

Computing Diffeomorphic Paths for Large Motion Interpolation*

Dohyung Seo
 Department of Electrical and
 Computer Engineering
 University of Florida
 dhseo@ufl.edu

Jeffrey Ho Baba C. Vemuri
 Department of Computer and
 Information Sciences and Engineering
 University of Florida
 {jho, vemuri}@cise.ufl.edu

Abstract

In this paper, we introduce a novel framework for computing a path of diffeomorphisms between a pair of input diffeomorphisms. Direct computation of a geodesic path on the space of diffeomorphisms $\text{Diff}(\Omega)$ is difficult, and it can be attributed mainly to the infinite dimensionality of $\text{Diff}(\Omega)$. Our proposed framework, to some degree, bypasses this difficulty using the quotient map of $\text{Diff}(\Omega)$ to the quotient space $\text{Diff}(M)/\text{Diff}(M)_\mu$ obtained by quotienting out the subgroup of volume-preserving diffeomorphisms $\text{Diff}(M)_\mu$. This quotient space was recently identified as the unit sphere in a Hilbert space in mathematics literature, a space with well-known geometric properties. Our framework leverages this recent result by computing the diffeomorphic path in two stages. First, we project the given diffeomorphism pair onto this sphere and then compute the geodesic path between these projected points. Second, we lift the geodesic on the sphere back to the space of diffeomorphisms, by solving a quadratic programming problem with bilinear constraints using the augmented Lagrangian technique with penalty terms. In this way, we can estimate the path of diffeomorphisms, first, staying in the space of diffeomorphisms, and second, preserving shapes/volumes in the deformed images along the path as much as possible. We have applied our framework to interpolate intermediate frames of frame-sub-sampled video sequences. In the reported experiments, our approach compares favorably with the popular Large Deformation Diffeomorphic Metric Mapping framework (LDDMM).

1. Introduction

Image and shape matching using 2D diffeomorphisms (diffeomorphic matching) is a popular technique in many computer vision and medical imaging applications such as

atlas construction [1] and shape analysis with a long and venerable history [1, 2, 3, 4, 5, 6]. Recall that a diffeomorphism is a smooth bijective mapping between two (image) domains, and many important types of transformations and deformations used in computer vision can be taken as 2D/3D diffeomorphisms between appropriate domains. A computational problem that has a fundamental importance in diffeomorphic matching is to determine a smooth path of diffeomorphisms that connects/interpolates a given pair of diffeomorphisms. More specifically, let Ω denote the 2D (image) domain and $\text{Diff}(\Omega)$ the space of diffeomorphisms, $\rho : \Omega \rightarrow \Omega$. Given two diffeomorphisms ρ_0, ρ_1 , a smooth path in $\text{Diff}(\Omega)$ connecting ρ_0, ρ_1 is a smooth map

$$\phi : [0, 1] \times \Omega \rightarrow \Omega,$$

such that for each $t \in [0, 1]$, $\phi(t, \cdot) \in \text{Diff}(\Omega)$ is a diffeomorphism and $\phi(0, \cdot) = \rho_0, \phi(1, \cdot) = \rho_1$. In this paper, we present a novel framework for computing such a smooth path for a given pair of diffeomorphisms ρ_0 and ρ_1 . However, evaluating diffeomorphic path between two given diffeomorphisms is a challenging problem. This is mainly due to first, the infinite dimensionality of the space of diffeomorphisms and second, the lack of knowledge of the metric on this space.

