

Hamiltonian Dynamics for Real-World Shape Interpolation

Marvin Eisenberger, Daniel Cremers

Technical University of Munich, Garching, Germany
`marvin.eisenberger@in.tum.de`

Abstract. We revisit the classical problem of 3D shape interpolation and propose a novel, physically plausible approach based on Hamiltonian dynamics. While most prior work focuses on synthetic input shapes, our formulation is designed to be applicable to real-world scans with imperfect input correspondences and various types of noise. To that end, we use recent progress on dynamic thin shell simulation and divergence-free shape deformation and combine them to address the inverse problem of finding a plausible intermediate sequence for two input shapes. In comparison to prior work that mainly focuses on small distortion of consecutive frames, we explicitly model volume preservation and momentum conservation, as well as an anisotropic local distortion model. We argue that, in order to get a robust interpolation for imperfect inputs, we need to model the input noise explicitly which results in an alignment based formulation. Finally, we show a qualitative and quantitative improvement over prior work on a broad range of synthetic and scanned data. Besides being more robust to noisy inputs, our method yields exactly volume preserving intermediate shapes, avoids self-intersections and is scalable to high resolution scans.

Keywords: Shape interpolation, registration, 3D computer vision




Fig. 1. An example interpolation (middle) on real scans from the FAUST dataset [7] and the final overlap (right). Here, the input correspondences were computed with Deep Functional Maps [36] (left).

1 Introduction

Modeling realistic deformations of 3D shapes is at the heart of many computer vision applications. The central motivation in this context is to give meaning to sparse observations of a dynamically moving 3D object. Depending on the application, these measurements are given in the form of a point cloud, a triangle mesh, a voxel grid or a signed distance function.

In many cases, the sampling is not consistent over time and finding commonalities between observations is not a trivial task. While there are a lot of approaches that try to fuse scanned data for 3D reconstruction [45, 53, 54], relatively few work was dedicated to modeling the temporal transformation of the observed object directly. In this work, we revisit this classical challenge of 3D shape interpolation. Although there exists a multitude of elegant formulations, we will show that a lot of these approaches are mainly designed for synthetic shapes and therefore lack robustness to noisy real-world measurements.

The classical formulation is to define an interpolation as a sequence of shapes with minimal local distortion between consecutive frames [12, 28, 32]. While this is undoubtedly a reasonable assumption, it does not suffice in practice to account for the peculiarities of real-world data. For synthetic 3D objects, the ground-truth correspondences are typically known. For real scans, on the other hand, we need to first estimate them, e.g. by using a shape matching method. In practice, the resulting correspondences are not perfect and contain both outliers and fine-scale noise. This is problematic for an interpolation method that minimizes the local distortion between neighboring frames, because the noise from the faulty correspondences tends to distort the local geometry throughout the whole sequence. Moreover, most classical approaches do not model the global geometry of an object which can lead to artifacts like self-intersections.

Contribution We propose a novel framework for real-world shape interpolation that is systematically derived from Hamiltonian dynamics. It resolves the above challenges by introducing additional, physically plausible modeling assumptions like volume preservation and momentum conservation. More specifically, we formulate shape interpolation as the inverse problems of a dynamic thin shell simulation. The Eulerian time-varying deformation fields are represented in a low rank manner which allows us to build volume preservation directly into our model. In qualitative and quantitative experiments, we demonstrate that our method gives rise to high-quality interpolations for real-world inputs.

2 Related work

Shape interpolation has a long tradition in computer graphics. Originally it was developed for planar shapes [2, 42, 52] with [16] being a more recent formalism. A common approach for 3D surfaces is to define an interpolating trajectory as a geodesic in some higher dimensional shape space [12, 26, 27, 60, 61]. Most of these methods use some kind of deformation measure and then optimize for

a sequence such that the local distortion between any two consecutive shapes is low. In [32] this is done with an as-Killing-as-possible energy and in [28] with a discrete shell energy motivated by [24]. Other popular examples of non-linear shape deformation are PriMo [9] and as-rigid-as-possible [56]. For a more thorough introduction to shape spaces, we refer the reader to the book of Younes [63].

An alternative approach to shape interpolation is to interpolate intrinsic quantities like dihedral angles before reconstructing the extrinsic geometry [2, 6, 62]. One class of such intrinsic quantities are rotation-invariant or differential coordinates [1, 34, 35, 51].

Sometimes shape deformation is stated as the time-dependent gradient flow wrt. some surface functional. Typically these functionals promote a smooth flow [14, 15, 18, 58] but most of these methods focus on shape matching with less emphasis on the quality of the intermediate shapes.

Recently, more and more work was dedicated to processing collections of shapes in order to make interpolation more efficient. This can e.g. be achieved by constructing a low-dimensional subspace of admissible poses [3, 21, 29, 57, 65]. In practice, this greatly helps to reduce the computational cost of shape interpolation and even allows for interactive applications [47].

A common assumption of interactive shape deformation modeling is volume preservation. This can be obtained by defining a deformation as the flow of a divergence-free Eulerian vector field [4, 22]. Recently, [19] extended this idea by constructing a divergence-free vector field basis that can be used to interpolate 3D objects. We will make use of this vector field representation and additionally formulate shape interpolation as the inverse problem of a dynamical thin shell simulation. The forward simulation corresponding to this is a well-known problem in computer graphics [44] with applications like cloth [23] or fluid [37] simulation. A recent formulation of this problem that is akin to our approach is projective dynamics [10, 11]. Here, the Lagrangian gradient flow of a dynamical system is restated using the variational form of implicit Euler integration from [38] which leads to an efficient and extremely robust thin shell simulation.

