

Geodesics in Shape Space via Variational Time Discretization




Fig. 1. Time-discrete geodesics between a cat and a lion and the letters A and B . Geodesic distance is measured on the basis of viscous dissipation inside the objects (color-coded in the middle row from blue, low dissipation, to red, high dissipation), which is induced by a pairwise 1-1 deformation map between consecutive shapes along the discrete geodesic path. Shapes are represented via level set functions, whose level lines are texture-coded in the bottom row for the 2D example.

Benedikt Wirth[†], Leah Bar[‡], Martin Rumpf[†], and Guillermo Sapiro[‡]

[†]Institute for Numerical Simulation, University of Bonn, Germany

[‡]Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, U.S.A.

Abstract. A variational approach to defining geodesics in the space of implicitly described shapes is introduced in this paper. The proposed framework is based on the time discretization of a geodesic path as a sequence of pairwise matching problems, which is strictly invariant with respect to rigid body motions and ensures a 1-1 property of the induced flow in shape space. For decreasing time step size, the proposed model leads to the minimization of the actual geodesic length, where the Hessian of the pairwise matching energy reflects the chosen Riemannian metric on the shape space. Considering shapes as boundary contours, the proposed shape metric is identical to a physical dissipation in a viscous fluid model of optimal transportation. If the pairwise shape correspondence is replaced by the volume of the shape mismatch as a penalty functional, for decreasing time step size one obtains an additional optical flow term controlling the transport of the shape by the underlying motion field. The implementation of the proposed approach is based on a level set representation of shapes, which allows topological transitions along the geodesic path. For the spatial discretization a finite element approximation is employed both for the pairwise deformations and for the level set representation. The numerical relaxation of the energy is performed via an efficient multi-scale procedure in space and time. Examples for 2D and 3D shapes underline the effectiveness and robustness of the proposed approach.

1 Introduction

This paper deals with the computation of geodesic paths and distances between (possibly non-rigid) shapes represented via *level sets* in 2D and 3D. Such computations are fundamental for problems ranging from computational anatomy to object recognition, warping, and matching. The aim is to reliably and effectively evaluate distances between non-parametrized geometric shapes of possibly different topology. We investigate the close link between abstract geometry on the infinite-dimensional space of shapes and the continuum mechanical viewpoint of shapes as being boundary contours of physical objects, e.g. identifying the Riemannian metric on shape space with the physical dissipation — the loss of energy due to friction. Thereby, we simultaneously address the following major challenges:

- a physically sound modeling of the geodesic flow of shapes given as boundary contours of objects on a void background,
- the need for a coarse time discretization which is nevertheless invariant with respect to rigid body motions, ensures a 1-1 object correspondence, and relates to the corresponding continuous geodesic path,
- a numerically effective multi-scale treatment of the resulting time and space discrete energy.

Our approach is closely linked to the concept of optimal transportation [1]. The motion field v governing the flow in shape space vanishes outside the object bounded by the corresponding shape contour. The field is optimal in the sense that it minimizes an accumulated physical dissipation — a quadratic functional depending on the first order local variation of a flow field, representing the rate at which mechanical energy is converted into heat in a viscous fluid per unit volume. Thus, the Riemannian metric on the shape space is defined to coincide with this rate of dissipation. If we assume frame indifference as first principle (rigid body motion invariance), then the dissipation depends only on the symmetric part $\epsilon[v] = \frac{1}{2}(\mathcal{D}v^T + \mathcal{D}v)$ of the Jacobian $\mathcal{D}v$ of the underlying motion field v . Under the additional assumption of isotropy, a typical model for the local rate of dissipation is given by $\mathbf{Diss}[v] = \int_0^1 \int_{\mathcal{O}} g(v, v) dx dt$ with

$$g(v, v) = \frac{\lambda}{2} (\text{tr}\epsilon[v])^2 + \mu \text{tr}(\epsilon[v]^2) \quad (1)$$

(cf. Fuchs et al. [2]), where \mathcal{O} describes the deformed object. Here $\text{tr}(\epsilon[v]^2)$ measures the averaged local change of length and $(\text{tr}\epsilon[v])^2$ the local change of volume (obviously $\text{div } v = \text{tr}(\epsilon[v]) = 0$ represents an incompressible flow), induced by the transport by v . In their pioneering paper Miller et al. [3] exploited the fact that in case of sufficient Sobolev regularity for the motion field v on the whole surrounding domain, the induced flow consists of a family of diffeomorphisms. A straightforward time discretization of a geodesic flow would neither guarantee local rigid body motion invariance for the time discrete problem nor a 1-1 mapping property between objects at consecutive time steps.

In this paper, we present a time discretization of the squared path length in shape space which is based on a pairwise matching of intermediate shapes cor-

responding to subsequent time steps. In fact, such a discretization of a path as concatenation of short connecting lines between consecutive points on the path is most natural with regard to the variational definition of a geodesic and for instance underlies the algorithm by Schmidt et al. [4]. Our approach is inspired both by work in mechanics [5] and in geometry [6]. Here, a suitable deformation energy will measure the deformation between subsequent shapes. This can be regarded as the infinite-dimensional counterpart of the following time discretization for a geodesic between two points s_A and s_B on a finite-dimensional Riemannian manifold: Consider a sequence of points $s^A = s_0, s_1, \dots, s_K = s_B$ connecting two fixed points s_A and s_B and minimize $\sum_{k=1}^K \text{dist}^2(s_{k-1}, s_k)$, where $\text{dist}(\cdot, \cdot)$ is a suitable approximation of the Riemannian distance. In our case, the squared approximate distance is replaced by the deformation energy, for which we will employ a particular class of so-called polyconvex energies [7] to ensure both exact frame indifference (observer independence and thus rigid body motion invariance) and a global 1-1 property. We will also discuss the corresponding continuous problem when the time discretization step vanishes.