One of the most influential framework in this context is the Large Deformation Diffeomorphic Metric Mapping framework (LDDMM) introduced in [2]. However, in this framework (applied to image registration), the diffeomorphic path ϕ is not explicitly computed; instead, a time-dependent vector field $v(t, \cdot)$ is computed by minimizing the cost functional (using an appropriate Sobolev norm on v)

$$E = \int_0^1 \|v(t, \phi(t, \cdot))\|_V^2 dt + \lambda \|I_1 - I_2 \circ \phi(t = 1, \cdot)\|_{L_2}^2, \quad (1)$$

where I_1, I_2 are the two input images, and the diffeomorphic path ϕ is obtained by integrating the time-dependent

*This research was in part supported by the NIH grant NS066340 to BCV

vector field v via the following ODE

$$\frac{d}{dt}\phi(t, \cdot) = v(t, \phi(t, \cdot)). \quad (2)$$

Beg *et al.* [2] have shown that the diffeomorphic paths determined by LDDMM stay in the space of diffeomorphisms. The LDDMM framework has been further elaborated in the context of evolution equations on groups of diffeomorphisms and momentum conservation [3]. For example, Younes *et al.* introduced geodesic equations on the group of diffeomorphisms based on the LDDMM in [3] and Sommer *et al.* introduced a multi-scale kernel bundle for LDDMM to increase the accuracy of LDDMM in [4]. The latter method has been named the Large Deformation Diffeomorphic Kernel Bundle Mapping (LDDKBM), and [5] introduced the evolution equations for LDDKBM.

These frameworks involve two different optimizations: image matching and smoothing of the diffeomorphic paths in the space of diffeomorphisms. However, there are certain issues in these joint frameworks that we will elaborate on now. First, the Sobolev metric in Eq.(1) may cause undesirable geometric distortion in the space of diffeomorphisms (Fig.1). Second, the smoothness constraint over a diffeomorphic path may lead to unsatisfactory correspondences between input images. Further, tuning the set of parameters for a convergent result is heavily data-dependent, and is often cumbersome and tedious.

One way to overcome these issues is to separate registration stage and the stage of estimating diffeomorphic paths. The key contribution in this paper is that we design a framework for interpolating the two given diffeomorphisms, which is formulated separately from the registration task. Instead of interpolating the diffeomorphisms in their space directly, we project diffeomorphisms on to the space of densities, whose geodesics are readily computed, and then lift the geodesics back to the space of diffeomorphisms efficiently.

More specifically, let M denote a compact n -dimensional Riemannian manifold, and $Diff(M)$ and $Diff_\mu(M)$ respectively denote the infinite-dimensional group of diffeomorphisms of M and its infinite-dimensional subgroup of volume-preserving diffeomorphisms. The quotient space $Diff(M)/Diff(M)_\mu$ was studied in depth and identified with the space of densities functions $Dens(M)$ on M in [6]. The latter space can then be canonically embedded into a sphere in the Hilbert space (See [6] for details); therefore, the geodesics on $Dens(M)$ can be computed readily. In LDDMM, the diffeomorphic path is computed directly in $Diff(M)$ without referencing the quotient space $Dens(M) = Diff(M)/Diff(M)_\mu$. However, in our method, the diffeomorphic path is first computed on the quotient space $Diff(M)/Diff(M)_\mu$, and the full path is then computed in $Diff(M)$ by lifting the geodesic




Figure 1: Top: from left to right, a cylindrical can, 50 degree rotated can, the color plot of diffeomorphic deformation fields deforming the original can, and the colormap scale respectively. Second row: the ground truth of the rotation sequence. Third row: the path of images warped by the diffeomorphisms along the path estimated using our method. Bottom: the path of images obtained using LDDMM.

path in $Dens(M)$ to $Diff(M)$. We remark that it is precisely because the geometry of the quotient space is available (and simple), the geodesic path can first be determined on the quotient space. The second step of lifting the path in $Dens(M)$ back to $Diff(M)$ is then formulated as a quadratic programming problem with bilinear constraints that minimizes point-wise L_2 -norm of deformation fields along the diffeomorphic path. The resulting optimization problem can be solved efficiently using the numerical algorithm in [7] that is based on the augmented Lagrangian method.

With the approach presented in this paper, we obtain a novel framework for interpolating diffeomorphisms, which is less sensitive to tuning parameters and more computationally efficient than LDDMM-based methods. This also means that our method provides a more general approach than the LDDMM-base methods. We validate the proposed framework by applying it to the problem of filling-in missing frames in a video footage. We choose this application since it can provide ground-truth comparisons and show potential for further applications.

