GPU Based Image Geodesics for Optical Coherence Tomography

Benjamin Berkels[1], Michael Buchner[2], Alexander Effland[2], Martin Rumpf[2], Steffen Schmitz-Valckenberg[3]

[1]AICES Graduate School, RWTH Aachen University
[2]Institute for Numerical Simulation, University of Bonn
[3]Department of Ophthalmology, University of Bonn
michael.buchner@uni-bonn.de

Abstract. Within a manifold framework, the interpolation of tomographic image time series is investigated. To this end, the metamorphosis model of a manifold of images is taken into account. Based on a variational time discretization, discrete geodesic paths in this space of images are computed. The space discretization is based on finite elements spanned by tensor product cubic B-splines. An efficient implementation is obtained by utilizing graphics hardware and a proper combination of GPU and CPU computation. First results for time series of optical coherence tomography images of a macular degeneration demonstrate the applicability of this geometric concept.

1 Introduction

This paper deals with the interpolation of images considered as objects in a Riemannian manifold $\mathcal{M}$ of images. In this context, image interpolation can naturally be phrased as computing a geodesic path between the input images. Geodesics on Riemannian manifolds are minimizers of the path energy, which in particular implies that they also minimize the path length and are arclength parametrized. The path energy of a path $u : [0,1] \to \mathcal{M}$, $t \mapsto u(t)$, is given by $\mathcal{E}[u] = \int_0^1 g_u(\dot{u},\dot{u})\,\mathrm{d}t$, where $\dot{u}$ is the time derivative of the curve. In our case, $t \mapsto u(t)$ is a curve of images and the metric $g_u(v,v)$ is a bilinear form measuring the cost of an infinitesimal variation $u+\delta v$ of an image u. To define the metric, we follow the *metamorphosis* model, which was analyzed by Trouvé and Younes [1]. The associated metric is defined as an integral over the image domain $\Omega \subset \mathbb{R}^2$ and reflects

- the cost caused by viscous Newtonian and multipolar dissipation $|Dv|^2 + \gamma|D^2v|^2$ due to friction, where Dv, D^2v are the Jacobian and Hessian of v,
- the cost of the intensity modulation $(\dot{u} + v \cdot \nabla u)^2$ along transport paths described via the so-called material derivative $\dot{u} + v \cdot \nabla u$.

To evaluate the metric, we take into account the flow field v, which causes the minimal cost. Altogether, we obtain for a fixed $\delta > 0$ the path energy

$$\mathcal{E}[u] = \int_0^1 \min_{\text{flow fields } v} \int_\Omega |Dv|^2 + \gamma|D^2v|^2 + \frac{1}{\delta}(\dot{u} + v \cdot \nabla u)^2 \,\mathrm{d}x\,\mathrm{d}t\,. \tag{1}$$

Geodesic paths in the space of images, i.e. minimizers of this path energy, provide smooth interpolations between the input images. In [2], Berkels et al. introduced a variational time discretization of the path energy and proved Γ-convergence to the time continuous path energy under slightly stronger assumptions. This in particular implies the convergence of the minimizers of the discrete path energy to (time continuous) geodesics. Furthermore, they proposed an algorithm to numerically compute time discrete geodesics, for which a spatial discretization of the underlying deformations and image intensities by piecewise affine finite elements is employed.

Here, we improve the robustness and approximation quality taking into account cubic spline spaces. The algorithm is based on an alternating descent scheme, during which multiple images are updated. To speed up these registration subproblems, which are the computational bottleneck, a GPU implementation (cf. [3]) significantly improves the performance of the proposed method.

In this paper, we demonstrate the applicability of this approach to optical coherence tomography (OCT) images in age-related macular degeneration, the most common cause of irreversible visual loss in industrial countries. In this retinal disease, in-vivo imaging by OCT allows to detect and monitor progressive microstructural changes of the outer retina that lead to photoreceptor cell degeneration and thus functional loss (cf. [4]). As shown below, the approach generates suitable interpolations equipped with explicitly computed pointwise motion fields and intensity modulations. Furthermore, we compare a piecewise geodesic path between images recorded annually against geodesic paths spanning a four year period to test the prediction quality of the geodesic interpolation.

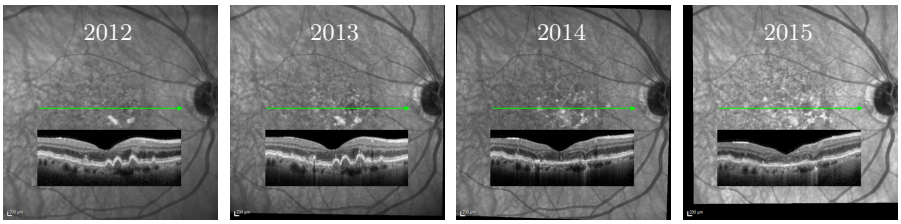

Fig. 1. Ophthalmoscopy (background) and a slice (perpendicular to the green line) from an optical coherence tomography of a human eye (foreground) show an age-related macular degeneration in four consecutive years.