Even though the functionals are borrowed from nonlinear elasticity, the underlying physics is only related to elasticity in the sense that a viscous deformation can be regarded as the limit of infinitely small elastic deformations with subsequent stress relaxation. Indeed, different from elasticity, none of the shapes is in a stressed configuration since local stresses are immediately absorbed via dissipation, which in a physical context reflects a local heating.





Fig. 2. Discrete geodesics between a straight and a rolled up bar, from first row to fourth row based on 1, 2, 4, and 8 time steps. The light gray shapes in the first row show the linear interpolation of the deformations connecting the dark gray shapes. The shapes from the finest time discretization are overlaid over the others as thin black lines. In the last row the rate of viscous dissipation is rendered on the shape domains $\mathcal{O}_1, \dots, \mathcal{O}_{K-1}$ from the previous row, color-coded as .

Both the built-in exact frame indifference and the 1-1 mapping property ensure that fairly coarse time discretizations already lead to an accurate approximation of geodesic paths (cf. Fig. 2). The actual convergence is dealt with later in this paper. Careful consideration is required with respect to the effective minimization of the time discrete path length. Already in the case of low dimensional Riemannian manifolds the need for efficient minimization strategies is apparent. To give a conceptual sketch of the proposed algorithm on the actual shape space, Fig. 3 depicts the proposed procedure in the case of \mathbb{R}^2 considered as the stereographic projection of the two-dimensional sphere and outlines the advantage of our proposed optimization framework.




Fig. 3. Different refinement levels of discrete geodesics ($K = 1, 2, 4, \dots, 256$) from Johannesburg to Kyoto in the stereographic projection (right) and backprojected on the globe (left). A single-level nonlinear Gauss-Seidel on the finest resolution with successive relaxation of the different vertices requires 917235 elementary relaxation steps, whereas in a cascadic relaxation from coarse to fine resolution in time, only 2593 of these elementary minimization steps are needed.

1.1 Related work

Conceptually, in the last decade, the distance between shapes has been studied on the basis of a general framework of a space of shapes and its intrinsic structure. The notion of shape space has been introduced already in 1984 by Kendall [8].

An isometrically invariant distance measure between two objects \mathcal{S}_A and \mathcal{S}_B in (different) metric spaces is the Gromov–Hausdorff distance, which is (in a simplified form) defined as the minimizer of $\frac{1}{2} \sup_{y_i=\phi(x_i), \psi(y_i)=x_i} |d(x_1, x_2) - d(y_1, y_2)|$ over all maps $\phi : \mathcal{S}_A \rightarrow \mathcal{S}_B$ and $\psi : \mathcal{S}_B \rightarrow \mathcal{S}_A$, matching point pairs (x_1, x_2) in \mathcal{S}_A with pairs (y_1, y_2) in \mathcal{S}_B . It evaluates — globally and based on an L^∞ type functional — the lack of isometry between two different shapes. Mémoli and Sapiro [9] introduced this concept into the shape analysis community and discussed efficient numerical algorithms based on a robust notion of intrinsic distances $d(\cdot, \cdot)$ on the shapes given by point clouds. Bronstein et al. incorporate the Gromov–Hausdorff distance concept in various classification and modeling approaches in geometry processing [10].

Charpiat et al. [11] discuss shape averaging and shape statistics based on the notion of the Hausdorff distance of sets. They propose to use smooth approximations of the Hausdorff distance based on a comparison of the signed distance functions of shapes. The approach by Eckstein et al. [12] is conceptually related. They consider regularized geometric gradient flow for the warping of surfaces.

There are a variety of approaches which consider shape space as an infinite-dimensional Riemannian manifold. Michor and Mumford [13] gave a corresponding definition exemplified in the case of curves. Younes [14] considered a left invariant Riemannian distance between planar curves. Miller and Younes generalized this concept to the space of images [15]. Klassen and Srivastava [16] proposed a framework for geodesics in the space of arclength parametrized curves and suggested a shooting type algorithm for the computation, whereas Schmidt et al. [4] presented an alternative variational approach.

Dupuis et al. [17] and Miller et al. [18] defined the distance between shapes based on a flow formulation in the embedding space. The underlying motion fields v are globally defined, and as Riemannian metric they considered $\int_{\Omega} Lv \cdot v dx$, where L was chosen as a higher order elliptic operator [19, 14]. This operator ensures sufficient regularity along paths of finite length and thus implies a diffeomorphic property for the flow map ϕ generated via integration of the motion fields v .

Fuchs et al. [2] proposed a Riemannian metric on shape space motivated by linearized elasticity, leading to the same quadratic form (1), which is in their approach evaluated on a displacement field. They used a B-spline parametrization of the shape contour together with a finite element approximation for the displacements on an accompanying triangulation of one of the two objects. Due to the linearization this approach is not rigid body motion invariant, and they do not consider a hierarchical treatment. The explicitly parametrized shapes on a geodesic path share the same topology. A Riemannian metric in the space of surface triangulation in 3D of fixed mesh topology has been investigated by Kilian et al. [20], where an inner product of deformations fields as the underlying metric measures the local distance from a rigid body motion.