3 Background

We briefly review important preliminary work on shape deformation and interpolation of non-rigidly deforming 3D objects. In this work, we focus on surface-based models like point clouds and 3D meshes. This allows for a compact representation and is in coherence with the output of real-world sensors. In particular, the set of observations $p = (p_1, \dots, p_n)^\top \in \mathbb{R}^{n \times 3}$ consists of n points sampled from a two-dimensional Riemannian manifold \mathcal{X} . Depending on the application, these points are either part of a triangle mesh or embedded in a (knn-)graph.

3.1 Physical assumptions for shape deformation

In order to find similarities between two non-rigid poses of an object, it helps to model geometric assumptions about the expected deformations directly. We

review two common assumptions, namely small local distortions and volume preservation.

Local distortion A popular deformation energy to quantify the distortion between p and a deformed counterpart $p^* \in \mathbb{R}^{n \times 3}$ is the as-rigid-as-possible (arap) energy [56]:

$$\mathcal{W}_{\text{arap}}(p, p^*; (R_i)_{1 \leq i \leq n}) = \frac{1}{2} \sum_{i=1}^n \sum_{j \in \mathcal{N}(i)} \|R_i(p_j - p_i) - (p_j^* - p_i^*)\|_2^2. \quad (1)$$

The assumption behind this functional is that the local deformation of the geometry in the neighborhood $\mathcal{N}(i)$ of every vertex p_i is approximately rigid. I.e. one can find a rotation matrix $R_i \in SO(3)$ that approximately captures the transformation of the neighboring edges $p_j - p_i$. In turn, deviations of the deformation p^* from the approximate rigidity are penalized. The neighborhood $\mathcal{N}(i)$ for a given vertex i is defined as some set of adjacent vertices j to i .

There are multiple popular alternatives with the same flavor as $\mathcal{W}_{\text{arap}}$, including PriMo [9], discrete shells [24] and as-Killing-as-possible [55]. Most techniques penalize deformations of the local geometry and each one of them has certain advantages. In our formulation, we choose the arap energy because it is applicable directly for point clouds and because the optimization for p^* and R_i can be done efficiently in closed form.

Volume preservation Another common assumption for shape deformation is that the volume of the observed object is preserved over time [22]. This can be obtained by prescribing that the deformation is the flow induced by an underlying Eulerian deformation field $v : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ which is divergence-free $\text{div}(v) = 0$. Recently, [19] proposed a formulation of a coarse-to-fine vector field basis that has the volume preservation built in as a hard constraint. A flow field v is then obtained as the linear combination of a finite subset of those divergence-free basis functions:

$$v(x; c) = \sum_{k=1}^K c_k \phi_k(x), \text{ where } \text{div}(\phi_k) = 0. \quad (2)$$

These deformation fields v are exactly volume preserving because the divergence is a linear operator, see [19] for more details. In practice, a relatively small number $K \approx 1000$ of coefficients $c = (c_1, \dots, c_K)^\top \in \mathbb{R}^K$ suffices to represent arbitrary smooth, volume preserving vector fields v . We make use of this compact representation in this work. However, while [19] only considered stationary vector fields $v(x)$, in this work we consider time-dependent vector fields $v(t, x)$ in order to account for more complex shape variation.

3.2 Shape interpolation

Computing an interpolation of two 3D objects $p = p^{(0)}$ and $q = p^{(T)}$ is a common problem in computer graphics and vision. In general, it is not a well-defined

problem because there are typically infinitely many conceivable paths between $p^{(0)}$ and $p^{(T)}$. Therefore, we need to make additional assumptions about plausible sequences like small local distortions or volume preservation. The common way to do this is to define a deformation energy for the whole, time-discrete sequence $p^{(0)}, \dots, p^{(T)}$ of intermediate shapes [12, 27, 28, 31, 32]:

$$E(p^{(1)}, \dots, p^{(T-1)}) = \sum_{t=0}^{T-1} \mathcal{W}(p^{(t)}, p^{(t+1)}). \quad (3)$$

Here, \mathcal{W} is some local distortion measure like $\mathcal{W}_{\text{arap}}$ from Eq. (1). For symmetry reasons, the optimization is commonly done jointly for both the standard and the inverse sequence $p^{(T)}, \dots, p^{(0)}$. W.l.o.g. we will consider the time interval $[0, t_{\max}] = [0, 1]$ which leads to a discrete step size $\tau = \frac{1}{T}$.

4 Interpolation of real-world objects

The implicit assumption behind most shape interpolation approaches is that the exact point-to-point correspondences between the two input surfaces p and q are known. While there is a lot of synthetic data where this is feasible, for scanned data the sampling of two given objects is typically not consistent, even if they approximate the same real-world surface \mathcal{X} . Not even the number of points of the two surfaces $p \in \mathbb{R}^{n \times 3}$ and $q \in \mathbb{R}^{m \times 3}$ is necessarily the same in the most general case. In order to compute an interpolation for this type of input data, we need to first estimate the surface correspondences between p and q .

