

An Optimal Control Approach for the Registration of Image Time-Series

Marc Niethammer

Gabriel L. Hart

Christopher Zach

Abstract—This paper discusses an optimal control approach for the registration of image time-series (growth modeling). It combines and augments work on an optimal control formulation to optical flow with theory from large-displacement diffeomorphic image registration. The unification of the two viewpoints leads to (i) a more efficient computation of the gradient of the optimization problem, (ii) an easier numerical implementation, and (iii) an intuitive interpretation of the adjoint equation underlying the optimization problem. Further, a novel formulation for the unbiased estimation of image correspondences across time is proposed.

I. INTRODUCTION

The most basic and fundamental tools for signal processing are filtering, smoothing, and regression, followed by or combined with an interpretation of the so-processed signals: signal analysis. While these methods are very advanced for example for scalar-valued data, the theory and methodology is much less developed, but of equal importance, for the case of time-varying images.

One of the key components of the analysis of time-varying images is to establish dense point-correspondences between images. The problem is known as image registration in medical imaging or optical flow in computer vision. Its solution is based on a model of allowable deformations and a measure of image similarity. In the most general case, deformations are described non-parametrically, where every image element (pixel/voxel/...) is assigned an individual estimated displacement. This generality makes the estimation problem severely ill-posed. Well-posedness is achieved by imposing regularity constraints on the deformations, penalizing deformations which are perceived as too extreme. If regularization is performed directly on the displacements, one obtains elastic registration. An alternative viewpoint is to cast the image-to-image registration problem into a dynamic formulation, where the source image flows to the target image. Here, regularization is achieved by penalizing irregularities in the flow field, resulting in registration based on a model of fluid flow, and allowing for larger displacements than elastic registration approaches. The fluid flow model can be interpreted from an optimal control perspective, where the optimal control sought is the time-dependent velocity field flowing one image into the other. While a multitude of image-to-image registration methods exist, registration across time is less well explored, but of high importance

in areas such as visual tracking [29], [13], [20], [21], motion analysis [2], [1], [28], [24], cardiac motion estimation [15], [22], [19], and neuroimaging [18], [25].

A straightforward way to extend image-to-image registration approaches to time-series is to composite the registration solutions obtained from image pairs [11]. To obtain and make use of temporal consistency, optical flow approaches have been augmented by incorporating a temporal smoothness constraint [3], [10]. While these methods have shown excellent results at recovering local velocity information [10], they are not designed to estimate deformations over extended periods of time. Since no image is used as a fixed reference, but instead velocity information is only estimated locally, the overall system state (space deformation over time) becomes unobservable, causing estimation drift.

A fixed template may be used to serve as a reference and to counteract the estimation drift. Such a template can either be (i) a single representative image [4], [27], [19], [20], [21], [29], [13], [9] establishing the frame of reference for all other images in the image sequence, or (ii) a spatio-temporal template [25], thus reducing the time-dependent registration problem to registration in spatio-temporal space. The proposed approach falls into the former category, employing a fluid-flow model for the regularization of the time-dependent velocity fields to be estimated.

In this paper, we explore the connections between optimal control and the large-displacement diffeomorphic mapping registration method for growth models (i.e., the estimation of deformations of a template image across time) pioneered by Miller et al. [18]. We demonstrate the relation to the optimal control approach to optical flow by Borzi et al. [9], use it to derive a computationally efficient approach to solve the optimization problem and propose an unbiased formulation, which estimates the time-dependent deformation jointly with a template image representative of the image time-series.

Section II discusses the growth model and its optimality conditions obtained for continuous-time measurements. Section III explains the changes when moving to a more realistic discrete-time measurement model. Discretizations of the resulting equations are explained in Section IV, resulting in a novel hybrid formulation, which combines computational efficiency with numerical accuracy for growth models with sparse discrete-time measurement samples. Section V proposes an unbiased formulation of the growth model. We conclude and discuss future work in Section VII. Details on derivations can be found in the Appendix.

M. Niethammer, G. Hart, and C. Zach are with the Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA. M. Niethammer is also with the Biomedical Research Imaging Center, Medical School, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA {mn, hartg, cmzach}@cs.unc.edu

II. CONTINUOUS GROWTH MODELING

Miller et al. [19] proposed a continuous growth model for the registration of image time sequences minimizing

$$E(v) = \int_{t_0}^T \|v\|_V^2 dt + \frac{1}{\sigma^2} \int_{t_0}^T \|I_T \circ \Phi_{t,t_0} - I^m(t)\|_{L_2}^2 dt, \quad (1)$$

where v is a time-dependent velocity that flows the template image I_T (defined at $t = t_0$) across time to match the measured images $I^m(t)$ as well as possible. Here, $\|v\|_V = \langle v, v \rangle_V$, $\langle f, g \rangle_V := \langle Lf, Lg \rangle = \langle L^\dagger Lf, g \rangle$, L is a differential operator, L^\dagger its adjoint and $\langle \cdot, \cdot \rangle$ denotes the inner product¹.

The template image I_T deforms over time according to the map Φ_{t,t_0} , which in turn is induced by the velocity field v according to

$$\Phi_t + D\Phi v = 0, \quad \Phi(t_0) = id.$$

(Note, the map from time point u to time point t is denoted $\Phi_{t,u} = \Phi_{t,t_0} \circ \Phi_{u,t_0}^{-1}$.)