2. Model formulation

In this section, we present the novel framework for computing the path of diffeomorphisms. In what follows, M will denote the (image) domain and μ ($d\mu$) an associated volume form on Ω . The estimation of geodesics on the space of diffeomorphisms requires the knowledge of its metric. However, in this paper, instead of working on the manifold of diffeomorphisms whose metric is unknown, we map diffeomorphisms to a known space, the space of densities, and take advantage of the fact that geodesics for the canonical metric on the sphere are the great circles. Our algorithm consists of the following two steps:

- Map a given pair of diffeomorphisms to the space of densities, $Dens(M)$ and computing the geodesic between the two projected points.
- Lift the geodesic in $Dens(M)$ back to $Diff(M)$.

2.1. The Space of Densities and Diffeomorphisms

Given a pair of diffeomorphisms mapping Ω to Ω , and without loss of generality, we can assume one of them is the identity map id and denote the other diffeomorphism by ρ . The goal is to compute the one-parameter family of diffeomorphisms, $\phi(t)$, between id and ρ , such that $\phi(0) = id$ and $\phi(t = t_f) = \rho$ where, $0 \leq t \leq t_f$.

Given an orientation-preserving diffeomorphism, let $Jac_\mu\phi$ denote the determinant of the Jacobian of ϕ . It is a positive function, and in particular, if ϕ is volume-preserving, $Jac_\mu\phi$ is the constant function 1. Furthermore, since ϕ is a diffeomorphism of M , the integral transform

$$\mu(M) = \int_M d\mu = \int_M \phi^*(d\mu) = \int_M Jac_\mu\phi d\mu.$$

shows that the integral of the positive function $Jac_\mu\phi$ is a constant (the μ -volume of M) for all $\phi \in Diff(M)$. Therefore, after normalization, $Jac_\mu\phi$ can be considered as a density function on M . Furthermore, $Jac_\mu\phi$ has the interesting factorization property:

$$Jac_\mu(\phi \circ \rho)(x) = Jac_\mu\phi(\rho(x)) \cdot Jac_\mu\rho(x).$$

It follows that $Jac_\mu(\phi \circ \rho) = Jac_\mu\phi$ if ρ is a volume-preserving diffeomorphism. Therefore, using the square-root of the density function, we can define the map Φ from $Diff(M)$ to a sphere in a Hilbert space via

$$\Phi : \phi \rightarrow f = \sqrt{Jac_\mu\phi} \quad (3)$$

since

$$\int_M \Phi^2(\phi) d\mu = \mu(M) \quad (4)$$

shows that $\Phi(\phi)$ belongs to a sphere with radius $\sqrt{\mu(M)}$. Using normalization, we can consider $\Phi(\phi)$ as a point on the unit sphere and for the two diffeomorphism id, ρ mapped on to the sphere by Φ , the unique geodesic between them can be readily computed using the formula

$$f(t) = \frac{1}{\sin(\theta)} [\sin(\theta - t)f_1 + \sin(t)f_2] \quad (5)$$

where $f_1 = \sqrt{Jac_\mu id}$, $f_2 = \sqrt{Jac_\mu\rho}$.

We remark that if ρ is a volume-preserving diffeomorphism, by definition of Φ, ρ and id are projected to the same point on the sphere. Hence, the geodesic joining the two projected points is degenerate – i.e., it consists of just one point – and the lifted diffeomorphic path in $Diff(M)$ is then a

path joining ρ and id consisting of only volume-preserving diffeomorphisms.