1.2 Key contributions

Key contributions of our approach are: • The presented time discretization strictly ensures rigid body motion invariance and a 1-1 mapping property. • The implicit treatment of shapes via level sets allows for topological transitions and enables to compute geodesics in the context of partial occlusion. • Robustness and effectiveness of the algorithm is ensured via a cascadic multi-scale relaxation strategy. • The approach mathematically rigorously links consecutive pairwise shape matching and a flow perspective on a Riemannian shape space. • A formal connection between physics-motivated and geometry-motivated shape spaces is provided, with an intuitive physical interpretation of the framework.

2 Variational time discretization: Preamble and the discrete geodesic model

In this section, we present the time discretization of a geodesic path in shape space, whereas the induced Riemannian distance will be investigated in Sec. 3.

We do not consider a purely geometric notion of shapes as curves in 2D or surfaces in 3D. In fact, motivated by physics, we consider shapes \mathcal{S} as boundaries $\partial\mathcal{O}$ of sufficiently regular, open object domains $\mathcal{O} \subset \mathbb{R}^d$ for $d = 2, 3$.

Discrete path in shape space. Given two shapes $\mathcal{S}_A, \mathcal{S}_B$, we define a discrete path of shapes as a sequence of shapes $\mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_K$ with $\mathcal{S}_0 = \mathcal{S}_A$ and $\mathcal{S}_K = \mathcal{S}_B$. For the time step $\tau = \frac{1}{K}$ the shape \mathcal{S}_k is supposed to be an approximation of $\mathcal{S}(t_k)$ ($t_k = k\tau$), where $\mathcal{S}(t), t \in [0, 1]$, is a continuous path connecting $\mathcal{S}_A = \mathcal{S}(0)$ and $\mathcal{S}_B = \mathcal{S}(1)$, e. g. a geodesic between these two shapes (continuous results will be presented in the next section).

Pairwise deformations between consecutive shapes. Now, we introduce a matching deformation ϕ_k for each pair of consecutive shapes \mathcal{S}_{k-1} and \mathcal{S}_k such that $\phi_k(\mathcal{S}_{k-1}) = \mathcal{S}_k$, and a corresponding deformation energy

$$\mathcal{E}_{\text{deform}}[\phi_k, \mathcal{S}_{k-1}] = \int_{\mathcal{O}_{k-1}} W(\mathcal{D}\phi_k) dx, \quad (2)$$

where W is an energy density (cf. [21]). As in the axiom of elasticity, the energy is assumed to depend only on the local deformation, reflected by the Jacobian $\mathcal{D}\phi$. But different from elasticity, we suppose the material to relax immediately so that the object at the next time step is again in a stress-free configuration. Let us emphasize that the stored energy does not depend on the deformation history as in most plasticity models in engineering. If we postulate as fundamental assumption on the time discretization the invariance of the deformation energy with respect to rigid body motions,¹ i. e.

$$\mathcal{E}_{\text{deform}}[Q \circ \phi_k + b, \mathcal{S}_{k-1}] = \mathcal{E}_{\text{deform}}[\phi_k, \mathcal{S}_{k-1}] \quad (3)$$

for $Q \in SO(d)$ and $b \in \mathbb{R}^d$ (the axiom of frame indifference in continuum mechanics), one deduces that the energy density only depends on the first Cauchy–Green deformation tensor $\mathcal{D}\phi^T \mathcal{D}\phi$ (which geometrically represents the metric measuring the deformed length in the reference configuration), i. e. $W(A) = \bar{W}(A^T A)$ for some \bar{W} . Now, we assume that the deformation is chosen such that the Hessian at the identity coincides with the desired local dissipation rate or metric tensor (1) (cf. Sec. 3). For an isotropic material the energy can be rewritten as a function solely depending on the principal invariants of the Cauchy–Green tensor, namely $I_1 = \text{tr}(\mathcal{D}\phi^T \mathcal{D}\phi)$, controlling the local average change of length, $I_2 = \text{tr}(\text{cof}(\mathcal{D}\phi^T \mathcal{D}\phi))$ ($\text{cof}A := \det A A^{-T}$), reflecting the local average change of area, and $I_3 = \det(\mathcal{D}\phi^T \mathcal{D}\phi)$, which controls the local change of volume. Furthermore let us assume that the energy is a convex function of $\mathcal{D}\phi$, $\text{cof}\mathcal{D}\phi$, and $\det \mathcal{D}\phi$ and that isometries, i. e. deformations with $\mathcal{D}\phi^T(x) \mathcal{D}\phi(x) = \mathbb{1}$, are local minimizers [7] (Fig. 4 provides an example of good local isometry preservation). A template in this class of energy densities is $\bar{W} = \alpha_1 I_1^{\frac{p}{2}} + \alpha_2 I_2^{\frac{q}{2}} + \Gamma(I_3)$ with $p > 0, q \geq 0, \alpha_1, \alpha_2 > 0$, and Γ convex. Indeed, by straightforward computation one verifies that for any dissipation rate (1), there is a nonlinear energy density