Instead of directly minimizing Equation 1, we recast the optimization problem into a continuous (dynamically constrained) registration problem for time sequences, which can be interpreted as an optimal control problem with respect to the sought for time-dependent velocity field v , with the first-order system dynamics given by an advection equation for the template image. The optimal control v is the minimizer of

$$E(v) = \int_{t_0}^T \|v\|_V^2 dt + \frac{1}{\sigma^2} \int_{t_0}^T \|I(t) - I^m(t)\|_{L_2}^2 dt, \\ \text{s.t. } \frac{dI}{dt} = I_t + (DI)v = 0, \quad I(t_0) = I_T. \quad (2)$$

The associated optimality conditions are²

$$\begin{aligned} I_t + (DI)v &= 0, \quad I(t_0) = I_T, \\ -\lambda_t - \text{div}(\lambda v) &= -\frac{2}{\sigma^2}(I(t) - I^m(t)), \\ \lambda(T) &= 0, \\ 2L^\dagger Lv + (DI)^T \lambda &= 0, \end{aligned} \quad (3)$$

specifying the state equation, the co-state equation (its adjoint, a scalar conservation law), and the stationarity condition. The adjoint variable measures the accumulated error over time subject to spatial warping (i.e., the overall error is conserved throughout its transport), where the source term of the conservation law injects the error between measurement $I^m(t)$ and estimate $I_T(t)$. The stationarity condition corresponds to the gradient used to find the problem solution iteratively, i.e.,

$$\nabla_v E = 2L^\dagger Lv + (DI)^T \lambda.$$

The adjoint formulation of the growth model (Eq. 3) makes use of the fact that the gradients $\nabla_{v(t)} E$ (for each time point

¹Choosing L appropriately [18], assures that the resulting maps Φ are diffeomorphic, which is a highly desirable property for example for a variety of applications in medical image analysis.

²For space reasons, we do not give the derivation here. The derivation is similar to the one for the continuous-discrete growth model, of Section III, which is detailed in the Appendix.

t) are not independent. Adjoint methods are used for example for meteorological simulations and are the core of variational data assimilation methods and optimal control [8], [12], [16], [7]. Note, that just as in the original solution approach by Beg et al. [5], or for the evolution of Sobolev active contours [26], improvements in convergence by gradient descent in Sobolev space may be obtained for adjoint solution methods [23].

III. CONTINUOUS-DISCRETE GROWTH MODELING

While growth models (such as the one introduced in Section II) are typically formulated with continuous-time measurements, this is frequently not a suitable modeling choice for practical applications.

Notable exceptions, are the works by Khan et al. [14] and Borzi et al. [9], who proposed models with continuous-time dynamics and discrete time-measurements. In this paper, we combine their viewpoints to derive an efficient way of computing the gradient for the growth model optimization problem (which relies on a discrete form of time integrals in the formulation by Khan et al. [14]) combined with a simplified numerical solution method, circumventing the necessity for high-order accurate, nonlinear numerical schemes, as used by Borzi et al. [9].

The energy related to the growth model with discrete-time measurements is

$$E = \int_{t_0}^{t_{M-1}} \|v\|_V^2 dt + \frac{1}{\sigma^2} \sum_{i=0}^{M-1} \|I(t_i) - I^m(t_i)\|_{L_2}^2, \quad (4)$$

subject to the constraint

$$I_t + (DI)v = 0, \quad I(t_0) = I_T.$$

Due to internal measurement constraints, we have jump constraints for the adjoint variable λ associated with the unconstrained problem. This yields the optimality conditions

$$\begin{aligned} I_t + (DI)v &= 0, \\ -\lambda_t - \text{div}(\lambda v) &= 0, \quad \text{piecewise}, \\ 2L^\dagger Lv + (DI)^T \lambda &= 0, \\ I(t_0) &= I_T, \\ \lambda(t_{M-1}(-)) &= \frac{2}{\sigma^2}(I^m(t_{M-1}) - I(t_{M-1})), \\ \lambda(t_i(-)) &= \lambda(t_i(+)) + \frac{2}{\sigma^2}(I^m(t_i) - I(t_i)), \end{aligned} \quad (5)$$

where $(-)$ denotes left-sided and $(+)$ a right-sided values in time, which are distinct at the measurement points.

Since the equations need to be implemented discretely, we regard the continuous formulation as a special case of the continuous-discrete formulation with dense sampling in time. The derivation of the optimality conditions is given in the Appendix.

In contrast to Borzi et al. [9], formulation 4 does not impose explicit temporal smoothness constraints. This allows for simple solutions of the underlying advection equation for I and the scalar conservation law for λ . Further, by utilizing Borzi's optimal control formulation in the context

of large-displacement diffeomorphic growth modeling for images [18], [14], we obtain an alternative formulation of the solution problem, which is easy to derive, and leads directly to more efficient numerical implementations as discussed in Section IV.

IV. EFFICIENT DISCRETIZATION

The continuous and the continuous-discrete approaches operate directly on the image I and the co-state λ . This is sufficient in theory. However, for practical implementations the use of sophisticated numerics is recommended to be able to capture discontinuities in state and co-state, which not only are expected to occur, but are the essential driving force of the overall registration formulation. Borzi et al. [9] use a flux-delimiter to minimize numerical smoothing effects. In contrast, many registration methods in medical imaging do not operate on images directly, but instead update estimations of the map, which aligns image sets. Thus, image discontinuities are no longer problematic, since the object of interest (the map itself) is smooth. Unfortunately, the extensions of these methods to time-series have to-date resulted in numerical schemes with costly gradient computations.

While the continuous formulation of Section II requires the integration of an advection equation for I forward in time and of a scalar conservation law with source term for λ backward in time, the standard map-based formulation of Miller et al. [19] computes the gradient at every time-point as

$$\begin{aligned} \nabla_v E &= 2L^\dagger Lv \\ &- \frac{2}{\sigma^2} D(I_T \circ \Phi_{t,t_0}) \int_t^T (I^m \circ \Phi_{t,u} - I_T \circ \Phi_{t,t_0}) |D\Phi_{t,u}| du, \end{aligned}$$

which requires the evaluation of a time integral for every time point. For the continuous-discrete formulation, Khan et al. [14] obtain

$$\begin{aligned} \nabla_v E &= 2L^\dagger Lv - \frac{2}{\sigma^2} \sum_{i=0}^{M-1} 1_{[t_0, t_i]}(t) |D\Phi_{t,t_i}| \times \\ &(I_T \circ \Phi_{t,t_0} - I^m(t_i) \circ \Phi_{t,t_i})(D(I_T \circ \Phi_{t,t_0}))^T, \end{aligned}$$

where the indicator function $1_{[a,b]}(t) = 1$, $t \in [a, b]$, 0, otherwise. The same equation can be obtained from the optimality conditions for the continuous-discrete approach by pulling all measurements back to the original images (template or measurement) through the maps Φ . It is the equivalent of the time-integral in the continuous setting. While this formulation is numerically benign and accurate, it is computationally costly.