Computing shape correspondences is a problem in itself and there is a variety of methods that focus on shape matching, either in the classical sense [20, 33, 41, 46, 48, 59] or using machine learning [8, 25, 36, 39, 43, 50]. The output of those




Fig. 2. A qualitative comparison of our approach with other popular shape interpolation methods. Here, we display the intermediate shapes at $t = 0.5$ for an example pair from SCAPE [5] with correspondences from BCICP [48]. Like us, [19] solves for an approximate alignment formulated as an IVP but the stationary vector field leads to slight distortions of the geometry (e.g. at the head and right arm). The other methods [28, 32, 56] solve a BVP and in certain areas the high frequency noise of the correspondences from BCICP leads to a severely degenerate geometry.

methods is a point-to-point assignment of the surface $p \in \mathbb{R}^{n \times 3}$ to $q \in \mathbb{R}^{m \times 3}$ which can be represented with a matrix $\Pi \in \{0, 1\}^{n \times m}$. In principle, we can now transfer the points and neighborhood information from p to q and apply a classical interpolation method like [32] or [28] to $p \in \mathbb{R}^{n \times 3}$ and $\Pi q \in \mathbb{R}^{n \times 3}$. However, in practice the correspondences Π are not perfect and contain faulty or noisy matches. We found that most interpolation methods that assume perfect correspondences are not very robust to fine-scale noise, see Figure 2.

One possible way to make interpolation feasible for real scans is to acknowledge that the given matching Π is not perfect and to build this stochastic discrepancy directly into our model. In particular, we add Gaussian random noise η to the vertex position of the second shape Πq :

$$\tilde{q} := \Pi q + \eta, \text{ with } \eta \sim \mathcal{N}(0, \sigma). \quad (4)$$

Instead of finding intermediate shapes by solving a boundary value problem (BVP) as outlined in Eq. (3), we can then define an initial value problem (IVP) similar to [19]. In particular, we will optimize for a sequence $p^{(0)}, \dots, p^{(T)}$ with $p = p^{(0)}$ and $p^{(T)} = \tilde{q} \approx \Pi q$.

5 From Hamiltonian dynamics to Eulerian-Lagrangian shape interpolation

In this work, we model the motions of objects in an inertial frame of reference as a physical phenomenon that is governed by three aspects: internal forces, momentum conservation and volume preservation. Most existing interpolation techniques model internal forces in some way, yet they omit the momentum conservation and volume preservation. Without momentum conservation, the intermediate objects can be plausible but in many cases the motions lack temporal coherence. The volume preservation helps to constrain the optimization and prevents self-intersections, see Figure 3. Our formulation combines the strengths of volume preserving fields [19] and projective dynamics [10] with those of classical interpolation methods [28, 32].

5.1 Deformation model

We systematically derive the evolution of a surface as a physical system from the Hamiltonian energy given by:

$$\mathcal{H}(p, v) = \frac{1}{2} \|v\|_2^2 + \mathcal{W}(p). \quad (5)$$

This energy consists of a kinetic energy term that models momentum conservation (with unit mass per point) and some potential energy component \mathcal{W} that penalizes intrinsic distortions. The principles of Hamiltonian mechanics now prescribe how this system evolves over time:

$$\begin{cases} \dot{p} = \frac{d\mathcal{H}}{dv} = v. \\ \dot{v} = -\frac{d\mathcal{H}}{dp} = -\nabla \mathcal{W}(p). \end{cases} \quad (6)$$

We couple this with the volume preservation assumption by constraining v to the low rank vector field representation from Eq. (2). This allows us to model displacements of a shape $p(t) = (p_1(t), \dots, p_n(t))^T$ at time t with only $K \ll n$ degrees of freedom:

$$\dot{p}_i(t) = v(p_i(t); c(t)) = \sum_{k=1}^K c_k(t) \phi_k(p_i(t)). \quad (7)$$

Besides providing a compact representation, this approach builds volume preservation directly into the deformation model, because $\text{div}(v) = 0$. In [19], the authors model shape deformations in a similar way but with a stationary vector field $v(x; c(t)) = v(x; c)$. This leads to a well-constrained optimization problem with only K degrees of freedom c_1, \dots, c_K but it is also restrictive and lacks expressivity. Instead of using a constant vector field, following Eq. (6), we define a dynamic flow $v(t, x) = v(x; c(t))$:

$$\begin{cases} \dot{v}(t, p(t)) = -\nabla \mathcal{W}(p(t)). \\ \text{div}(v) = 0. \end{cases} \quad (8)$$

In our formulation, the internal forces are defined as the negative gradient of our anisotropic as-rigid-as-possible potential \mathcal{W} which we define in the next chapter.




Fig. 3. A pair of synthetic shapes with ground-truth correspondences from the KIDS dataset [50] for which we show the intermediate shapes at $t = 0.5$. This example shows that many classical methods like [32] or [28] cannot detect self-intersections of different subparts. Here, the optimal path that minimizes a local distortion metric makes the right arm of the kid move through itself. Our method, on the other hand, avoids self-intersections by design: All deformations are expressed as a divergence-free Eulerian field, therefore the resulting flow has to be globally consistent in the sense that two close parts cannot have contradictory motions.

5.2 Anisotropic as-rigid-as-possible deformation

For most 3D objects, not all parts are behaving similar in terms of local distortions. For example, regions near joints of a human body allow for more movement than most other parts of the surface. The classical as-rigid-as-possible potential that we reviewed in Eq. (1) penalizes distortions of the geometry uniformly in all directions and equal for all parts of the considered object. We generalize this idea and introduce an anisotropic as-rigid-as-possible energy:

$$\mathcal{W}(p(t); (R_i)_i, (\Sigma_i)_i) = \frac{1}{2} \sum_{i=1}^n \sum_{j \in \mathcal{N}(i)} \left\| (p_j(0) - p_i(0)) - R_i^\top (p_j(t) - p_i(t)) \right\|_{\Sigma_i}^2. \quad (9)$$

In this context, $\|\cdot\|_{\Sigma_i}$ denotes the standard Mahalanobis norm [40] with an unknown covariance matrix $\Sigma_i \in \mathbb{R}^{3 \times 3}$. This energy \mathcal{W} allows our model to adapt the appropriate local behavior during the optimization, see Figure 4 for an example. Moreover, the distortion is always computed in the reference frame of the first pose $p(0)$. This means that we only need to compute the distortion model of $p(0)$ and therefore we only need one local distortion matrix per vertex Σ_i for the whole sequence.