In this paper, diffeomorphisms are expressed as deformation vector fields. In the case of 2-D, $\phi(t) = (x + U(\mathbf{r}, t), y + V(\mathbf{r}, t))$ where $\mathbf{r} = (x, y)$, and $(U(\mathbf{r}, t), V(\mathbf{r}, t))$ denote the deformation vector fields. Then,

$$Jac_\mu\phi(t) = (1 + U_x)(1 + V_y) - U_y V_x = 1 + U_x + V_y + U_x V_y - U_y V_x, \quad (6)$$

and we have

$$f(t)^2 = Jac_\mu\phi(t)/\mu(M). \quad (7)$$

Eq.(6) gives the determinant of the Jacobian of ϕ , and U_i and V_i are the first-order derivatives with respect to $i \in \{x, y\}$, respectively. We note in passing that in 2D Eq.(7) represents a bilinear equation and in 3D, a trilinear equation.

2.2. Lifting the Geodesic Path to the Space of Diffeomorphisms

After a geodesic on the sphere is obtained, we have to lift this path back to $Diff(M)$. However, solving Eq.(7) does not yield a unique solution. This is because the equation is bilinear, and geometrically, this corresponds to the fact that Eq.(7) only requires $\phi(t)$ to lie on a $Diff(M)_\mu$ -orbit in $Diff(M)$, parametrized by the point $f(t)$ in the sphere.

Therefore, the computation of the lifted path must be regularized, and we propose using the L_2 smoothness of the deformation vector fields over time as the main regularization criterion. Even though there are other choices of norms for regularization—such as the Sobolev norm—the L_2 norm regularization was chosen because of its computational simplicity. The lifting problem now leads to a quadratic program with bilinear constraints for 2D domains

$$\min \int \left| \frac{dU(t)}{dt} \right|^2 + \left| \frac{dV(t)}{dt} \right|^2 d\mu dt \quad (8)$$

s.t. $Jac_\mu\phi(t) = f(t)^2 \mu(M)$.

The following section shows how to optimize Eq(8) numerically.

3. Numerical Solution

While Eq.(8) is formulated for continuous variables x, y and t , in practice, we have to work with discrete pixels and time, and its discrete version is given by

$$\min_{U_{ij}^t, V_{ij}^t} \sum_{i,j,t} \left[\left| \frac{U_{ij}^{t+1} - U_{ij}^{t-1}}{2\delta x} \right|^2 + \left| \frac{V_{ij}^{t+1} - V_{ij}^{t-1}}{2\delta y} \right|^2 \right] \delta t \delta x \delta y$$

s.t. $\mu(\Omega)(f_{ij}^t)^2 = (1 + U_{xij}^t + V_{yij}^t + U_{xij}^t V_{yij}^t - U_{yij}^t V_{xij}^t)$, (9)

where the superscript t and subscript i, j denote the discrete time index and pixels, respectively and the subscripts x and y in the constraints denote the first-order derivatives w.r.t. x and y , respectively. We set δx and δy to be one, and therefore the volume $\mu(\Omega)$ is the image size.

Geometrically, the determinant of the Jacobian $Jac_\mu\phi(t)$ is the ratio of the change of volume elements by $\phi(t)$ at domain points at time t . In our numerical scheme, we consider it as the change in area of a triangle in the mesh with four neighboring pixel points replaced by the deformation field at vertices. That is, $Jac_\mu\phi(t)$, can be treated as the area of a rectangle in the mesh, and it is formulated as follows.

$$\begin{aligned} Jac_\mu\phi_{ij}(t) &= 1 + 0.5(U_{xij}^F(t) + V_{yij}^F(t) + U_{xij}^F(t)V_{yij}^F(t) - U_{yij}^F(t)V_{xij}^F(t)) \\ &+ 0.5(U_{xij}^B(t) + V_{yij}^B(t) + U_{xij}^B(t)V_{yij}^B(t) - U_{yij}^B(t)V_{xij}^B(t)), \end{aligned} \quad (10)$$

where \bar{i} and \bar{j} denote $i + 1$ and $j + 1$ respectively, and $(i, j) \in [1, H - 1] \times [1, W - 1]$ with the domain size of $H \times W$. The superscripts F and B in Eq.(10) denote the forward and backward difference schemes of numerical derivatives respectively. After vectorizing all deformation vector fields, the optimization problem becomes a quadratic program with bilinear constraints given by:

$$\begin{aligned} \min \frac{1}{2}\mathbf{U}^\top \mathbf{R} \mathbf{U} + \frac{1}{2}\mathbf{V}^\top \mathbf{R} \mathbf{V} + \mathbf{B}_u^\top \mathbf{U} + \mathbf{B}_v^\top \mathbf{V} \\ \text{s.t. } \mathbf{h} = \mathbf{c} + \mathbf{C}_y \mathbf{V} + \mathbf{C}_x \mathbf{U} + \mathbf{D}_u(\mathbf{V})\mathbf{U} = 0, \\ \text{or } \mathbf{h} = \mathbf{c} + \mathbf{C}_x \mathbf{U} + \mathbf{C}_y \mathbf{V} + \mathbf{D}_v(\mathbf{U})\mathbf{V} = 0. \end{aligned} \quad (11)$$

In the above, $\mathbf{U}^\top = (\mathbf{U}^\top(t = 1), \dots, \mathbf{U}^\top(t = T - 1))$, and $\mathbf{B}_u^\top = (\mathbf{U}^\top(t = 1), \mathbf{0}^\top, \mathbf{U}^\top(t = T))$ where $\mathbf{0}$ is a vector of zeros with length $H \times W \times (T - 4)$ if we have a time sequence $t \in [1, T]$ including two boundary conditions. \mathbf{V} and \mathbf{B}_v are defined in the same way. In Eq.(11), \mathbf{R} is the matrix for the quadratic component and its size is $H \times W \times (T - 2)$, $\mathbf{C}_x \mathbf{U}$ and $\mathbf{C}_y \mathbf{V}$ correspond to the linear terms of \mathbf{U} and \mathbf{V} in Eq.(10), respectively. $\mathbf{D}_u(\mathbf{V})\mathbf{U}$ ($\mathbf{D}_v(\mathbf{U})\mathbf{V}$) corresponds to the bilinear terms in Eq.(10) after factoring out \mathbf{U} (\mathbf{V}). Finally, \mathbf{c} is the vectorization of $\mu(\Omega)(f_{ij}^t)^2$ in Eq.(9) plus a vector of ones.

We solve the optimization problem using the augmented Lagrangian method with penalty terms [7] by iteratively solving for the two blocks of variables \mathbf{V} , \mathbf{U} , fixing one while optimizing the other. With fixed \mathbf{V} , the problem is given by

$$\begin{aligned} \min \frac{1}{2}\mathbf{U}^\top \mathbf{R} \mathbf{U} + \mathbf{B}_u^\top \mathbf{U} + \lambda^\top \mathbf{h} + \frac{1}{2}c\|\mathbf{h}\|^2 \\ \text{s.t. } \mathbf{U} \in \mathbb{R}^m, \end{aligned} \quad (12)$$

where $\lambda \in \mathbb{R}^m$ is the Lagrange multiplier, $m = H \times W \times (T - 1)$. This is an unconstrained optimization problem and

it can be rewritten in a more transparent way

$$\begin{aligned} \min \frac{1}{2}\mathbf{U}^\top (\mathbf{R} + c\mathbf{H}_u^\top \mathbf{H}_u) \mathbf{U} \\ + \left(\mathbf{B}_u + \frac{1}{2}c((\mathbf{G}_u^\top \mathbf{H}_u)^\top + \mathbf{H}_u^\top \mathbf{G}_u) + (\lambda^\top \mathbf{H}_u)^\top \right)^\top \mathbf{U} \\ \text{s.t. } \mathbf{U} \in \mathbb{R}^m, \end{aligned} \quad (13)$$

where $\mathbf{H}_u = \mathbf{D}_u + \mathbf{C}_x$ and $\mathbf{G}_u = \mathbf{c}$, i.e., $\mathbf{h} = \mathbf{H}_u \mathbf{U} + \mathbf{G}_u = 0$.