¹ Our general framework can be extended to other invariances as well.

of the above type such that the dissipation rate is the corresponding Hessian. In our computations this energy was chosen so that $\frac{\lambda}{\mu} = 3$. A built-in penalization of volume shrinkage, i. e. $\Gamma \xrightarrow{I_3 \rightarrow 0} \infty$, ensures local injectivity [22], and thus the sequence of deformations ϕ_k linking objects \mathcal{O}_{k-1} and \mathcal{O}_k actually represents homeomorphisms (rigorously proved for deformations with finite energy under mild assumptions for sufficiently large p, q , certain growth conditions on Γ , and very soft instead of void material on $\Omega \setminus \mathcal{O}$ with Dirichlet boundary conditions on $\partial\Omega$). Let us remark that self-contact at the boundary is still possible, so that the mapping from $\mathcal{S}_{k-1} = \partial\mathcal{O}_{k-1}$ to $\mathcal{S}_k = \partial\mathcal{O}_k$ does not have to be homeomorphic. By interpreting such self-contact as a closing of the gap between two object edges in the sense that the viscous material flows together, this will indeed allow for certain topological transitions along a discrete path in shape space [7] (cf. Fig. 1 for an example). Based on these mechanical preliminaries we can now define a time discrete geodesic path.

Definition (Discrete Geodesic). A discrete path $\mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_K$ connecting two shapes \mathcal{S}_A and \mathcal{S}_B is a discrete geodesic, if there exists an associated family of deformations ϕ_k with $\phi_k(\mathcal{S}_{k-1}) = \mathcal{S}_k$ which minimizes the total energy $\sum_{k=1}^K \mathcal{E}_{\text{deform}}[\phi_k, \mathcal{S}_{k-1}]$.

Relaxed formulation of the consecutive matching. Computationally, the constraint $\phi_k(\mathcal{S}_{k-1}) = \mathcal{S}_k$ for a 1-1 matching of consecutive shapes is difficult to treat and non-robust (e. g. , not allowing for the handling of noise). Hence, we utilize a relaxed formulation adding a mismatch penalty

$$\mathcal{E}_{\text{match}}[\phi_k, \mathcal{S}_{k-1}, \mathcal{S}_k] = \text{vol}(\mathcal{O}_{k-1} \Delta \phi_k^{-1}(\mathcal{O}_k)), \quad (4)$$

where $A \Delta B = A \setminus B \cup B \setminus A$ defines the symmetric difference between two sets. One might want to further restrict the set of possible shapes \mathcal{S}_k along a discrete geodesic adding an additional surface energy term $\mathcal{E}_{\text{area}}[\mathcal{S}] = \int_{\mathcal{S}} da$. Finally, we end up with the total discrete energy

$$\begin{aligned} \mathcal{E}_{\tau}[(\phi_k, \mathcal{S}_{k-1}, \mathcal{S}_k)_{k=1, \dots, K}] \\ = \sum_{i=1}^K \left(\frac{1}{\tau} \mathcal{E}_{\text{deform}}[\phi_k, \mathcal{S}_{k-1}] + \eta \mathcal{E}_{\text{match}}[\phi_k, \mathcal{S}_{k-1}, \mathcal{S}_k] + \nu \tau \mathcal{E}_{\text{area}}[\mathcal{S}_k] \right), \quad (5) \end{aligned}$$

where η, ν are parameters, and a minimizer of this energy describes a *relaxed discrete geodesic path* between two shapes \mathcal{S}_A and \mathcal{S}_B .



Fig. 4. Discrete geodesic for two different examples from [2] and [23] where the local rate of dissipation is color-coded as . In the right example the local preservation of isometries is clearly visible, whereas in the left example stretching is the major effect.

3 Viscous fluid model in the limit for $\tau \rightarrow 0$

We now investigate the relation of the above-introduced relaxed discrete geodesic paths and a time continuous model for geodesics in shape space.

At first, let us derive from a time discrete sequence of deformations $(\phi_k)_{k=1,\dots,K}$ and shapes $(\mathcal{S}_k)_{k=0,\dots,K}$ a time continuous deformation field ϕ_τ , a corresponding motion field v_τ , and a continuous path \mathcal{S}_τ in shape space:

$$v_\tau(t) := \frac{1}{\tau}(\phi_k - \mathbb{1}), \quad (6)$$

$$\phi_\tau(t) := (\mathbb{1} + (t - t_{k-1})v_\tau) \circ \phi_{k-1} \circ \dots \circ \phi_1, \quad (7)$$


$$\mathcal{S}_\tau(t) := (\mathbb{1} + (t - t_{k-1})v_\tau)(\mathcal{S}_{k-1}), \quad (8)$$

for $t \in [t_{k-1}, t_k]$. If we now let $\tau \rightarrow 0$ and assume that $\mathcal{S}_\tau(t) \rightarrow \mathcal{S}(t)$ for a regular family of shapes $(\mathcal{S}(t))_{0 \leq t \leq 1}$ and that $v_\tau(t) \rightarrow v(t)$ with an induced sufficiently regular flow $(\phi(t))_{0 \leq t \leq 1}$ with $\dot{\phi} = v$, the following limit behavior can be observed: The first term in the global discrete energy representing the deformation energy (2) turns into a time continuous dissipation functional,