Fig. 4. An example from TOSCA where we color code the log-determinant of the covariance matrix $\log \det(\Sigma_i)$ for each point p_i . Red stands for a low value which corresponds to a high local distortion. This shows how our anisotropic as-rigid-as-possible energy (9) automatically adapts to objects consisting of inhomogeneous parts. Certain regions like joints allow for more local distortion throughout the sequence than others. Notice the difference between the hind legs, the head and the rest of the body.

5.3 Time discretization

In the time-discrete setting, we can approximate Eq. (7) and Eq. (8) using an implicit Euler integration scheme:

$$\begin{cases} p^{(t+1)} = p^{(t)} + \tau v^{(t+1)}. & (10a) \\ v^{(t+1)} = v^{(t)} - \tau \nabla \mathcal{W}(p^{(t+1)}). & (10b) \\ \operatorname{div}(v^{(t+1)}) = 0. & (10c) \end{cases}$$

This is a Eulerian-Lagrangian scheme: The velocity field is represented on the surface $v^{(t)} \in \mathbb{R}^{n \times 3}$ but the divergence-free condition $\operatorname{div}(v^{(t+1)}) = 0$ is Eulerian. In order to make this interaction tractable, we will use the divergence-free vector field representation from Eq. (2) and combine it with the variational form of implicit Euler integration introduced in [38]. This allows us to restate this scheme as an optimization problem in terms of the vector field coefficients $c \in \mathbb{R}^K$:

$$\begin{cases} c^{(t+1)} = \arg \min_{c, R} \left\| v(p^{(t)}; c) - \bar{v}^{(t)} \right\|_F^2 + \mathcal{W}\left(p^{(t)} + \tau v(p^{(t)}; c); R, \Sigma\right). & (11a) \end{cases}$$

$$\begin{cases} v_i^{(t+1)} = v(p_i^{(t)}; c^{(t+1)}) = \sum_{k=1}^K c_k^{(t+1)} \phi_k(p_i^{(t)}). & (11b) \end{cases}$$

$$\begin{cases} p^{(t+1)} = p^{(t)} + \tau v^{(t+1)}. & (11c) \end{cases}$$

$$\begin{cases} \bar{v}^{(t+1)} = 2v^{(t+1)} - v^{(t)}. & (11d) \end{cases}$$

We refer the interested reader to [38] and [10] for more details on how this scheme is derived. The update of the coefficients c in (11a) can be computed using Gauss-Newton optimization. We use an additional extrapolation step (11d) to get a better prediction of the velocity $v^{(t+2)}$ which we justify in the following:

Theorem 1. *For continuously differentiable vector fields, the extrapolation step (11d) of Algorithm 11 yields an estimate $\bar{v}^{(t+1)}$ of $v^{(t+2)}$ with an error of order $\mathcal{O}(\tau^2)$. For the alternative scheme without step (11d) it is $\mathcal{O}(\tau)$.*

This result implies that (11d) leads to a qualitative improvement because a better estimate $\bar{v}^{(t+1)} \approx v^{(t+2)}$ provides a more faithful approximation in the next update step (11a) of c . See Appendix A for a proof of Thm. 1.

5.4 Interpolation algorithm

We will now use the scheme (11) from last chapter to define an interpolation algorithm for two given shapes p and q . In each iteration, we initialize the scheme with $p^{(0)} := p$ and the unknown variables $c^{(0)} := \hat{c}$ and $(\hat{\Sigma}_i)_{1 \leq i \leq n}$. We then compute the deformed shapes $p^{(0)}, \dots, p^{(T)}$ according to our scheme (11). Overall, this forward pass can be summarized as the differentiable solution operator \mathcal{S} :

$$\mathcal{S} : \begin{cases} \mathbb{R}^K \times \mathbb{R}^{n \times 3 \times 3} \rightarrow \mathbb{R}^{n \times 3 \times (T+1)}. \\ (\hat{c}, \hat{\Sigma}) \mapsto (p^{(0)}, \dots, p^{(T)}). \end{cases} \quad (12)$$

The goal is now to find the input parameters \hat{c} and $\hat{\Sigma}$ that lead to a tight alignment of the deformed shape $p^{(T)}$ with q in accordance with Eq. (4). Together with our regularizer \mathcal{W} from Eq. (9) this leads to the following energy:

$$E(p^{(0)}, \dots, p^{(T)}; \Pi) := \frac{1}{2\sigma^2} \|p^{(T)} - \Pi q\|^2 + \sum_{t=0}^T \mathcal{W}(p^{(t)}). \quad (13)$$

Putting everything together, we can derive the following algorithm:

Algorithm 1 Volume preserving shape interpolation.