4. Experimental Results

We evaluate the effectiveness and efficiency of the proposed framework using the difficult problem image interpolation for missing frames in videos. Specifically, for the experiment, an image sequence is extracted from a video and a few frames are selected as the key frames. We then apply our method to compute a diffeomorphic path connecting two adjacent key frames, and the intermediate frames are computed and warped by the appropriate diffeomorphism along the diffeomorphic path. These filled-in image sequences are compared with the original sequences by measuring the pSNR of the computed frames w.r.t. the true frames. In addition, we compared our results with those from the Large Deformation Diffeomorphic Metric Mapping (LDDMM) method.

An overall description of this filling-in process is as follows: First, a source and target images (key frames) are selected, and denoted by I_s and I_t respectively. Second, a registration map— or diffeomorphism denoted by ρ —between the two images is computed using an image registration method. We then set id as the diffeomorphism which maps the source image to itself, and in this way, we have the two diffeomorphisms to be interpolated, id and ρ . Next, we compute a diffeomorphic path, $\phi(t)$ joining id and ρ , using the proposed algorithm. Finally, with this computed diffeomorphic path $\phi(t)$, we can simulate an image sequence $I(t)$ between I_s and I_t simply by warping the image $I(t) = I_s \circ \phi(t)$.

A simple example is presented in Fig.1. We have two cylinders: one is positioned upright and the other one is the 50-degree-rotated version of the upright cylinder. We evaluate deformation fields that deform the upright cylinder to the rotated one keeping the boundary fixed. Next we generate intermediate images between the two images. The ground truth is generated by rotating the upright cylinder by 50/19 degree consecutively so that we have 20 sequential images. We present the sequences in Fig.1 (down-sampled for computational efficiency). Fig.1 provides the visual comparison, and they show that our framework produces results that are visually more appealing and acceptable than LDDMM. In addition, we quantify the comparison using pSNR score between the ground truth and the computed images. The




Figure 2: pSNR score plot: red lines and blue lines represent scores for our method and LDDMM respectively. Beginning at the leftmost plot, rotating cylinder experiment, bending motion experiment and dancing motion experiment.

Table 1: The averaged pSNR scores between ground truths and computed image sequences using our method and LDDMM

Method	Our method	LDDMM
Rotation test	46.24 dB	26.59 dB
Bending woman	48.25 dB	46.2 dB
Dancing woman	45.94 dB	No convergence

Table 2: Running time comparison between the proposed method and LDDMM. The reported times for the proposed method include registration and interpolation stages. For the dancing motion sequence, we were not able to find suitable parameters for LDDMM and LDDMM does not produce convergent results.

Method	The proposed method	LDDMM
Rotation test	24min 40sec	42min 20sec
Bending woman	16min 7sec	41min 57sec
Dancing woman	12min 30sec	No convergence

pSNR scores for the sequences are plotted in Fig.2, and the average scores are reported in Table 1.

The following subsections consider more practical examples. All resulting sequences and original video footages are provided as movie files in the supplemental material. The original footages are obtained from <http://www.greenscreenfilms.com/>

4.1. Bending motion sequence

In this experiment, we present a video sequence in which a person is depicting a bending motion. The video’s frame rate is 29.75 fps, and we have extracted 34 frames from videos. The images from the sequence are shown in Fig.3. However, due to space limit, we present only down-sampled sequence in the figure. We choose the first, the last and the seventeenth frames as keyframes so that there are fifteen missing frames between keyframes. In Fig.3, the two groups of missing frames in each sequence are highlighted with rectangular boxes, with the three remaining key frames. Next we compute two diffeomorphic deformation fields: one from the first to the seventeenth frame, and one from the seventeenth frame to the last frame. The missing frames are computed as described in the previous sec-

tion. The resulting frames are compared with the ground truth and also with results produced by LDDMM: visual comparison shown in Fig.3 and quantitative comparison reported in Table 1 and Table 2. Visually from Fig.3, there are no marked differences between the proposed method and LDDMM. The reason could be that the motion in the sequence is not particularly complicated. However, the results in Table 1, Table 2 and Fig.2 show that our method is more accurate and with a considerable shorter running time than LDDMM. Additionally, LDDMM requires more key frames (between the first and seventeenth frames) in order to ensure its convergence.