$$\mathbf{Diss}[v] = \int_0^1 \int_{\mathcal{O}(t)} \mathbf{C}\epsilon[v] : \epsilon[v] \, dx \, dt, \quad (9)$$

where $\mathbf{C} = 2 \text{Hess } \bar{W}$, $\epsilon[v] = \frac{1}{2}(\mathcal{D}v^T + \mathcal{D}v)$, and $A:B = \text{tr}(A^T B)$ for $A, B \in \mathbb{R}^{d \times d}$. To see this, we observe that $\mathcal{D}\phi_\tau(t)^T \mathcal{D}\phi_\tau(t) = \mathbb{1} + 2(t - t_k)\epsilon[v_\tau(t)] + O(\tau^2)$ for $t \in [t_k, t_{k+1})$, and by second order Taylor expansion of \bar{W} at the identity, $\bar{W}(\mathcal{D}\phi^T \mathcal{D}\phi) = \tau^2 \mathbf{C}\epsilon[v] : \epsilon[v] + O(\tau^3)$. Here, we have used the fact that \bar{W} attains its minimum 0 at the identity. Thus, the resulting Riemannian structure given by the rate of dissipation is indeed associated with the Hessian of our in general nonlinear deformation energy at the identity. For the well-known exemplary metric (1), the length control based on the first invariant I_1 of $\mathcal{D}\phi_\tau$ turns into the infinitesimal length control via $\text{tr}(\epsilon[v]^2)$, and the volume control based on the third invariant I_3 of $\mathcal{D}\phi_\tau$ turns into the control of compression via $\text{tr}(\epsilon[v])^2$ (cf. Fig. 5 for the impact of these two terms on the shapes along a geodesic path). In the limit the term for the mismatch energy (4) converges to an optical flow



Fig. 5. Two geodesic paths between dumb bell shapes varying in the size of the ends. In the left example the ratio λ/μ between the parameters of the dissipation is 0.01 (leading to rather independent compression and expansion of the ends since the associated change of volume implies relatively low dissipation), and 100 in the right example (now mass is actually transported from one end to the other). The underlying texture on the shape domains $\mathcal{O}_0, \dots, \mathcal{O}_{K-1}$ is aligned to the transport direction, and the absolute value of the velocity v is color-coded as .

type energy

$$\mathcal{E}_{\text{OF}} = \int_{\bigcup_{t \in [0,1]}(t, \mathcal{S}(t))} |(1, v(t))^{\text{T}} \cdot n_{\mathcal{S}(t)}| \, da, \quad (10)$$

where $n_{\mathcal{S}(t)}$ denotes the normal on the shape tube $\bigcup_{t \in [0,1]}(t, \mathcal{S}(t))$ in space time and $(1, v(t))$ is the underlying space time motion field (cf. L^1 type optical flow functionals like in [25]). To see this, we have to consider $\text{vol}(\mathcal{O}_{k-1} \Delta \phi_k^{-1}(\mathcal{O}_k))$ as the time discrete mismatch induced by a motion field v_τ which is not consistent with the actual time discrete flow of the shape \mathcal{S}_τ . Indeed, $(1, v_\tau)^{\text{T}} \cdot n_{\mathcal{S}_\tau}$ is the local rate with which $(\mathbb{1} + (t - t_{k-1})v_\tau)(\mathcal{S}_{k-1})$ and the tube of shapes $\bigcup_{t \in [0,1]}(t, \mathcal{S}(t))$ diverge on the time interval $[t_{k-1}, t_k]$.

The third term of the global energy measuring the shape perimeter turns into the time integral over the perimeter. Finally, as $\mathcal{S}_\tau(t) \rightarrow \mathcal{S}(t)$ and $v_\tau(t) \rightarrow v(t)$, the energy converges against

$$\mathcal{E}[v, \mathcal{S}] = \mathbf{Diss}[v] + \eta \mathcal{E}_{\text{OF}}[v, \mathcal{S}] + \nu \int_0^1 \mathcal{E}_{\text{area}}[\mathcal{S}(t)] \, dt. \quad (11)$$

The convergence of the time-discrete energy functional (5) for $\tau \rightarrow 0$ in the sense of Γ -convergence involves further considerations and is not treated here. In the limit $\eta \rightarrow \infty$, the optical flow term will act as a mere penalty which ensures that the family of shapes $\mathcal{S}(t)$ is exactly generated by the flow associated with $v(t)$. In this case and if we set $\nu = 0$, $\mathbf{Diss}[v]$ indeed represents the first fundamental form for the desired Riemannian metric. Thus, the notion of our time discrete geodesics is consistent with this both geometrically and physically sound time continuous geodesic path model in a Riemannian shape space.