Require: $p \in \mathbb{R}^{n \times 3}$, $q \in \mathbb{R}^{m \times 3}$
 $\hat{c} \leftarrow 0 \in \mathbb{R}^K$
 $\hat{\Sigma}_i \leftarrow \text{Id}_3 \in \mathbb{R}^{3 \times 3}$
 $\Pi \leftarrow \text{match_shapes}(p, q) \in \{0, 1\}^{n \times m}$
for $i = 1, \dots, N_{\text{it}}$ **do**
 $(\hat{c}, \hat{\Sigma}) \leftarrow (\hat{c}, \hat{\Sigma}) - \gamma \nabla E(\mathcal{S}(\hat{c}, \hat{\Sigma}); \Pi)$
end for
return $(p^{(0)}, \dots, p^{(T)}) := \mathcal{S}(\hat{c}, \hat{\Sigma})$

In our implementation, we use a modern automatic differentiation toolbox to compute the gradient $\nabla E \circ \mathcal{S}$ wrt. $(\hat{c}, \hat{\Sigma})$ in Algorithm 1. The choice of algorithm to compute the input correspondences Π is not further specified here because it is more or less arbitrary. We show various different possibilities in our experiments.




Fig. 5. Two interpolated sequences for real scans of a puppet from SHREC’19 Isometry [17] and a human from FAUST [7] with input correspondences from Smooth Shells [20] and Deep Functional Maps [36] respectively.

6 Experiments

We verify the generality of our method on four different datasets with increasing complexity. The first two are the synthetic datasets TOSCA [13] and SCAPE [5] where we use the ground truth correspondences for the former and correspondences from BCICP [48] for the latter. The last two datasets SHREC’19 Isometry [17] and FAUST [7] contain real scans of a puppet and different humans respectively, see Figure 5. For those, we use correspondences from Smooth Shells [20] and FMNet [36]. Our experiments show that our formulation is applicable to a wide range of inputs with varying levels of noise. Figure 7 summarizes our quantitative evaluations on all datasets with comparisons to four other popular interpolation methods. The other methods are Geometric Modeling in Shape Space [32], Time-Discrete Geodesics in the Space of Shells [28], Divergence-Free Shape Correspondence by Deformation [19] and As-Rigid-As-Possible Surface Modeling [56]. Although the latter does not describe an interpolation algorithm explicitly, it is trivial to employ its shape deformation procedure in an interpolation pipeline by using Eq. (3). On the surface, our method is similar to [19] in the sense that both approaches compute divergence-free fields in a low rank basis. The decisive difference is that our method is based on a physically plausible formulation which, among other things, allows for time-dependent vector fields $v(t, x)$. This makes our method more expressive, see Figure 6 for an example.

Error metrics In order to quantify the precision of a shape interpolation, we compute three different metrics for each pair of input shapes and plot the resulting cumulative curves in Figure 7. In particular, we measure the conformal distortion [64, Eq. (6)] and volume change [30, Eq. (3)] of intermediate shapes and the Chamfer distance to the target shapes in % of the diameter for our




Fig. 6. A comparison of our method and divergence-free interpolation [19] on a pair of synthetic shapes (green). Both methods preserve the volume but for this large scale deformation the stationary vector field in [19] is too restrictive which leads to a distorted geometry for $t \geq 0.5$

method and the second alignment based method [19]. If we are strict, the notion of volume change is only meaningful for watertight meshes, which typically does not hold for real scans. Our argument regarding this is that in theory, a flow induced by a divergence-free deformation field is exactly volume preserving in terms of the underlying watertight real-world manifold \mathcal{X} . Remarkably, in this way we can even make sense of the notion of volume for a point cloud, assuming that it was sampled from a closed, continuous surface.

Implementation details The low rank vector field representation of divergence-free fields in our Scheme 11 is entirely decoupled of the input resolutions n and m . Moreover, the vector fields are represented in a spatially dense Eulerian basis