The approach proposed in this paper can be viewed as an alternative solution, which combines the efficiency of the image-based adjoint formulation with the numerical benefits gained from a map-based solution strategy. As can be seen from the optimality conditions 5, gradients at different time-points are not independent. Thus computational complexity can be reduced by simply propagating the adjoint backward in time, where integrations in between measurement points

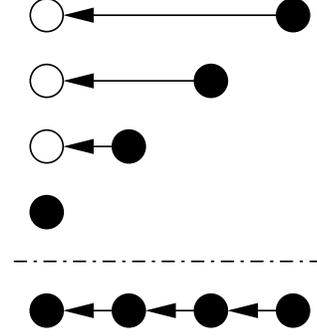


Fig. 1. Illustration of gradient computations. The standard approach (top) is quadratic in the number of image interpolations needed. The new approach (bottom) scales linearly.

are performed using maps to reduce numerical dissipation effects. Of note, this hybrid scheme is only necessary for the adjoint λ , the images $I(t)$ can efficiently be related back to the template image I_T defined at t_0 .

Two types of equations play a central role in the optimal control formulation for registration: The advection equation flowing image intensities along the velocity field v and the adjoint equation flowing the matching error along v backward in time, while conserving its overall mass. Given a current estimate of the velocity field v , the current estimates for the forward map Φ_{t,t_0} and the backward map $\Phi^{-1} = \Phi_{t,t_{M-1}}$ are governed by

$$\begin{aligned} \Phi_t + D\Phi v &= 0, \quad \Phi(t_0) = id, \\ -\Phi_t^{-1} - D\Phi^{-1}v &= 0, \quad \Phi^{-1}(t_{M-1}) = id. \end{aligned}$$

Since the image I simply undergoes a space transformation and λ a mass-preserving transformation, solutions of $I(t)$ and $\lambda(t)$ can be computed through the solutions of $\Phi(t)$ and $\Phi^{-1}(t)$ as

$$I(t) = I_T \circ \Phi(t); \quad \lambda(t) = |D\Phi^{-1}(t)| \lambda(t_{M-1}) \circ \Phi^{-1}(t).$$

These relations allow for the gradient computations for the growth model (3) by evolutions on the maps Φ and Φ^{-1} . The updates of state I and co-state λ can be computed entirely through successive summations and mappings, where the adjoint is updated at every measurement point through a discrete error-injection term. For two time-points only, the approach simplifies to standard large displacement diffeomorphic registration. An overview of the algorithm (with the unbiased extension discussed in Section V) is shown in Figure 2. Figure 1 illustrates the difference between the standard gradient computation (including the time integral) and the one proposed in this paper.

V. UNBIASED FORMULATIONS

Growth models use a template image that is registered across time to derive an image-based growth trajectory. The template is frequently the first image measurement ($I^m(t_0)$). This may be suboptimal from an image registration perspective, because it biases the result obtained with respect

to $I^m(t_0)$. It is also sometimes undesirable: e.g., structural MR brain images only exhibit little contrast in the early stages of neurodevelopment, and show increasingly better contrast with age. One could argue that in such case, a longitudinal analysis should be performed in reverse (from old to young). However, this does not remove the overall bias. Recently, a number of methods have been developed to allow for unbiased image-to-image registration and unbiased atlas building (image-averaging). To the best of our knowledge, unbiased growth models have not been investigated to-date.

Almost none of the approaches for unbiased registration and unbiased atlas-building can be adapted to define an unbiased growth model. Beg and Khan [6] give an overview of symmetric data attachment terms for large deformation image registration. The methods described are designed for image-pairs and the symmetry is achieved in the simplest form by bi-directional image matching, estimating the deformations and penalizing deformation differences (between the forward and the backward map). Another approach subdivides the image-to-image registration problem into two time intervals of equal length, matching images in the middle. Due to their innate dependency on image pairs, the proposed approaches do not extend to the case of more than two images.

We unbiased the estimation problem by regarding the template image I_T as a free parameter and estimate it jointly with the velocity field v . For the basic two image case, this results in penalized matching of both images and is thus symmetric³.

While the image-based formulation of the growth model is advantageous to derive efficient schemes to compute the gradient with respect to v , the map-based formulation

$$E = \left(\sum_{i=0}^{M-2} \int_{t_i(+)}^{t_{i+1}(-)} \|v\|_V^2 + \langle \lambda, \Phi_t + (D\Phi)v \rangle dt \right) + \frac{1}{\sigma^2} \sum_{i=0}^{M-1} \|I_T \circ \Phi(t_i) - I^m(t_i)\|_{L^2}^2,$$

allows for the easy derivation of the best template image (from the viewpoint of optimal control with free initial and free final state).

Computing the variation, we note that the gradient of the energy with respect to the template image is

$$\nabla_{I_T} E = \frac{2}{\sigma^2} \sum_{i=0}^{M-1} |D\Phi_{t_0, t_i}| (I_T - I^m(t_i) \circ \Phi_{t_0, t_i}),$$

whose steady state condition is

$$\sum_{i=0}^{M-1} |D\Phi_{t_0, t_i}| I_T = \sum_{i=0}^{M-1} |D\Phi_{t_0, t_i}| I^m(t_i) \circ \Phi_{t_0, t_i}$$

³Of note, while this approach is unbiased in theory, for practical implementations, with coarse spatial discretizations, the domain of definition for the template image may matter: If large size differences are expected across time, it is advisable to define the template image at a time point of large spatial extent.