4 Regularized level set approximation

To numerically solve the minimization problem for the energy (5), we assume the object domains \mathcal{O} to be represented by zero super level sets $\{x \in \Omega : u(x) > 0\}$ of a scalar function u . Similar representations of shapes have been used for shape matching and warping in [26, 11]. We follow the approximation proposed by Chan and Vese [27] and encode the partition of the domain into object and background in the different energy terms via a regularized Heaviside function $H_\epsilon(u_k)$. As in [27] we consider the function $H_\epsilon(x) := \frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{x}{\epsilon}\right)$, where ϵ is a scale parameter representing the width of the smeared-out shape contour. Hence, the mismatch energy is replaced by the approximation

$$\mathcal{E}_{\text{match}}^\epsilon[\phi_k, u_{k-1}, u_k] = \int_{\Omega} (H_\epsilon(u_k(\phi_k)) - H_\epsilon(u_{k-1}))^2 \, dx, \quad (12)$$

and the area of the k th shape \mathcal{S}_k is replaced by the total variation $\mathcal{E}_{\text{area}}^\epsilon[u_k] = \int_{\Omega} |\nabla H_\epsilon(u_k)| \, dx$ of $H_\epsilon \circ u_k$. With respect to the deformation energy we assume that the whole computational domain is deformed, but with a material which

is several orders of magnitude softer on the complement set $\Omega \setminus \mathcal{O}_k$ than inside \mathcal{O}_k . Hence, the elastic energy (2) is replaced by the energy

$$\mathcal{E}_{\text{deform}}^{\epsilon, \delta}[\phi_k, u_{k-1}] = \int_{\Omega} ((1-\delta)H_{\epsilon}(u_{k-1}) + \delta) W(\mathcal{D}\phi_k) dx, \quad (13)$$

where $\delta = 10^{-4}$ in our implementation. Let us emphasize that in the energy minimization algorithm, the guidance of the initial zero level lines towards the final shapes relies on the nonlocal support of the regularized Heaviside function (cf. [28]). Finally, we end up with the approximation of the total energy,

$$\mathcal{E}_{\tau}^{\epsilon, \delta}[(\phi_k, u_k)_k] = \sum_{k=1}^K \left(\frac{1}{\tau} \mathcal{E}_{\text{deform}}^{\epsilon, \delta}[\phi_k, u_{k-1}] + \eta \mathcal{E}_{\text{match}}^{\epsilon}[\phi_k, u_{k-1}, u_k] + \nu \tau \mathcal{E}_{\text{area}}^{\epsilon}[u_k] \right). \quad (14)$$

In our applications we have chosen $\eta = 200$ and $\nu = 0$ except for Fig. 7, where $\nu = 0.005$. The essential formulas for the variation of the energy can be found in the appendix.

5 Finite element discretization in space

For the spatial discretization of the energy $\mathcal{E}_{\tau}^{\epsilon, \delta}$ in (14) the finite element method has been applied. The level set functions u_k and the different components of the deformations ϕ_k are represented by continuous, piecewise multilinear (trilinear in 3D and bilinear in 2D) finite element functions U_k and Φ_k on a regular grid superimposed on the domain $\Omega = [0, 1]^d$. For the ease of implementation we consider dyadic grid resolutions with $2^L + 1$ vertices in each direction and a grid size $h = 2^{-L}$. In 2D we considered $L = 7, \dots, 10$ and in 3D $L = 7$.

Single level minimization algorithm. For fixed time step τ and fixed spatial grid size h , let us denote by $\mathcal{E}_{\tau, h}^{\epsilon, \delta}[(\Phi_k, U_k)_k]$ the discrete total energy depending on the set of K discrete deformations Φ_1, \dots, Φ_K and $K + 1$ discrete level set functions U_0, \dots, U_K , where U_0 and U_K describe the shapes \mathcal{S}_A and \mathcal{S}_B and are fixed. This is a nonlinear functional both in the discrete deformations Φ_k (due to the concatenation $U_k \circ \Phi_k$ with the discrete level set function U_k and the nonlinear integrand $W(\cdot)$ of the deformation energy $\mathcal{E}_{\text{deform}}^{\epsilon, \delta}$) and in the discrete level set functions U_k (due to the concatenation with the regularized Heaviside function $H_{\epsilon}(\cdot)$). In our energy relaxation algorithm for fixed time step and grid size, we consider a gradient descent approach. We constantly alternate between performing a single gradient descent step for all deformations and one for all level set functions. The step sizes are chosen according to Armijo's rule. This simultaneous relaxation with respect to the whole set of discrete deformations and discrete level set functions, respectively, already outperforms a simple nonlinear Gauss-Seidel type relaxation (cf. Fig. 3). Nevertheless, the capability to identify a globally optimal shortest path between complicated shapes depends on an effective multi-scale relaxation strategy (see below).

Numerical quadrature. Integral evaluations in the energy descent algorithm are

performed by Gaussian quadrature of third order on each grid cell. For various terms we have to evaluate pushforwards $U \circ \Phi$ of a discretized level set function U or a test function under a discretized deformation Φ . In our algorithm, this evaluation is performed exactly at the quadrature points.

Cascadic multi-scale algorithm. The variational problem considered here is highly nonlinear, and for fixed time step size the proposed scheme is expected to have very slow convergence; also it might end up in some nearby local minimum. Here, a multi-level approach (initial optimization on a coarse scale and successive refinement) turns out to be indispensable in order to accelerate convergence and not to be trapped in local minima far from the global minimum. Due to our assumption of a dyadic resolution $2^L + 1$ in each grid direction, we are able to build a hierarchy of grids with $2^l + 1$ nodes in each direction for $l = L, \dots, 0$. Via a simple restriction operation we restrict every finite element function to any of these coarse grid spaces. Starting the optimization on a coarse grid, the results from coarse scales are successively prolonged onto the next grid level for a refinement of the solution [29]. Hence, the construction of a multigrid hierarchy allows to solve coarse scale problems in our multi-scale approach on coarse grids. Since the width ϵ of the diffusive shape representation $H_\epsilon \circ u_k$ should naturally scale with the grid width h , we choose $\epsilon = h$.