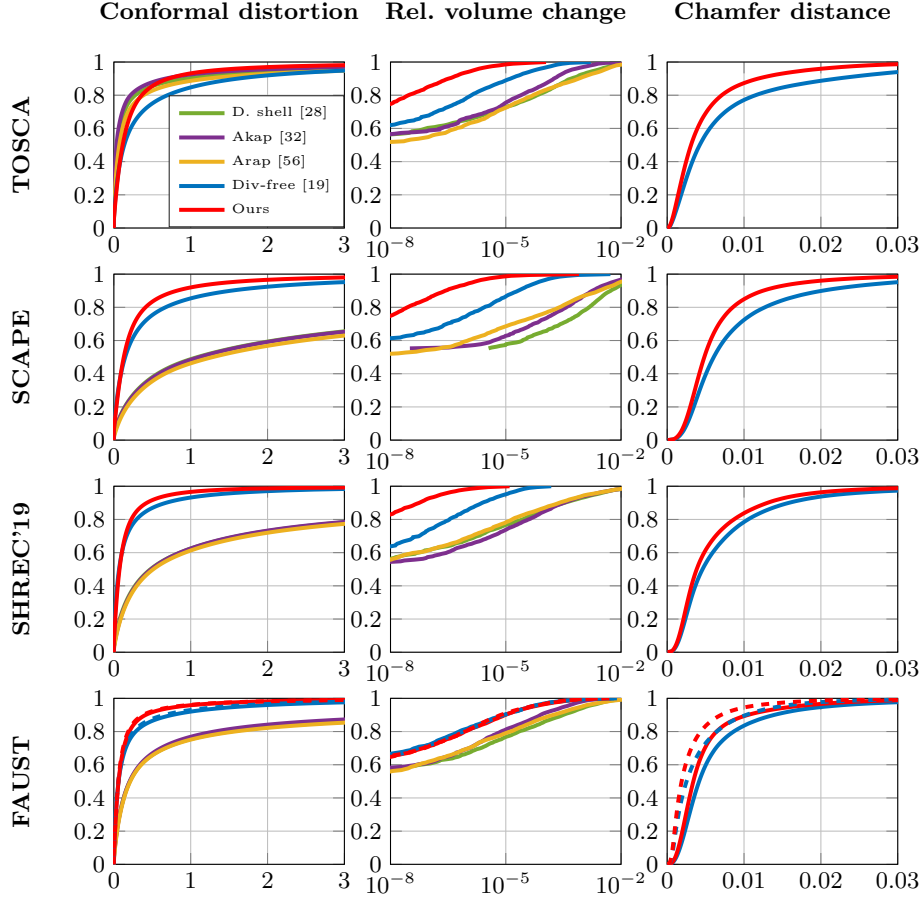


Fig. 7. Quantitative results and comparisons with other methods on four benchmarks. For FAUST, dashed lines correspond to the results on the high resolution scans. Those were only computed for our method and [19] because for the other methods the resolution of around 200k vertices is prohibitively high.

which means that at any discrete time t , the resulting vector field $v(x; c^{(t)})$ can be computed for arbitrary points x in our domain, see [19] for more details. This allows us to efficiently perform the optimization in Algorithm 1 on a subsampled version of the input shapes p and q with a fixed resolution of $2k$ points. Afterwards, the computed vector field can be applied to the full resolution in a single forward pass without any skinning strategy or the like. For once, this makes our approach significantly faster but it also allows for an interpolation of very high resolution objects like those from FAUST ($\sim 200k$ vertices). Many other classical interpolation methods use some multiscale scheme to allow for higher resolutions [28, 32], but there are still upper limits for them as to what is feasible in terms of computation cost. Our interpolation Algorithm 1 is directly applicable to point clouds, therefore we simply subsample both input shapes using Euclidean farthest point sampling. However, other subsampling strategies like remeshing are also possible if one wants to work directly with meshes. Finally, we use the same set of parameters for all experiments, see our implementation for details.



Fig. 8. We show how our method can be used to refine an imperfect shape correspondence. Using the input matching from Deep Functional Maps [36], we compute an interpolation (left half) and use it to recover the improved correspondences using the final alignment at $t = 1$ (right half). We display the matching with a texture map from the first input shape (3rd human from right) to the final pose with both methods.




Fig. 9. An example of how our approach can be used to extrapolate the motion prescribed by the two input frames $t = 0$ and $t = 1$. The sequences obtained with our method are physically plausible and remain stable over a long period of time. The cat keeps raising its paw until at $t = 2$, driven by the regularizer (9), the motion reverses.

Additional evaluations As a proof of concept, we show that our physically plausible formulation allows for a broad range of applications beyond shape interpolation. For once, we can use our alignment at $t = 1$ to refine a shape matching which we show for a real scan of FAUST in Figure 8. Furthermore, we can compute plausible shape extrapolations by simply simulating the forward integration for a longer period of time than $t = 1$. Remarkably, this can be done without any additional optimization, we simply compute an interpolation between p at $t = 0$ and q at $t = 1$ and then integrate our Scheme 11 until $t > 1$, see Figure 9. Finally, we show that our method allows for input objects where only parts of the geometry are available, see Figure 10. This is only feasible for an alignment based method, because the classical formulation as a BVP requires that every vertex has a corresponding point on the other surface. Partial shape interpolation is an important preliminary result for many real world applications like scanning of dynamically moving 3D objects.




Fig. 10. An example interpolation of a pair of partial shapes from the synthetic TOSCA cuts [49] dataset with our method.

7 Conclusion

We presented a general and flexible approach to shape interpolation that is systematically derived from a formulation of Hamiltonian dynamics. For this, we employ recent advances in dynamic thin shell simulation to get a robust deformation model and solve its inverse problem by optimizing over the initial motion and anisotropic surface properties. We demonstrated that, in comparison to prior work, our approach is able to compute high quality, physically plausible interpolations of noisy real world inputs. In future work, we will apply our setup to a broader range of applications like 3D scanning of actions or mesh compression.

Acknowledgements We would like to thank Zorah Löhner and Aysim Toker for useful discussions. We gratefully acknowledge the support of the ERC Consolidator Grant "3D Reloaded".