2 Materials and Methods

2.1 Time discrete geodesics in the metamorphosis model

This section summarizes the time discrete metamorphosis model from [2], which builds on a variational time discretization. To approximate the path energy (1),

we consider a time discrete curve $\mathbf{u} = (u_0, \dots, u_K)$ in the space of images (with $u_k : \Omega \to \mathbb{R}$ being a gray valued intensity map) and define a dissimilarity measure

$$\mathcal{W}^D[u, \tilde{u}, \phi] := \int_\Omega |D(\phi - \mathrm{Id})|^2 + \gamma |D^2\phi|^2 + \frac{1}{\delta} |\tilde{u} \circ \phi - u|^2 \, \mathrm{d}x$$

for two images u, $\tilde{u}$, and a deformation $\phi : \Omega \to \Omega$. By minimizing w.r.t. all deformations ϕ with $\phi(x) = x$ for all $x \in \partial\Omega$, we actually solve a very simple (elastic) matching problem between the template image $\tilde{u}$ and the reference image u. Now, we define $\mathcal{W}[u, \tilde{u}] := \min_\phi \mathcal{W}^D[u, \tilde{u}, \phi]$ as the optimal matching cost and use this to introduce a discrete path energy $\mathbf{E}^K[\mathbf{u}] = K \sum_{k=1}^K \mathcal{W}[u_{k-1}, u_k]$ summing over the (minimal) dissimilarity measure of consecutive image pairs. Indeed, [2] shows that the first two terms in $\mathcal{W}^D$ approximate the viscous friction reflected in the first two terms of the path energy (1), whereas the last term approximates the squared material derivative term appearing in (1). Here, the associated time step is $\tau = \frac{1}{K}$.

Now, a *discrete geodesic* connecting two images u_A and u_B is a discrete curve in the space of images that minimizes $\mathbf{E}^K$ over all discrete curves $\mathbf{u}$ with $u_0 = u_A$ and $u_K = u_B$. For a curve $\mathbf{u}$ and deformations $\boldsymbol{\phi} = (\phi_1, \dots, \phi_K)$, we set

$$\mathbf{E}^{K,D}[\mathbf{u}, \boldsymbol{\phi}] = K \sum_{k=1}^K \mathcal{W}^D[u_{k-1}, u_k, \phi_k].$$

Thus, a discrete geodesic from u_A to u_B is obtained by minimizing $\mathbf{E}^{K,D}[\mathbf{u}, \boldsymbol{\phi}]$ with respect to $\mathbf{u}$ and $\boldsymbol{\phi}$ while fixing $u_0 = u_A$ and $u_K = u_B$. Minimizing with respect to $\boldsymbol{\phi}$ for fixed images $\mathbf{u}$ results in K independent image registration problems (registering u_{k-1} to u_k), while minimizing with respect to $\mathbf{u}$ for fixed deformations $\boldsymbol{\phi}$ leads to a linear system of equations for $\mathbf{u}$, i.e. for $k = 1, \dots, K-1$

$$\left(1 + (\det D\phi_k)^{-1} \circ \phi_k^{-1}\right) u_k = u_{k+1} \circ \phi_{k+1} + (u_{k-1} \circ \phi_k^{-1}) \left((\det D\phi_k)^{-1} \circ \phi_k^{-1}\right).$$

2.2 Spatial discretization and computation of geodesics

We consider the unit square $[0, 1]^2$ as our computational domain Ω. To reduce the complexity, we take into account a fine grid $\mathcal{G}_h$ for the image discretization and a coarse grid $\mathcal{G}_H$ for the deformation discretization – both regular and quadrilateral. Grid elements of $\mathcal{G}_h$ and $\mathcal{G}_H$ are denoted by e_h and e_H, respectively. For the space of discrete images, we use the piecewise bilinear finite element space. The space of discrete deformations is the Cartesian product space of (vector valued) cubic splines.