On a 3 GHz Pentium 4, still without runtime optimization, 2D computations for $L = 8$ and $K = 8$ require ~ 1 h. Based on a parallelized implementation we observed almost linear scaling.

6 Further results and generalizations

We have computed discrete geodesic paths for 2D and 3D shape contours. The method is both robust and flexible due to the underlying implicit shape description via level sets (cf. Fig. 1 and Fig. 6). Indeed, neither topologically equivalent meshes on the initial shapes are required, nor need the shapes themselves be topologically equivalent. In addition, we can easily restrict the approach to the submanifold of 2D area or 3D volume preserving objects based on a predictor corrector scheme. Fig. 7 shows an example of two different geodesics between the letters X and M , demonstrating the impact of the term $\mathcal{E}_{\text{area}}$ controlling the $d - 1$ dimensional area of the shapes.





Fig. 6. Geodesic path between the cat and the lion, with the local rate of dissipation on the shapes S_0, \dots, S_{K-1} color-coded as  (top) and a transparent slicing plane with texture-coded level lines (bottom).




Fig. 7. Geodesic paths between an X and an M , without a contour length term ($\nu = 0$), allowing for crack formation, (top rows) and with this term damping down cracks and rounding corners (bottom rows). In the bottom rows we additionally enforced area preservation along the geodesic.



Fig. 8. A discrete geodesic connecting different poses of a matchstick man can be computed (from left to right starting with the second), even though part of one arm and one leg of \mathcal{S}_0 (left) are occluded.

In many shape classification applications, one would like to evaluate the distance of a partially occluded shape from a given template shape. As a proof of concept, Fig. 8 depicts a corresponding discrete geodesic path. This requires a minor modification of our model, i.e. solely for $k = 0$ in $\mathcal{E}_{\text{match}}^{\epsilon}$ we insert a smooth function as a mask for \mathcal{S}_0 .




Fig. 9. Left: Pairwise geodesic distances between (also topologically) different letter shapes. Obviously, the B s and X s form clusters, and these two clusters are closer to each other than the significantly distant M . Right: Pairwise geodesic distances between different scanned 3D feet (data courtesy of adidas). Despite being geometrically fairly close, the computed geodesic distance allows to single out the fourth foot as being significantly farther away from the other three, which are almost at equal distance, even though feet 1 and 4 are of equal volume and feet 2 and 3 have 13% less volume.

Furthermore, we evaluated distances between different 2D letters based on the discrete geodesic path length. The resulting clustering is shown in Fig. 9 left. Finally, in Fig. 9 right we studied distances between four different foot level sets converted from 3D scans. Surprisingly, the observed clustering is different from the criterion based on the enclosed volume.

7 Conclusions and future work

We have proposed a novel variational time discretization of geodesics in shape space. The key ingredients are the 1-1 mapping property between consecutive time steps and the rigid body motion invariance. The approach is physically motivated and based on measuring flow-induced dissipation in the interior of shape contours. The proposed formulation allows to weight the effect of the local change of length and volume separately, leading to significantly different geodesic paths. Both physically and with respect to the shape description, geodesic paths can undergo certain topological transitions. A cascading multi-scale relaxation strategy renders the computation robust and effective. Future generalization of the model might deal with the incorporation of prior statistical knowledge and the space of general image morphologies. Furthermore, we would like to rigorously investigate the time discrete to time continuous limit via the concept of Γ convergence.

Acknowledgement. This work has been partially supported by NSF, ARO, ONR, NGA, and DARPA. Benedikt Wirth was supported by BIGS Mathematics and the Hausdorff Center for Mathematics at Bonn University. Furthermore, the authors thank Sergio Conti for valuable hints on the physical model.

Appendix. Here, we give explicit formulas for the variation of the different energy contributions in directions of the unknown functions u_k ($k = 1, \dots, K - 1$) and ϕ_k ($k = 1, \dots, K$), required in the numerical implementation. Let us denote by $\langle \delta_w \mathcal{E}, \vartheta \rangle$ a variation of an energy \mathcal{E} with respect to a parameter function w in a direction ϑ . Using straightforward differentiation, for sufficiently smooth u_k and ϕ_k we obtain

$$\begin{aligned} \langle \delta_{\phi_k} \mathcal{E}_{\text{match}}^\epsilon[\phi_k, u_{k-1}, u_k], \psi \rangle &= 2 \int_{\Omega} (H_\epsilon(u_k \circ \phi_k) - H_\epsilon(u_{k-1})) \delta_\epsilon(u_k \circ \phi_k) \nabla u_k \circ \phi_k \cdot \psi \, dx, \\ \langle \delta_{u_{k-1}} \mathcal{E}_{\text{match}}^\epsilon[\phi_k, u_{k-1}, u_k], \vartheta \rangle &= -2 \int_{\Omega} (H_\epsilon(u_k \circ \phi_k) - H_\epsilon(u_{k-1})) \delta_\epsilon(u_{k-1}) \vartheta \, dx, \\ \langle \delta_{u_k} \mathcal{E}_{\text{match}}^\epsilon[\phi_k, u_{k-1}, u_k], \vartheta \rangle &= 2 \int_{\Omega} (H_\epsilon(u_k \circ \phi_k) - H_\epsilon(u_{k-1})) \delta_\epsilon(u_k \circ \phi_k) \vartheta \circ \phi_k \, dx, \\ \langle \delta_{\phi_k} \mathcal{E}_{\text{deform}}^{\epsilon, \delta}[\phi_k, u_{k-1}], \psi \rangle &= \int_{\Omega} ((1 - \delta) H_\epsilon(u_{k-1}) + \delta) W_{,A}(\mathcal{D}\phi_k) : \mathcal{D}\psi \, dx \end{aligned}$$

for test functions ϑ and test displacements ψ , where $W_{,A}$ denotes the derivative of W with respect to its matrix argument. For the variation of $\mathcal{E}_{\text{area}}^\epsilon[u_k]$ we refer to [27].