References

1. Alexa, M.: Differential coordinates for local mesh morphing and deformation. *The Visual Computer* **19**(2-3), 105–114 (2003)
2. Alexa, M., Cohen-Or, D., Levin, D.: As-rigid-as-possible shape interpolation. In: *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. pp. 157–164. SIGGRAPH '00 (2000)
3. Allen, B., Curless, B., Popović, Z., Hertzmann, A.: Learning a correlated model of identity and pose-dependent body shape variation for real-time synthesis. In: *Proceedings of the 2006 ACM SIGGRAPH/Eurographics symposium on Computer animation*. pp. 147–156. Eurographics Association (2006)
4. Angelidis, A., Cani, M.P., Wyvill, G., King, S.: Swirling-sweepers: Constant-volume modeling. *Graphical Models* **68**(4), 324–332 (2006)
5. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. In: *ACM transactions on graphics (TOG)*. vol. 24, pp. 408–416. ACM (2005)
6. Baek, S.Y., Lim, J., Lee, K.: Isometric shape interpolation. *Computers & Graphics* **46**, 257–263 (2015)
7. Bogo, F., Romero, J., Loper, M., Black, M.J.: FAUST: Dataset and evaluation for 3D mesh registration. In: *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Piscataway, NJ, USA (Jun 2014)
8. Boscaini, D., Masci, J., Rodolà, E., Bronstein, M.: Learning shape correspondence with anisotropic convolutional neural networks. In: *Advances in neural information processing systems*. pp. 3189–3197 (2016)
9. Botsch, M., Pauly, M., Gross, M.H., Kobbelt, L.: Primo: coupled prisms for intuitive surface modeling. In: *Symposium on Geometry Processing*. pp. 11–20. No. CONF (2006)
10. Bouaziz, S., Martin, S., Liu, T., Kavan, L., Pauly, M.: Projective dynamics: fusing constraint projections for fast simulation. *ACM Transactions on Graphics (TOG)* **33**(4), 1–11 (2014)
11. Brandt, C., Eisemann, E., Hildebrandt, K.: Hyper-reduced projective dynamics. *ACM Transactions on Graphics (TOG)* **37**(4), 1–13 (2018)
12. Brandt, C., von Tycowicz, C., Hildebrandt, K.: Geometric flows of curves in shape space for processing motion of deformable objects. In: *Computer Graphics Forum*. vol. 35, pp. 295–305. Wiley Online Library (2016)
13. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: *Numerical geometry of non-rigid shapes*. Springer (2008), http://tosca.cs.technion.ac.il/book/resources_data.html
14. Charpiat, G., Faugeras, O., Keriven, R.: Approximations of shape metrics and application to shape warping and empirical shape statistics. *Foundations of Computational Mathematics* **5**(1), 1–58 (2005)
15. Charpiat, G., Maurel, P., Pons, J.P., Keriven, R., Faugeras, O.: Generalized gradients: Priors on minimization flows. *International journal of computer vision* **73**(3), 325–344 (2007)
16. Chen, R., Weber, O., Keren, D., Ben-Chen, M.: Planar shape interpolation with bounded distortion. *ACM Transactions on Graphics (TOG)* **32**(4), 1–12 (2013)
17. Dyke, R., Stride, C., Lai, Y., Rosin, P.: Shrec-19: Shape correspondence with isometric and non-isometric deformations (2019)
18. Eckstein, I., Pons, J.P., Tong, Y., Kuo, C.C., Desbrun, M.: Generalized surface flows for mesh processing. In: *Proceedings of the fifth Eurographics symposium on Geometry processing*. pp. 183–192. Eurographics Association (2007)

19. Eisenberger, M., Löhner, Z., Cremers, D.: Divergence-free shape correspondence by deformation. In: *Computer Graphics Forum*. vol. 38, pp. 1–12. Wiley Online Library (2019)
20. Eisenberger, M., Löhner, Z., Cremers, D.: Smooth shells: Multi-scale shape registration with functional maps. *arXiv preprint arXiv:1905.12512* (2019)
21. Fletcher, P.T., Lu, C., Pizer, S.M., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE transactions on medical imaging* **23**(8), 995–1005 (2004)
22. von Funck, W., Theisel, H., Seidel, H.P.: Vector field based shape deformations. In: *ACM Transactions on Graphics (TOG)*. vol. 25, pp. 1118–1125. ACM (2006)
23. Goldenthal, R., Harmon, D., Fattal, R., Bercovier, M., Grinspun, E.: Efficient simulation of inextensible cloth. In: *ACM SIGGRAPH 2007 papers*, pp. 49–es (2007)
24. Grinspun, E., Hirani, A.N., Desbrun, M., Schröder, P.: Discrete shells. In: *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*. pp. 62–67. Eurographics Association (2003)
25. Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: 3d-coded: 3d correspondences by deep deformation. In: *The European Conference on Computer Vision (ECCV)* (September 2018)
26. Heeren, B., Rumpf, M., Schröder, P., Wardetzky, M., Wirth, B.: Exploring the geometry of the space of shells. In: *Computer Graphics Forum*. vol. 33, pp. 247–256. Wiley Online Library (2014)
27. Heeren, B., Rumpf, M., Schröder, P., Wardetzky, M., Wirth, B.: Splines in the space of shells. *Computer Graphics Forum* **35**(5), 111–120 (2016)
28. Heeren, B., Rumpf, M., Wardetzky, M., Wirth, B.: Time-discrete geodesics in the space of shells. In: *Computer Graphics Forum*. vol. 31, pp. 1755–1764. Wiley Online Library (2012)
29. Heeren, B., Zhang, C., Rumpf, M., Smith, W.: Principal geodesic analysis in the space of discrete shells. In: *Computer Graphics Forum*. vol. 37, pp. 173–184. Wiley Online Library (2018)
30. Hormann, K., Greiner, G.: Mips: An efficient global parametrization method. Tech. rep., Erlangen-Nuernberg University (Germany) Computer Graphics Group (2000)
31. Huber, P., Perl, R., Rumpf, M.: Smooth interpolation of key frames in a riemannian shell space. *Computer Aided Geometric Design* **52**, 313–328 (2017)
32. Kilian, M., Mitra, N.J., Pottmann, H.: Geometric modeling in shape space. In: *ACM Transactions on Graphics (TOG)*. vol. 26, p. 64. ACM (2007)
33. Kim, V.G., Lipman, Y., Funkhouser, T.A.: Blended intrinsic maps. *Transactions on Graphics (TOG)* **30**(4) (2011)
34. Lipman, Y., Sorkine, O., Cohen-Or, D., Levin, D., Rossi, C., Seidel, H.P.: Differential coordinates for interactive mesh editing. In: *Proceedings Shape Modeling Applications*, 2004. pp. 181–190. IEEE (2004)
35. Lipman, Y., Sorkine, O., Levin, D., Cohen-Or, D.: Linear rotation-invariant coordinates for meshes. *ACM Transactions on Graphics (TOG)* **24**(3), 479–487 (2005)
36. Litany, O., Remez, T., Rodolà, E., Bronstein, A., Bronstein, M.: Deep functional maps: Structured prediction for dense shape correspondence. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 5659–5667 (2017)
37. Macklin, M., Müller, M.: Position based fluids. *ACM Transactions on Graphics (TOG)* **32**(4), 1–12 (2013)
38. Martin, S., Thomaszewski, B., Grinspun, E., Gross, M.: Example-based elastic materials. In: *ACM SIGGRAPH 2011 papers*, pp. 1–8 (2011)