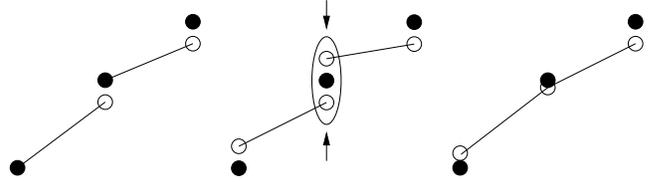


Fig. 3. Different ways of computing growth. Left: chaining registrations together, keeping one boundary fixed. Middle: chaining unbiased registrations together. Right: Complete growth model. Standard registration (left) matches one image exactly and penalizes discrepancies to the other. Unbiased image-to-image registration removes this bias between images (middle), unbiased growth modeling removes it for a series of images.

and thus

$$I_T = \frac{I^m(t_0) + \sum_{i=1}^{M-1} |D\Phi_{t_0, t_i}| I^m(t_i) \circ \Phi_{t_0, t_i}}{1 + \sum_{i=1}^{M-1} |D\Phi_{t_0, t_i}|},$$

which is a weighted image average, with weights proportional to image size changes with respect to their original domain of definition (i.e., if an image is locally compressed to be aligned in template space it will carry greater weight locally than an image that gets locally expanded). The overall unbiased algorithm for registration of time-series is summarized in Figure 2. Figure 3 illustrates the difference between registration schemes.

VI. RESULTS AND DISCUSSION

To demonstrate the behavior of the hybrid implementation scheme for the adjoint equation, we investigated its behavior for the solution of a simple advection equation $u_t + u_x = 0$ and compared it to results obtained from a purely map-based solution (i.e., the map is evolved with forward Euler time discretization and upwinding for the spatial derivatives), from the direct solution of the equation through a total variation diminishing, fifth order weighted-essentially-non-oscillatory (WENO) implementation, as well as from a direct solution obtained from first-order upwinding and forward Euler time discretization. Results are shown in Figure 4. As expected, direct solutions blur the final result, with the WENO scheme doing qualitatively better than the first-order upwind approach, which leads to strong signal blurring. The map solution result is virtually indistinguishable from the sought for final profile. The proposed hybrid solution yields results of high quality for reasonable numbers of intermediate steps. Of note, the 100 steps of the example problem would correspond to 100 discrete-time measurements. Extremely large numbers of steps (1000 in this case) yield mis-estimations. This case corresponds to an application of the hybrid scheme to effectively continuous-time measurements, where a fully map-based scheme is recommended.

Figures 5 and 7 show the results of time-series registration for a set of three distinct synthetic shapes for strong and weak regularizations respectively. Shapes are generally matched well, though changes regarded as too extreme are not recovered, as can be seen from the warping of image 1 to 2 in Figure 5. Forward and inverse maps are smooth (diffeomorphic), but significantly more extreme for weak

Algorithm 1: Unbiased hybrid continuous-discrete growth-model**Data:** $\{I^m(t_i)\}$, $i = 0(1)M - 1$, t_0 , t_{M-1} , σ , L **Result:** v Initialization: $v = 0$, $\lambda = 0$, $I_T = I(t_0)$;**repeat**Flow map forward: $\Phi_t + (D\Phi)v = 0$, $\Phi(t_0) = id$;Flow map backward: $-\Phi_t^{-1} - (D\Phi^{-1})v = 0$, $\Phi^{-1}(t_{M-1}) = id$; $\lambda(t_{M-1}(+)) = 0$; $i = M - 1$;**while** $i > 0$ **do** $\lambda(t_i(-)) = \lambda(t_i(+)) + \frac{2}{\sigma^2}(I^m(t_i) - I_T \circ \Phi(t_i))$;Flow temporary map, T^{-1} , backward, $t \in [t_{i-1}(+), t_i(-)]$: $-T^{-1} - (DT^{-1})v = 0$, $T^{-1}(t_i(-)) = id$;Compute adjoint updates, $t \in [t_{i-1}(+), t_i(-)]$: $\lambda(t) = \lambda(t_i(-)) \circ T^{-1}(t) |DT^{-1}(t)|$; $i \leftarrow i - 1$;**end**Update velocity fields: $v(t)_\theta = -(2v(t) + (L^\dagger L)^{-1}(D\Phi(t))^T \lambda(t))$;Update template image: $I_T = \frac{I^m(t_0) + \sum_{i=1}^{M-1} |D\Phi_{t_0, t_i}| I^m(t_i) \circ \Phi_{t_0, t_i}}{1 + \sum_{i=1}^{M-1} |D\Phi_{t_0, t_i}|}$;**until** convergence ;

Fig. 2. Algorithmic description of the hybrid continuous-discrete growth-model

regularizations, as shown in Figure 7 (which trades off a better image match with regularity of the recovered maps). Figures 6 and 8 show the template estimations throughout the registration process for strong and weak regularizations respectively. They both converge after only one iteration step. The estimated templates constitute a compromise between mean shape and average intensity across all shapes. Since object area factors into the computation of the unbiased template, the outer part of the template images is estimated as dark gray (consistent with the area covered by dark colors for images 2 and 3). The shape interior approximates a circular shape for strong regularization, which constitutes a good compromise of the interior shapes of the three images, while avoiding structural blurring. Weak regularization yields an interior shape dominated by the interior shape of image 2 which has the strongest contrast.

Figure 9 shows unbiased time-series registration results for a set of real magnetic resonance axial brain image slices of subjects between 38 and 81 years of age. Images were obtained through the OASIS project [17]. We used the skull-stripped, normalized brain volumes provided. Subjects are distinct and are all female. They were selected for presentational purposes. The most prominent difference between the five selected subjects is the increase of the ventricle sizes with age. Cortical structure is distinct. Registration successfully follows the expansion of the ventricles. The evolving template image, shown in Figure 10, demonstrates overall preservation of structure, while smoothing inconsistent cortical areas.