References

1. Zhu, L., Yang, Y., Haker, S., Allen, T.: An image morphing technique based on optimal mass preserving mapping. *IEEE T. Image Process.* **16**(6) (2007) 1481–1495
2. Fuchs, M., Jüttler, B., Scherzer, O., Yang, H.: Shape metrics based on elastic deformations. *Journal of Mathematical Imaging and Vision* **to appear** (2009)
3. Miller, M.I., Younes, L.: Group actions, homeomorphisms and matching: a general framework. *International Journal of Computer Vision* **41**(1-2) (2001) 61–84
4. Schmidt, F., Clausen, M., Cremers, D.: Shape matching by variational computation of geodesics on a manifold. In: *Pattern Recognition, LNCS 4174*. (2006) 142–151

5. Zhao, H.K., Chan, T., Merriman, B., Osher, S.: A variational level set approach to multiphase motion. *J. Comp. Phys.* **127** (1996) 179–195
6. Luckhaus, S., Sturzenhecker, T.: Implicit time discretization for the mean curvature flow equation. *Calc. Var.* **3** (1995) 253–271
7. Ciarlet, P.G.: Three-dimensional elasticity. Elsevier Science Publishers B. V. (1988)
8. Kendall, D.G.: Shape manifolds, procrustean metrics, and complex projective spaces. *Bull. London Math. Soc.* **16** (1984) 81–121
9. Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.* **5** (2005) 313–347
10. Bronstein, A., Bronstein, M., Kimmel, R.: Numerical Geometry of Non-Rigid Shapes. Monographs in Computer Science. Springer (2008)
11. 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) (2005) 1–58
12. Eckstein, I., Pons, J., Tong, Y., Kuo, C., Desbrun, M.: Generalized surface flows for mesh processing. In: Eurographics Symposium on Geometry Processing. (2007)
13. Michor, P.W., Mumford, D.: Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc.* **8** (2006) 1–48
14. Younes, L.: Computable elastic distances between shapes. *SIAM J. Appl. Math.* **58** (1998) 565–586
15. Miller, M.I., Younes, L.: Group actions, homeomorphisms and matching: a general framework. Technical report, John Hopkins University, Maryland (1999)
16. Klassen, E., Srivastava, A., Mio, W., Joshi, S.: Analysis of planar shapes using geodesic paths on shape spaces. *IEEE T. Pattern Anal.* **26**(3) (2004) 372–383
17. Dupuis, D., Grenander, U., Miller, M.: Variational problems on flows of diffeomorphisms for image matching. *Quarterly of Applied Mathematics* **56** (1998) 587–600
18. Miller, M., Trounev, A., Younes, L.: On the metrics and Euler-Lagrange equations of computational anatomy. *Ann. Rev. Biomed. Eng.* **4** (2002) 375–405
19. Sundaramoorthi, G., Yezzi, A., Mennucci, A.: Sobolev active contours. *International Journal of Computer Vision.* **73**(3) (2007) 345–366
20. Kilian, M., Mitra, N.J., Pottmann, H.: Geometric modeling in shape space. In: *ACM Transactions on Graphics. Volume 26.* (2007) #64, 1–8
21. Droske, M., Rumpf, M.: Multi scale joint segmentation and registration of image morphology. *IEEE Trans. Pattern Anal.* **29**(12) (2007) 2181–2194
22. Ball, J.: Global invertibility of Sobolev functions and the interpenetration of matter. *Proc. Roy. Soc. Edinburgh* **88A** (1981) 315–328
23. Charpiat, G., Maurel, P., Pons, J.P., Keriven, R., Faugeras, O.: Generalized gradients: Priors on minimization flows. *Int. J. Comput. Vision* **73**(3) (2007) 325–344
24. Kornprobst, P., Deriche, R., Aubert, G.: Image sequence analysis via partial differential equations. *Journal of Mathematical Imaging and Vision* **11** (1999) 5–26
25. Black, M.J., Anandan, P.: A framework for the robust estimation of optical flow. In: Fourth International Conference on Computer Vision, ICCV-93. (1993) 231–236
26. Kapur, T., Yezzi, L., Zöllei, L.: A variational framework for joint segmentation and registration. *IEEE CVPR - MMBIA* (2001) 44–51
27. Chan, T.F., Vese, L.A.: Active contours without edges. *IEEE Transactions on Image Processing* **10**(2) (2001) 266–277
28. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. *International Journal of Computer Vision* **22**(1) (1997) 61–79
29. Bornemann, F., Deuffhard, P.: The cascadic multigrid method for elliptic problems. *Num. Math.* **75**(2) (1996) 135–152