Given $K + 1$ discrete images $\mathbf{U} = (U_0, \dots, U_K)$ with $U_0 = \mathcal{I}_h u_A$ and $U_K = \mathcal{I}_h u_B$ ($\mathcal{I}_h$ is the nodal interpolation operator onto the discrete image space), and K discrete deformations $\boldsymbol{\Phi} = (\Phi_1, \dots, \Phi_K)$, we use a numerical quadrature

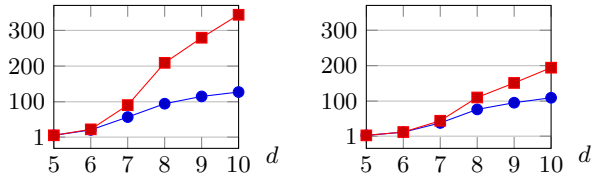

Fig. 2. Left: speedup for assembly of the energy without (blue) and with (red) image warping at image size $(2^d+1)\times(2^d+1)$. Right: Same data shown for gradient assembly.

scheme to compute the discrete path energy

$$\mathbf{E}_{h,H}^{K,D}[\mathbf{U},\mathbf{\Phi}] = K\sum_{k=1}^{K}\Bigg(\sum_{e_H\in\mathcal{G}_H}\sum_{Q=0}^{3}\omega_Q^{e_H}\left(|D\Phi_k-\mathrm{Id}|^2(x_Q^{e_H})+\gamma\,|D^2\Phi_k|^2(x_Q^{e_H})\right) + \frac{1}{\delta}\sum_{e_h\in\mathcal{G}_h}\sum_{q=0}^{8}\omega_q^{e_h}\left(U_k\circ\Phi_k(x_q^{e_h})-U_{k-1}(x_q^{e_h})\right)^2\Bigg).$$

Here, we use Gauss-Legendre quadrature of order 3 in 2D on elements of $\mathcal{G}_H$ to compute the prior with quadrature points $x_Q^{e_H}$ and weights $\omega_Q^{e_H}$. The fidelity term is computed using a tensor product Simpson quadrature on $\mathcal{G}_h$ with quadrature points $x_q^{e_h}$ and weights $\omega_q^{e_h}$. Alternatively, we also use cubic splines for the image intensity discretization on $\mathcal{G}_h$ and adapted the quadrature accordingly.

2.3 GPU accelerated computation of geodesics

In the numerical applications, it became apparent that the time-critical step in the computation of geodesics was the solution of the registration subproblem. The assembly of the discrete energy as well as the gradient exhibit data parallelism as the same integrand has to be evaluated on all elements during the numerical integration. As GPUs are particularly well suited for the computation of the data parallel parts of the program, the energy and gradient assembly are implemented on the graphics card. Elements in the mesh are identified with a thread on the GPU. For the assembly of the discrete energy, a reduction scheme is employed to perform the integration on each element embedded into streams to overlap computation on the GPU. As the assembly of the energy and the gradient involves the evaluation of deformed images, the assembly can be further accelerated by a pre-computing of deformed images via existing GPU image warping tools (cf. Figure 2 for some speed up results).

3 Results

We applied the proposed method to two OCT image sequences, each consisting of four images taken in consecutive years (cf. Figure 1 for one of the sequences) and computed a piecewise geodesic curve interpolating these input images consisting

of $3 \cdot K + 1$ images ($K = 27$). Figure 3 shows the input images in red boxes and the interpolated intermediate time steps u_9 and u_{18} in between. Furthermore, the corresponding velocity fields $\frac{1}{\tau}(\phi_k$–Id) and the intensity modulation given by the discrete material derivative $\frac{1}{\tau}(u_k \circ \phi_k - u_{k-1})$ are shown. Finally, selected images from a direct geodesic connection ($K = 27$) between the 2012 and 2015 images are shown together with the associated velocity fields and intensity modulations. This in particular enables a comparison of this wide span temporal geodesic interpolation with the images recorded for the years in between.

4 Discussion

Based on the tool of discrete geodesics, we are able to compute interpolation paths for a given time series of images. The method also provides information on a probable motion field reflecting the actual deformation process of tissue structures in tomographic images. In our application, a detailed analysis of dynamic disease evolution in age-related macular degeneration may serve for a better understanding of the underlying pathogenetic processes, an identification of prognostic biomarkers for progression and for the evaluation of new therapeutical strategies. In addition, we compared the piecewise geodesic interpolation with key frames for every year with the geodesic interpolation of a four year span. A qualitative comparison of recorded and interpolated images allows a validation of the physical model underlying the image manifold. Indeed, the obtained interpolation properly shows the progressive thinning of the outer retinal layers, while subtle dynamic hyperreflective dots–presumably reflecting migration of retinal pigment epithelium cells–are less accurately detectable by the long range interpolation. For the assessment of velocity fields, an accurate alignment of images at different visits using the hyperreflective outer band of the retinal pigment epithelium as a reference should be incorporated. Motivated by the quality of the obtained interpolation–even over wider time spans–we aim for the computation of image extrapolation to predict the progression of the decease based on in-vivo imaging. This might be of particular importance for future interventional clinical trials that aim to prevent blinding retinal diseases.