VII. CONCLUSIONS AND FUTURE WORK

A. Conclusions

We presented a novel optimal control perspective on image-based growth modeling, resulting in a computationally

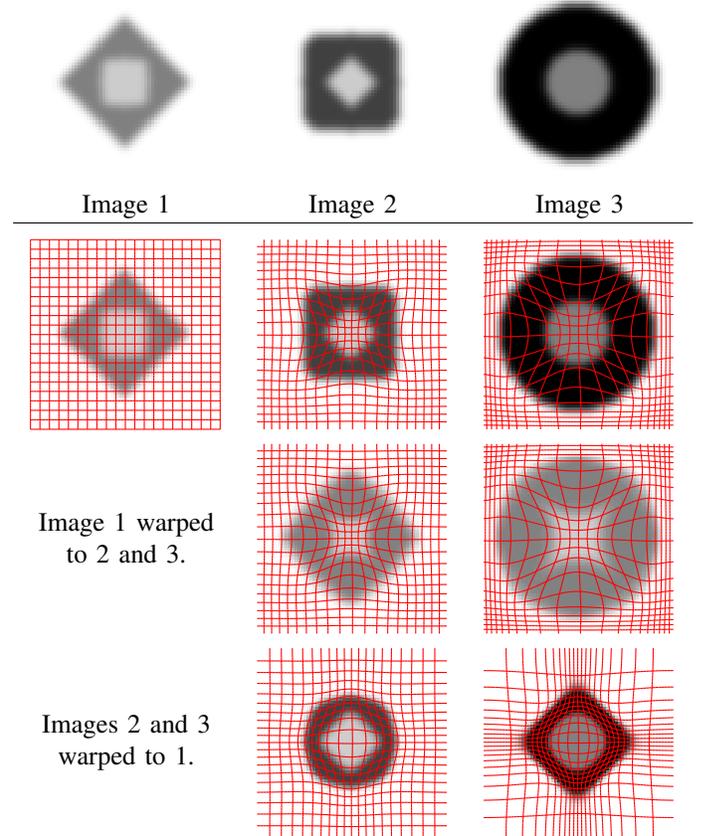


Fig. 5. Unbiased growth-modeling for a synthetic example with strong regularization. Image sequence (top row), computed deformation fields (second row), and forward and inverse image warpings (bottom two rows). Ten temporal discretization steps between images. $L = -0.3\nabla^2 + 1$, $\sigma = 0.1$.

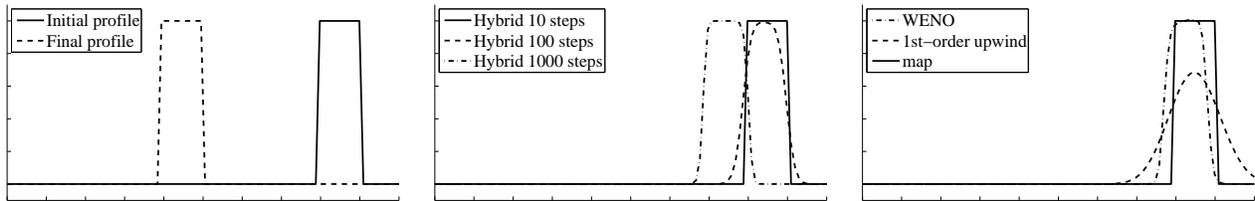


Fig. 4. Comparison of numerical schemes for a simple advection equation: $u_t + u_x = 0$. The domain is of length 1 discretized into 100 elements. Evolution time $t = 0.4$. Evolving the map Φ with upwinding and Euler forward produces a virtually indistinguishable solution (right) from the true final profile (left). A 5-th order WENO scheme captures the shape well, but introduces some signal distortion, whereas first-order upwinding leads to severe blurring of the initial profile (right). The hybrid scheme (middle) produces high-quality results for reasonable numbers of intermediate interpolation steps (middle).