39. Masci, J., Boscaini, D., Bronstein, M., Vandergheynst, P.: Geodesic convolutional neural networks on riemannian manifolds. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 37–45 (2015)
40. McLachlan, G.J.: Mahalanobis distance. *Resonance* **4**(6), 20–26 (1999)
41. Melzi, S., Ren, J., Rodola, E., Ovsjanikov, M., Wonka, P.: Zoomout: Spectral upsampling for efficient shape correspondence. arXiv preprint arXiv:1904.07865 (2019)
42. Michor, P.W., Mumford, D.: Riemannian geometries on spaces of plane curves. arXiv preprint math/0312384 (2003)
43. Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., Bronstein, M.M.: Geometric deep learning on graphs and manifolds using mixture model cnns. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5115–5124 (2017)
44. Müller, M., Heidelberger, B., Hennix, M., Ratcliff, J.: Position based dynamics. *Journal of Visual Communication and Image Representation* **18**(2), 109–118 (2007)
45. Newcombe, R.A., Fox, D., Seitz, S.M.: Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 343–352 (2015)
46. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)* **31**(4), 30 (2012)
47. von Radziewsky, P., Eisemann, E., Seidel, H.P., Hildebrandt, K.: Optimized subspaces for deformation-based modeling and shape interpolation. *Computers & Graphics* **58**, 128–138 (2016)
48. Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graph.* **37**(6), 248:1–248:16 (Dec 2018)
49. Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. In: *Computer Graphics Forum*. vol. 36, pp. 222–236. Wiley Online Library (2017)
50. Rodolà, E., Rota Bulò, S., Windheuser, T., Vestner, M., Cremers, D.: Dense non-rigid shape correspondence using random forests. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014)
51. Sassen, J., Heeren, B., Hildebrandt, K., Rumpf, M.: Geometric optimization using nonlinear rotation-invariant coordinates. arXiv preprint arXiv:1908.11728 (2019)
52. Sederberg, T.W., Gao, P., Wang, G., Mu, H.: 2-d shape blending: an intrinsic solution to the vertex path problem. In: Proceedings of the 20th annual conference on Computer graphics and interactive techniques. pp. 15–18 (1993)
53. Slavcheva, M., Baust, M., Cremers, D., Ilic, S.: Killingfusion: Non-rigid 3d reconstruction without correspondences. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1386–1395 (2017)
54. Slavcheva, M., Baust, M., Ilic, S.: Sobolevfusion: 3d reconstruction of scenes undergoing free non-rigid motion. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2646–2655 (2018)
55. Solomon, J., Ben-Chen, M., Butscher, A., Guibas, L.: As-killing-as-possible vector fields for planar deformation. In: *Computer Graphics Forum*. vol. 30, pp. 1543–1552. Wiley Online Library (2011)
56. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: *Symposium on Geometry processing*. vol. 4, pp. 109–116 (2007)
57. Sumner, R.W., Zwicker, M., Gotsman, C., Popović, J.: Mesh-based inverse kinematics. *ACM transactions on graphics (TOG)* **24**(3), 488–495 (2005)

58. Sundaramoorthi, G., Yezzi, A., Mennucci, A.C.: Sobolev active contours. *International Journal of Computer Vision* **73**(3), 345–366 (2007)
59. Vestner, M., Lähner, Z., Boyarski, A., Litany, O., Slossberg, R., Remez, T., Rodolà, E., Bronstein, A.M., Bronstein, M.M., Kimmel, R., Cremers, D.: Efficient deformable shape correspondence via kernel matching. In: *International Conference on 3D Vision (3DV)* (October 2017)
60. Wirth, B., Bar, L., Rumpf, M., Sapiro, G.: A continuum mechanical approach to geodesics in shape space. *International Journal of Computer Vision* **93**(3), 293–318 (2011)
61. Wirth, B., Bar, L., Rumpf, M., Sapiro, G.: A continuum mechanical approach to geodesics in shape space. *International Journal of Computer Vision* **93**(3), 293–318 (Jul 2011)
62. Xu, D., Zhang, H., Wang, Q., Bao, H.: Poisson shape interpolation. In: *Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling*. pp. 267–274 (2005)
63. Younes, L.: *Shapes and diffeomorphisms*, vol. 171. Springer (2010)
64. Zhang, C., Chen, T.: Efficient feature extraction for 2d/3d objects in mesh representation. In: *Proceedings 2001 International Conference on Image Processing (Cat. No. 01CH37205)*. vol. 3, pp. 935–938. IEEE (2001)
65. Zhang, C., Heeren, B., Rumpf, M., Smith, W.A.: Shell pca: Statistical shape modelling in shell space. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1671–1679 (2015)