References

1. Trouvé A, Younes L. Local Geometry of Deformable Templates. SIAM J Math Anal. 2005;37(1):17–59.
2. Berkels B, Effland A, Rumpf M. Time Discrete Geodesic Paths in the Space of Images. SIAM J Imaging Sci. 2015;8(3):1457–1488.
3. Cheng J, Grossmann M, McKercher T. Professional CUDA C Programming. John Wiley & Sons, Inc.; 2014.
4. Holz FG, Schmitz-Valckenberg S, Fleckenstein M. Recent developments in the treatment of age-related macular degeneration. J Clin Invest. 2014;124(4):1430–1438.

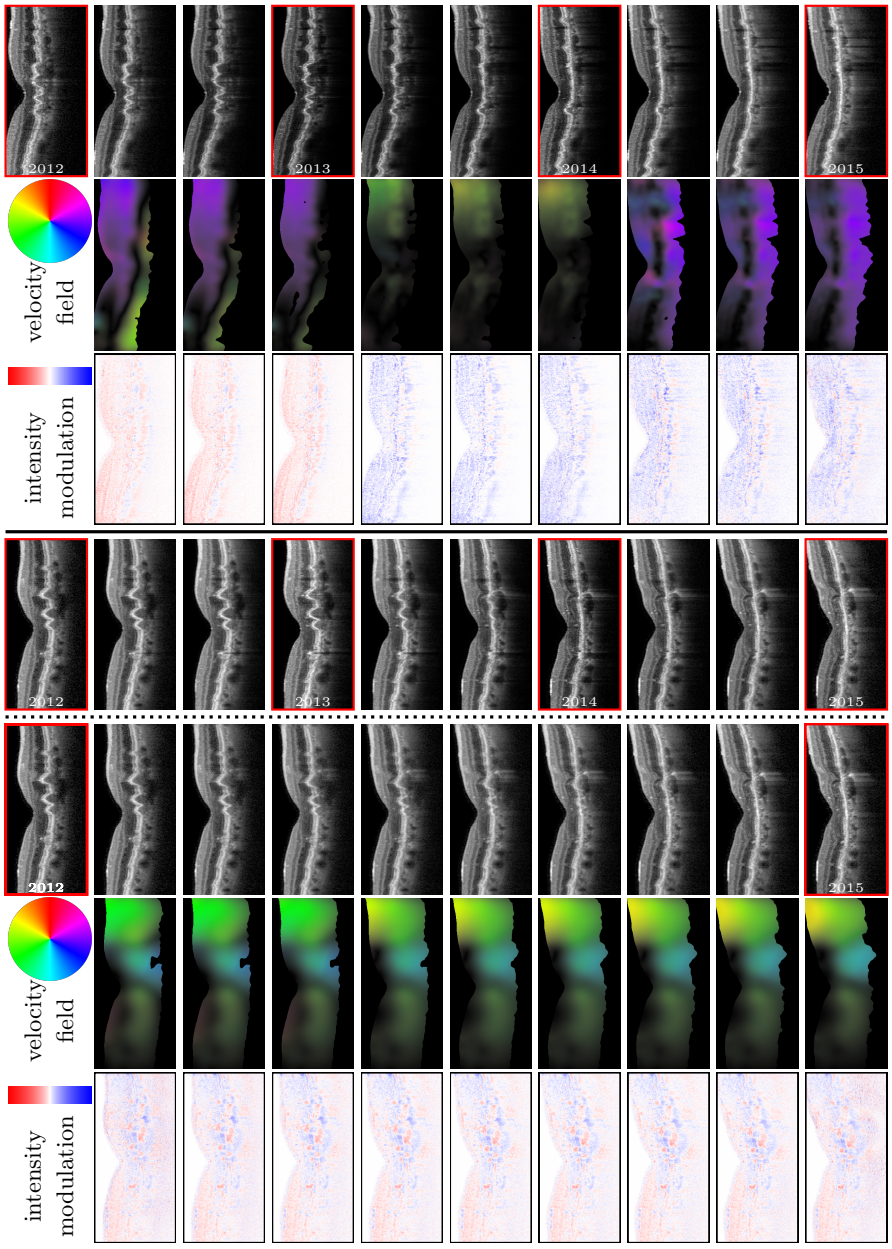

Fig. 3. First/fourth row: piecewise geodesic interpolation for four consecutive years of human eyes for two patients (input data in red boxes). Second/third row: the associated discrete velocity field (the hue refers to the direction and the intensity is proportional to the norm) and the intensity modulation of the first patient. Fifth to seventh row: geodesic interpolation between the years 2012 and 2015 of the second patient and the associated discrete velocity field and intensity modulation, respectively.