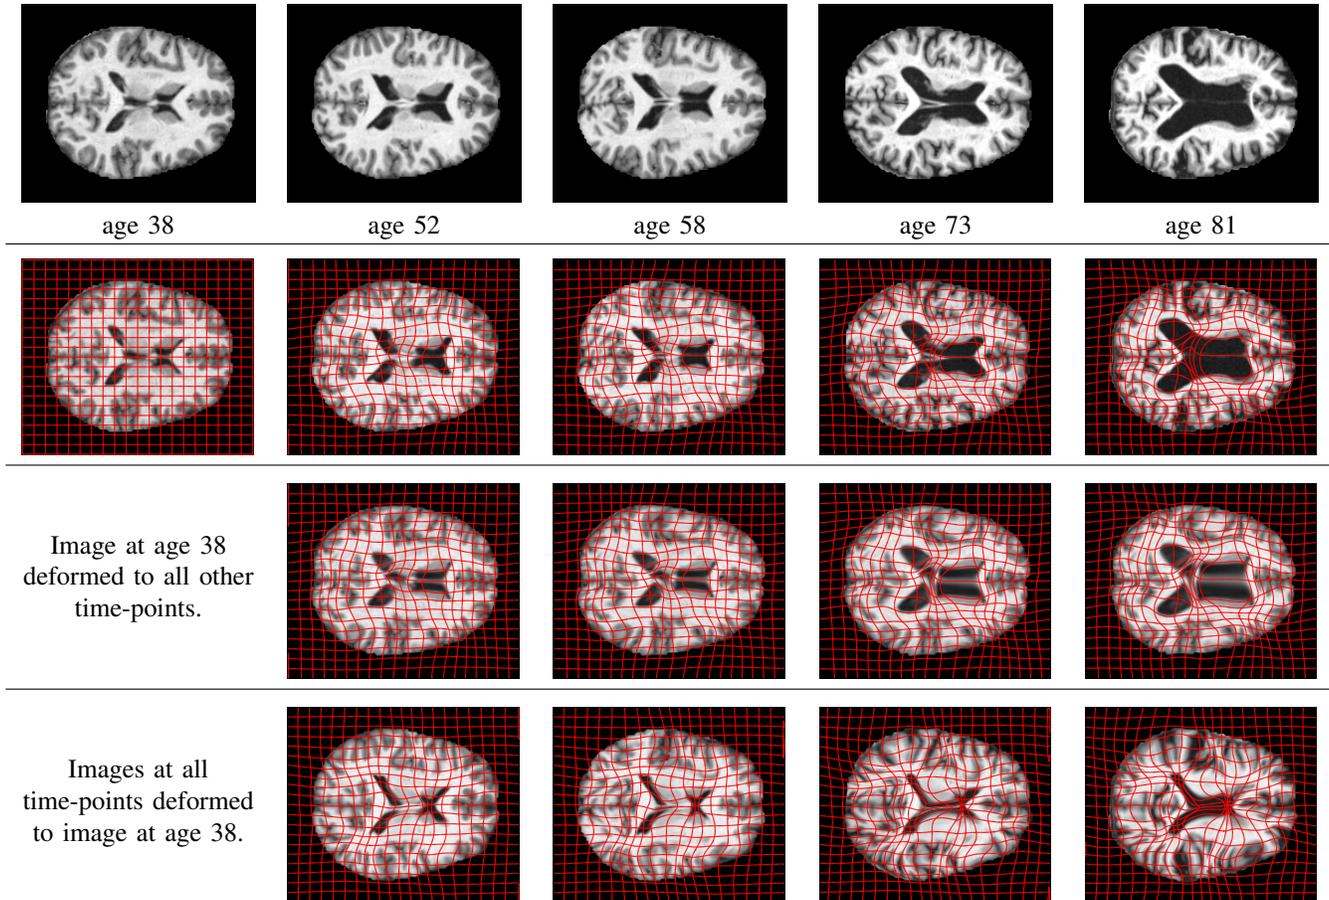


Fig. 9. Unbiased growth-modeling for a set of brain images of different ages. Image sequence (top row), computed deformation fields (second row), and forward and inverse image warpings (bottom two rows). One temporal discretization step per year. $L = -0.02\nabla^2 + 1$, $\sigma = 0.5$

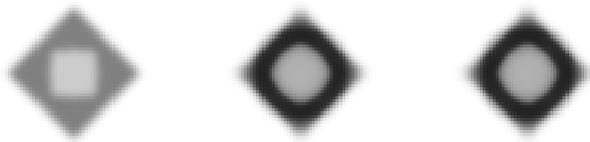


Fig. 6. Synthetic template evolution with strong regularization. Convergence is achieved after one iteration step. The result shows a compromise between the shapes used to construct the template image, both in terms of shape as well as in terms of intensity distributions. Overall structure is maintained. $L = -0.3\nabla^2 + 1$, $\sigma = 0.1$.

efficient algorithm for sparse discrete-time measurements. The optimality conditions obtained highlight the properties of the optimization process, where image mismatches are propagated backward in time through a conservation law preserving the overall error measure subject to space deformations. The derivations allow for three different implementations (two of them new) for growth modeling: (i) A direct implementation based on the image and its adjoint; the approach then does no longer require an explicit coordinate system and can in principle operate on any manifold (but may require sophisticated numerics). (ii) A hybrid version, which

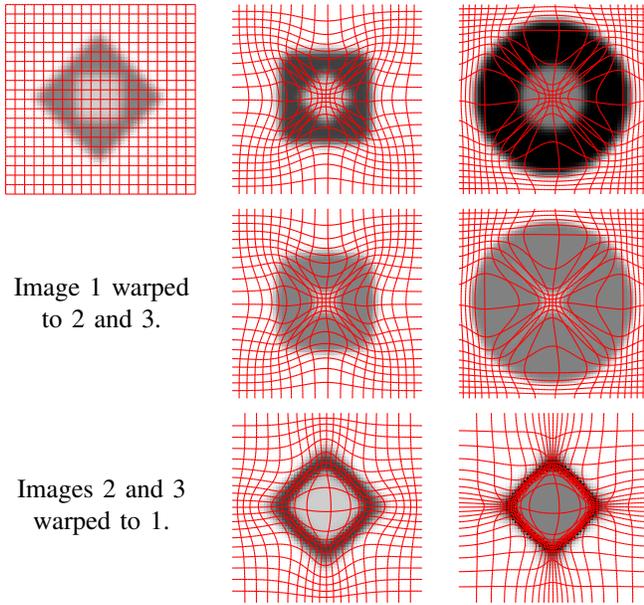


Fig. 7. Unbiased growth-modeling for the synthetic example with weak regularization. Image sequence (top row), computed deformation fields (second row), and forward and inverse image warpings (bottom two rows). Ten temporal discretization steps between images. $L = -0.075\nabla^2 + 1$, $\sigma = 0.1$.

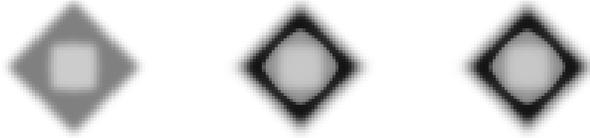


Fig. 8. Template evolution for the synthetic example, with weak regularization. Convergence is achieved after one iteration step. The result shows a compromise between the shapes used to construct the template image, both in terms of shape as well as in terms of intensity distributions. Overall structure is maintained. $L = -0.075\nabla^2 + 1$, $\sigma = 0.1$.

