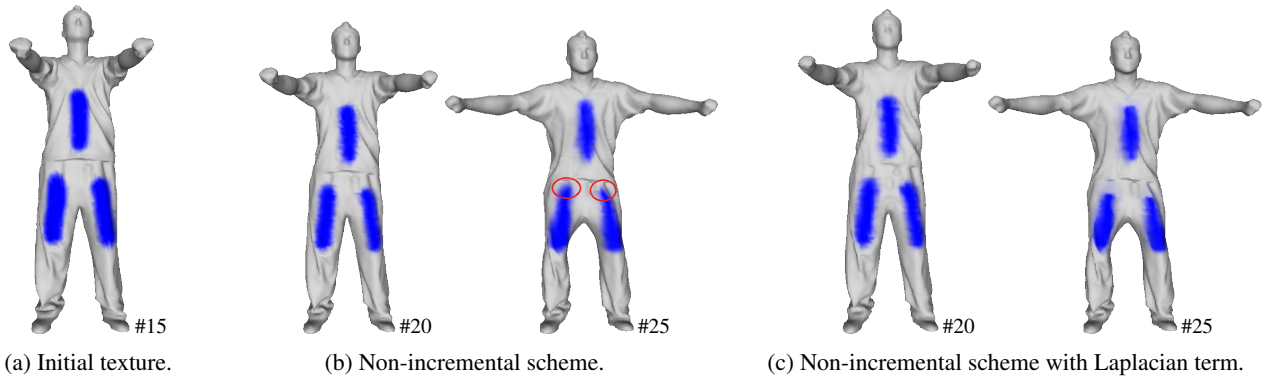


Figure 6: Contribution of the Laplacian eigencolour term. Blue regions remain stable while only the arms move. While there is certain colour diffusion with both schemes when the torso and legs also move, the Laplacian term restricts it better.

While the motion is only in the arms, both schemes keep the blue colour fixed. However, once the legs start bending, the version without Laplacian eigencolour (Fig. 6b) visibly diffuses more than the proposed one (Fig. 6c). Therefore, the lowest-frequency Laplacian eigenfunction aids in implicitly matching corresponding regions of the shapes.

6. Future Work

The presented implicit correspondence energy term is based on the Laplacian of a voxelized 3D shape, as it is suitable for direct inclusion in a variational level set framework. While our system can preserve surface colour during moderate movements, it does not succeed under very large motion, as seen in Figure 4. This is due to the fact that the Laplacian operator is not entirely invariant to isometric deformations [28], which might be a desirable property for applications such as non-rigid shape retrieval, but hinders handling big displacements. On the other hand, the Laplace-Beltrami operator is invariant [28], but we have avoided using it, since it is most straightforward to compute it through the mesh representation. In future work, we will explore possibilities to calculate it as a projection of a surface represented in implicit form [1], which would allow us to handle even larger deformations directly through the evolution equation. Further, Figure 3 suggests that a combination of several low-frequency eigenfunctions might be better suited to describe the shape than a single one, but previous works usually use a seemingly arbitrary number to build signatures for descriptor matching. Thus, we plan to study into what subset of the spectrum is required to robustly represent a shape while still discarding its high-frequency details, and use a sum over the corresponding eigenfunctions. Moreover, this work has focused on implicitly tracking complete 3D models, while in future we would like to tackle dense reconstruction from a single depth stream [18], which will entail further challenges in handling partial overlap.

7. Conclusion

We have presented a variational framework for signed distance field evolution, which seeks to preserve correspondences implicitly across volumetric representations of subjects undergoing moderate articulated motions. It is based on their lowest-frequency Laplacian eigenfunctions, as they encode information about the natural deformation patterns, and consequently the non-rigid isometries, of the underlying shape. We have demonstrated the ability of this energy formulation to reduce colour diffusion and preserve texture during level set evolution, while keeping geometric accuracy at the same order of magnitude as the classical level set formulation.

References

- [1] M. Bertalmio, L.-T. Cheng, S. Osher, and G. Sapiro. Variational Problems and Partial Differential Equations on Implicit Surfaces. *Journal of Computational Physics*, 174(2):759–780, 2001. 2, 8
- [2] F. Bogo, J. Romero, M. Loper, and M. J. Black. FAUST: Dataset and Evaluation for 3D Mesh Registration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2, 6
- [3] F. Bogo, J. Romero, G. Pons-Moll, and M. J. Black. Dynamic FAUST: Registering Human Bodies in Motion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [4] A. M. Bronstein. Spectral Descriptors for Deformable Shapes. Technical report, School of Electrical Engineering, Tel Aviv University, 2011. 2
- [5] B. Curless and M. Levoy. A Volumetric Method for Building Complex Models from Range Images. In *23rd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '96*, pages 303–312, 1996. 3
- [6] D. Enright, S. Marschner, and R. Fedkiw. Animation and Rendering of Complex Water Surfaces. *ACM Transactions on Graphics (TOG)*, 21(3):736–744, 2002. 2, 6