4.2. Dancing motion sequence

In this experiment, we work on a video sequence with a more complicated motion namely, dancing. Twenty-seven image frames have been sampled from the original video, and twenty-two frames have been removed, leaving five key frames. These keyframes and missing frames are shown in Fig.4. Due to the amount of motion variation, we have retained more key frames (five) than the previous experiment on bending motion (three). We compute the same number of missing frames using our method as in the previous experiment. Fig.4 shows the ground truth images in the sequence and the ones computed with our method. The pSNR scores are plotted in Fig.2, and the averaged pSNR score is reported in Table 1. We remark that for this experiment we were not able to find suitable parameters for LDDMM to produce convergent results.

5. Conclusion

In this paper, we have introduced a novel framework for computing a diffeomorphic path that interpolates the two given diffeomorphisms, and we have demonstrated its applicability to image interpolation problems for reconstructing missing or corrupted images in videos. Computing diffeomorphic paths is often challenging, and the difficulty mainly arises from the nonlinearity and the infinite dimensionality of the space. Furthermore, the lack of known and computationally accessible metric is also a substantial hurdle. However, some of these difficulties can be circumvented by appealing to the fact that the quotient space of $Diff(M)$ by (the subgroup of) volume-preserving diffeomorphisms can be identified with a convex subset of the unit sphere in a Hilbert space. Using this decomposition (loosely speaking), the geometry of the sphere, which is well-known and computationally straightforward, provides us with geodesics that can be readily computed. In particular, for the problems that require treating the volume-preserving diffeomorphisms as the nuisance parameters, the path computed by our method is exactly the geodesic paths for the problems. Comparing with LDDMM, our method does not require specifying a Riemannian metric and has fewer tunable parameters. Fur-




Figure 3: The original missing frames are highlighted with blue bounding rectangles, and the frames computed using our method and LDDMM are highlighted with red and green bounding rectangles, respectively. The three key frames are shown individually without bounding rectangles.




Figure 4: Missing frames from the video are highlighted with red bounding rectangles and the computed frames are highlighted using blue bounding rectangles. The remaining five key frames are also shown in the sequence.

thermore, its running time is considerably shorter than that of LDDMM. Preliminary experiments, with favorable comparisons to LDDMM, have demonstrated the effectiveness and efficiency of the proposed method.

References

- [1] S. Joshi, B. Davis, M. Jomier, and G. Gerig, Unbiased diffeomorphic atlas construction for computational anatomy, *NeuroImage*, 23, S151-S160, (2004)
- [2] M. F. Beg, M. I. Miller, and A. Trouvé and L. Younes, Computing large deformation metric mappings via geodesic flows of diffeomorphisms, *IJCV* 61(2), 139-157, (2005)
- [3] L. Younes, F. Arrate, and M. I. Miller, Evolutions equations in computational anatomy, *Neuroimages* 45, S40-S50, (Mar 2009)
- [4] S. Sommer, M. Nielsen, F. Lauze, and X. Pennec, A multi-scale kernel bundle for LDDMM: towards sparse deformation description across space and scales, *IPMI*, 624-635, (2011)
- [5] S. Sommer, F. Lauze, M. Nielsen, and X. Pennec, Kernel Bundle EPDiff: Evolution Equations for Multi-scale Diffeomorphic Image Registration, *SSVM*, 677-688, (2011)
- [6] Geometry of diffeomorphism groups, complete integrability and geometric statistics, arXiv:1105.0643, 41pp (2011)
- [7] L. Carotenuto and G. Raiconi, On the minimization of quadratic functions with bilinear constraints via augmented Lagrangians, *Journal of optimization theory and applications*, 55, 1, 23-36, (Oct. 1987)
- [8] A. Srivastava, I. Jermyn and S. H. Joshi Riemannian analysis of probability density functions with application in vision, *Proc. CVPR*, 18-23 (2007)