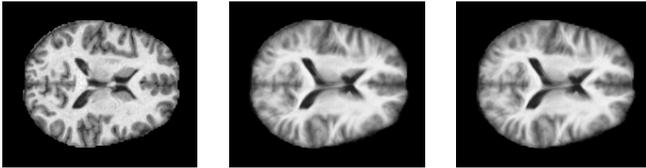


Fig. 10. Template evolution for the set of brain images. Convergence is achieved after one iteration step, showing a compromise between the shapes used to construct the template image, both in terms of shape as well as in terms of intensity distributions. Overall structure is maintained.

is designed for the continuous-discrete case as encountered frequently in practice. The approach combines the efficient optimality conditions obtained through the optimal control perspective, with a numerical scheme which yields good accuracy, while keeping the complexity of the gradient computations low. (iii) We recovered the continuous-discrete approach by Khan et al. [14] for a completely map-based implementation. However, we provide an alternative perspective

through an alternate derivation, giving additional insight into the behavior and properties of the optimization. Further, we proposed, to the best of our knowledge, the first unbiased formulation for time-series registration (the growth model), obtained through joint estimation of a free initial state (the template image) and the optimal control (the time-dependent velocity field).

B. Future Work

We will explore memory-efficient implementations, initial-value formulations, as well as generalizations of the approach to higher order dynamics. The latter would allow for the estimation of smooth image trajectories, avoiding the currently present velocity jumps at measurement instants. This is of particular interest, for processes varying smoothly in time, which can only be sampled coarsely in time.

VIII. ACKNOWLEDGMENTS

We gratefully acknowledge support by NSF (ECCS-0925875) and the OASIS project (Open Access Series of Imaging Studies; <http://www.oasis-brains.org/>), whose freely available data we used for our brain registration experiments [17].

APPENDIX I

VELOCITY GRADIENT AND OPTIMALITY CONDITIONS

Due to internal measurement constraints we will have jump constraints for the adjoint variable λ associated with the unconstrained problem

$$E(v, I, \lambda, \gamma) = \int_{t_0}^{t_{M-1}} \|v\|_V^2 + \langle \lambda, I_t + (DI)v \rangle dt + \langle \gamma, I(t_0) - I_T \rangle + \frac{1}{\sigma^2} \sum_{i=0}^{M-1} \|I(t_i) - I^m(t_i)\|_{L_2}^2.$$

Due to the jump constraints, the resulting dynamics equations will only be valid piecewise, in-between measurement intervals. We thus rewrite the energy as

$$E = \left(\sum_{i=0}^{M-2} \int_{t_i(+)}^{t_{i+1}(-)} \|v\|_V^2 + \langle \lambda, I_t + (DI)v \rangle dt \right) + \langle \gamma, I(t_0) - I_T \rangle + \frac{1}{\sigma^2} \sum_{i=0}^{M-1} \|I(t_i) - I^m(t_i)\|_{L_2}^2.$$

The variation is

$$\delta E = \left(\sum_{i=0}^{M-2} \int_{t_i(+)}^{t_{i+1}(-)} 2\langle v, dv \rangle_V + \langle d\lambda, I_t + (DI)v \rangle + \langle \lambda, dI_t + (DdI)v + (DI)dv \rangle dt \right) + \langle d\gamma, I(t_0) - I_T \rangle + \langle \gamma, dI(t_0) \rangle + \frac{2}{\sigma^2} \sum_{i=0}^{M-1} \langle I(t_i) - I^m(t_i), dI(t_i) \rangle,$$

which can be written (after rearranging terms, integration by parts, and noting that $dI(t_i(-)) = dI(t_i(+))$ because of

continuity of the state variable I) as

$$\begin{aligned} \delta E = & \left(\sum_{i=0}^{M-2} \int_{t_i(+)}^{t_{i+1}(-)} \langle 2L^\dagger Lv + (DI)^T \lambda, dv \rangle \right. \\ & \left. + \langle d\lambda, I_t + (DI)v \rangle + \langle -\lambda_t - \text{div}(\lambda v), dI \rangle dt \right) \\ & + \sum_{i=0}^{M-2} \left(\langle \lambda(t_{i+1}(-)), dI(t_{i+1}) \rangle - \langle \lambda(t_i(+)), dI(t_i(+)) \rangle \right) \\ & + \frac{2}{\sigma^2} \sum_{i=0}^{M-1} \langle I(t_i) - I^m(t_i), dI(t_i) \rangle \\ & + \langle d\gamma, I(t_0) - I_T \rangle + \langle \gamma, dI(t_0) \rangle \end{aligned}$$

which simplifies (with $dI(t_0) = dI_T = 0$) to

$$\begin{aligned} \delta E = & \left(\sum_{i=0}^{M-2} \int_{t_i(+)}^{t_{i+1}(-)} \langle 2L^\dagger Lv + (DI)^T \lambda, dv \rangle \right. \\ & \left. + \langle d\lambda, I_t + (DI)v \rangle + \langle -\lambda_t - \text{div}(\lambda v), dI \rangle dt \right) \\ & + \langle d\gamma, I(t_0) - I_T \rangle + \langle \lambda(t_{M-1}(-)) \\ & + \frac{2}{\sigma^2} (I(t_{M-1}) - I^m(t_{M-1})), dI(t_{M-1}) \rangle + \sum_{i=1}^{M-2} \langle -\lambda(t_i(+)) \\ & + \lambda(t_i(-)) + \frac{2}{\sigma^2} (I(t_i) - I^m(t_i)), dI(t_i) \rangle. \end{aligned}$$

The optimality conditions follow.