- [7] F. Huguet and F. Devernay. A Variational Method for Scene Flow Estimation from Stereo Sequences. In *IEEE International Conference on Computer Vision (ICCV)*, 2007. 3
- [8] B. Levy. Laplace-Beltrami Eigenfunctions Towards an Algorithm That "Understands" Geometry. In *IEEE International Conference on Shape Modeling and Applications*, 2006. 2, 3
- [9] C. Li, C. Xu, C. Gui, and M. D. Fox. Level Set Evolution Without Re-initialization: A New Variational Formulation. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 3, 4
- [10] O. Litany, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Fully Spectral Partial Shape Matching. *Computer Graphics Forum (CGF)*, 36(2):247–258, 2017. 2
- [11] O. Litany, E. Rodolà, A. M. Bronstein, M. M. Bronstein, and D. Cremers. Non-Rigid Puzzles. *Computer Graphics Forum (CGF)*, 35(5):135–143, 2016. 2
- [12] R. Litman, A. M. Bronstein, and M. M. Bronstein. Stable Volumetric Features in Deformable Shapes. *Computers & Graphics*, 36(5):569–576, 2012. 2
- [13] B. C. Lucas, M. Kazhdan, and R. H. Taylor. SpringLS: A Deformable Model Representation to Provide Interoperability between Meshes and Level Sets. In *14th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2011. 2, 3, 6
- [14] B. C. Lucas, M. Kazhdan, and R. H. Taylor. Spring Level Sets: A Deformable Model Representation to Provide Interoperability between Meshes and Level Sets. *IEEE Transactions on Visualization and Computer Graphics (VCG)*, 19(5):852–865, 2013. 2, 3
- [15] D. Mateus, R. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated Shape Matching using Laplacian Eigenfunctions and Unsupervised Point Registration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008. 2, 3
- [16] V. Mihalef. *The Marker Level Set Method: Applications to Simulation of Liquids*. PhD thesis, 2007. 2, 3
- [17] V. Mihalef, D. Metaxas, and M. Sussman. Textured Liquids based on the Marker Level Set. *Computer Graphics Forum (CGF)*, 26(3):457–466, 2007. 2, 3
- [18] R. A. Newcombe, D. Fox, and S. M. Seitz. DynamicFusion: Reconstruction and Tracking of Non-rigid Scenes in Real-Time. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 8
- [19] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon. KinectFusion: Real-Time Dense Surface Mapping and Tracking. In *10th International Symposium on Mixed and Augmented Reality (ISMAR)*, 2011. 3
- [20] S. Osher and R. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*, volume 153 of *Applied Mathematical Science*. Springer, 2003. 1, 2, 3
- [21] S. Osher and J. Sethian. Fronts Propagating with Curvature-dependent speed: Algorithms based on Hamilton-Jacobi Formulations. *Journal of Computational Physics*, 79(1):12–49, 1988. 1, 2, 3
- [22] J.-P. Pons, G. Hermosillo, R. Keriven, and O. Faugeras. How to Deal with Point Correspondences and Tangential Velocities in the Level Set Framework. In *9th IEEE International Conference on Computer Vision (ICCV)*, 2003. 2, 3, 4
- [23] D. Raviv, M. M. Bronstein, A. M. Bronstein, and R. Kimmel. Volumetric Heat Kernel Signatures. In *ACM Workshop on 3D Object Retrieval*, 2010. 2
- [24] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-spectra as Fingerprints for Shape Matching. In *ACM Symposium on Solid and Physical Modeling*, 2005. 2
- [25] 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. 2, 3, 4
- [26] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial Functional Correspondence. *Computer Graphics Forum (CGF)*, 36(1):222–236, 2017. 2
- [27] R. M. Rustamov. Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation. In *Eurographics Symposium on Geometry Processing (SGP)*, 2007. 2
- [28] R. M. Rustamov. Interpolated Eigenfunctions for Volumetric Shape Processing. *The Visual Computer*, 27(11), 2011. 2, 3, 8
- [29] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006. 1
- [30] J. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge University Press, 1999. 2
- [31] M. Slavcheva, M. Baust, D. Cremers, and S. Ilic. Killing-Fusion: Non-rigid 3D Reconstruction without Correspondences. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3, 4, 7
- [32] J. Sun, M. Ovsjanikov, and L. Guibas. A Concise and Provably Informative Multi-scale Signature Based on Heat Diffusion. In *Proceedings of the Symposium on Geometry Processing (SGP)*, 2009. 2
- [33] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer. Elastically Deformable Models. *Computer Graphics*, 21(4):205–214, 1987. 2
- [34] G. Turk and J. O'Brien. Shape Transformation Using Variational Implicit Functions. In *26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '99*, 1999. 2
- [35] D. Vlastic, I. Baran, W. Matusik, and J. Popović. Articulated Mesh Animation from Multi-view Silhouettes. *ACM Transactions on Graphics (TOG)*, 27(3), 2008. 1, 6, 7
- [36] A. Wedel, C. Rabe, T. Vaudrey, T. Brox, U. Franke, and D. Cremers. Efficient Dense Scene Flow from Sparse or Dense Stereo Data. In *10th European Conference on Computer Vision (ECCV)*, 2008. 3
- [37] R. T. Whitaker. A Level-Set Approach to 3D Reconstruction from Range Data. *International Journal of Computer Vision (IJCV)*, 29(3):203–231, 1998. 2, 4
- [38] A. Zaharescu, E. Boyer, and R. Horaud. Topology-Adaptive Mesh Deformation for Surface Evolution, Morphing, and Multiview Reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):823–837, 2011. 1