APPENDIX II

TEMPLATE-IMAGE GRADIENT

Computing the variation with respect to the template image I_T is simplified, by using a map-based formulation. For v and Φ fixed. The I_T -dependent energy term is

$$E(I_T) = \frac{1}{\sigma^2} \sum_{i=0}^{M-1} \|I_T \circ \Phi_{t_i, t_0} - I^m(t_i)\|_{L_2}^2.$$

Its variation is

$$\delta E(I_T; dI_T) = \frac{2}{\sigma^2} \sum_{i=0}^{M-1} \langle I_T \circ \Phi_{t_i, t_0} - I^m(t_i), dI_T \circ \Phi_{t_i, t_0} \rangle,$$

which (after change of variables) becomes

$$\delta E = \frac{2}{\sigma^2} \sum_{i=0}^{M-1} \langle D\Phi_{t_0, t_i} (I_T - I^m(t_i) \circ \Phi_{t_0, t_i}), dI_T \rangle$$

and yields the gradient.

REFERENCES

- [1] J. K. Aggarwal and Q. Cai. Human motion analysis: A review. *Computer Vision and Image Understanding*, 73(3):428–440, 1999.
- [2] J. K. Aggarwal, Q. Cai, W. Liao, and B. Sabata. Nonrigid motion analysis: Articulated and elastic motion. *Computer Vision and Image Understanding*, 70(2):142–156, 1998.
- [3] L. Alvarez, C. A. Castano, M. Garcia, K. Krissian, L. Mazorra, A. Salgado, and J. Sanchez. Symmetric optical flow. In *Proceedings of EUROCAST*, pages 676–683, 2007.
- [4] K. A. Bartels, A. C. Bovik, and C. E. Griffin. Spatio-temporal tracking of material shape change via multi-dimensional splines. In *Proceedings of the IEEE Workshop on Biomedical Image Analysis*, pages 110–116, 1994.
- [5] M. Beg, M. Miller, A. Trounev, and L. Younes. Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157, 2005.
- [6] M. F. Beg and A. Khan. Symmetric data attachment terms for large deformation image registration. *IEEE Transactions on Medical Imaging*, 26(9):1179–1189, 2007.
- [7] L. Bengtsson, M. Ghil, and E. Källen, editors. *Dynamic Meteorology: Data Assimilation Methods*. Springer, 1981.
- [8] A. F. Bennett. *Inverse Methods in Physical Oceanography*. Cambridge University Press, 1992.
- [9] A. Borzi, K. Ito, and K. Kunisch. Optimal control formulation for determining optical flow. *SIAM Journal on Scientific Computing*, 24(3):818–847, 2002.
- [10] A. Bruhn, J. Weickert, and C. Schnörr. Lucas/Kanade meets Horn/Schunck: Combining local and global optical flow methods. *International Journal of Computer Vision*, 61(3):211–231, 2005.
- [11] I. Csapo, C. M. Holland, and C. R. G. Guttman. Image registration framework for large-scale longitudinal MRI data sets: strategy and validation. *Magnetic Resonance Imaging*, 25:889–893, 2007.
- [12] G. Evensen. *Data Assimilation: The Ensemble Kalman Filter*. Springer, 2007.
- [13] J. Jackson, A. Yezzi, and S. Soatto. Tracking deformable moving objects under severe occlusions. In *Proceedings of the Conference on Decision and Control*, volume 3, pages 2990–2995. IEEE, 2004.
- [14] A. R. Khan and M. F. Beg. Representation of time-varying shapes in the large deformation diffeomorphic framework. In *Proceedings of the International Symposium of Biomedical Imaging (ISBI)*, pages 1521–1524. IEEE, 2008.
- [15] M. J. Ledesma-Carbayo, J. Kybic, M. Desco, A. Santos, M. Sühling, P. Hunziker, and M. Unser. Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation. *IEEE Transactions on Medical Imaging*, 24(9):1113–1126, 2005.
- [16] J. L. Lions. *Optimal Control Systems Governed by Partial Differential Equations*. Springer, 1970.
- [17] D. S. Marcus, T. H. Wand, J. Parker, J. G. Csernansky, J. C. Morris, and R. L. Buckner. Open access series of imaging studies (OASIS): Cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *Journal of Cognitive Neuroscience*, 19:1498–1507, 2007.
- [18] M. Miller. Computational anatomy: shape, growth, and atrophy comparison via diffeomorphisms. *Neuroimage*, 23:S19–S33, 2004.
- [19] M. Miller, A. Trounev, and L. Younes. On the metrics and Euler-Lagrange equations of computational anatomy. *Annual Reviews in Biomedical Engineering*, 4(1):375–405, 2002.
- [20] N. Papadakis, T. Corpetti, and E. Mémin. Dynamically consistent optical flow estimation. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 1–7, 2007.
- [21] N. Papadakis and E. Mémin. A variational technique for time consistent tracking of curves and motion. *Journal of Mathematical Imaging and Vision*, 31(1):81–103, 2008.
- [22] D. Perperidis, R. H. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *Medical Image Analysis*, 9:441–456, 2005.
- [23] B. Protas, T. R. Bewley, and G. Hagen. A computational framework for the regularization of adjoint analysis in multiscale PDE systems. *Journal of Computational Physics*, 195:49–89, 2004.
- [24] W. Ren, S. Singh, M. Singh, and Y. S. Zhu. State-of-the-art on spatio-temporal information-based video retrieval. *Pattern Recognition*, 42:267–282, 2009.
- [25] D. Shen and C. Davatzikos. Measuring temporal morphological changes robustly in brain MR images via 4-dimensional template warping. *Neuroimage*, 21(4):1508–1517, 2004.
- [26] G. Sundaramoorthi, A. Yezzi, and A. C. Mennucci. Coarse-to-fine segmentation and tracking using Sobolev active contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(5):851–864, 2008.
- [27] R. Szeliski and J. Coughlan. Spline-based image registration. *International Journal of Computer Vision*, 22(3):199–218, 1997.
- [28] A. Veeraraghavan, A. K. Roy-Chowdhury, and R. Chellappa. Matching shape sequences in video with applications in human movement analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(12):1896–1909, 2005.
- [29] A. J. Yezzi and S. Soatto. Deformation: Deforming motion, shape average and the joint registration and approximation of structures in images. *International Journal of Computer Vision*, 53(2):153–167, 2003.