MIT Libraries | DSpace@MIT

MIT Open Access Articles

Spatio-Temporal Data Fusion for 3D+T Image Reconstruction in Cerebral Angiography

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: Copeland, Andrew D et al. “Spatio-Temporal Data Fusion for 3D+T Image Reconstruction in Cerebral Angiography.” IEEE Transactions on Medical Imaging 29.6 (2010): 1238–1251. Web. 5 Apr. 2012. © 2010 Institute of Electrical and Electronics Engineers

Published Version: https://doi.org/10.1109/TMI.2009.2039645

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Permanent Link: http://hdl.handle.net/1721.1/69959

Version: Final published version

Terms of use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

Massachusetts Institute of Technology

Spatio-Temporal Data Fusion for 3D+T Image Reconstruction in Cerebral Angiography

Andrew D. Copeland, Rami S. Mangoubi*, Mukund N. Desai, Sanjoy K. Mitter, and Adel M. Malek

***Abstract*—This paper provides a framework for generating high resolution time sequences of 3D images that show the dynamics of cerebral blood flow. These sequences have the potential to allow image feedback during medical procedures that facilitate the detection and observation of pathological abnormalities such as stenoses, aneurysms, and blood clots. The 3D time series is constructed by fusing a single static 3D model with two time sequences of 2D projections of the same imaged region. The fusion process utilizes a variational approach that constrains the volumes to have both smoothly varying regions separated by edges and sparse regions of nonzero support. The variational problem is solved using a modified version of the Gauss–Seidel algorithm that exploits the spatio-temporal structure of the angiography problem. The 3D time series results are visualized using time series of isosurfaces, synthetic X-rays from arbitrary perspectives or poses, and 3D surfaces that show arrival times of the contrasted blood front using color coding. The derived visualizations provide physicians with a previously unavailable wealth of information that can lead to safer procedures, including quicker localization of flow altering abnormalities such as blood clots, and lower procedural X-ray exposure. Quantitative SNR and other performance analysis of the algorithm on computational phantom data are also presented.**

***Index Terms*—Angiography, blood, brain, cerebral, flow, fusion, image, reconstruction, sparse, spatio-temporal, variational, vascular, 3D+T, 4D.**

I. Introduction

A. Motivation and Objective

HIGH resolution time sequences of 3D images that show the dynamics of blood flow would allow diagnostic capabilities previously unavailable in cerebral angiography. These sequences have the potential to allow image feedback during catheterization procedures that would facilitate the detection and observation of flow anomalies such as stenoses, aneurysms, and clots. Failure to detect partially or completely obstructive clots during procedures such as carotid stenting or aneurysm coiling could result in blood flow interruption to downstream brain tissue leading to ischemic stroke. Unfortunately current angiography systems do not provide the necessary sequences of images due to insufficient volumetric sampling rates. The methods presented in this paper employ images from current angiography systems and achieve the goal of producing an accurate time sequence of 3D blood flow intensity images.

We develop a variational energy approach, and obtain quantitative performance results on computer generated phantom data. The approach has also been applied to two real vasculature flow cases: one with and one without the presence blood clots. The results are encouraging, indicating the methodology is promising, and further clinical study based on real vasculature flow is now justified before the methodology can be applied to direct operators during procedures. The emphasis of this paper is on the 3D+T reconstruction from lower dimensional 2D+T data. For other parts of the work, such as those used to register the data, implementation speed up, a brief description is provided.

B. Background

Current angiography systems produce high-resolution 3D volumes by capturing a set of over a hundred X-ray images each from a different angle, spanning a total of greater than 180°. Using a reconstruction algorithm, these projections are combined to generate a single 3D image. The approximately 4 s it takes to capture a complete set of projections necessary for 3D reconstruction is clearly too long of a duration to observe the advancing front of the contrast agent within the blood at up to 100 cm/s [1]. Faster image sampling could capture time histories of the rapid blood flow in 3D, but would also expose patients to considerably greater X-ray doses. Currently, fusion of the 2D time series information and a single 3D data set must be performed mentally by skilled operators.

In [2], a method is presented for finding smoothly varying set of skeletons that show the dynamics of contrasted blood flow in coronary arterial trees. Additionally, [1], [3] provide an algorithm that uses 3D data along with 2D time sets. First, 2D arrival times are determined by using correlation with a template time series. The 3D volume information is then distilled into a tree-like structure that is combined with the 2D arrival times using simple flow constraints. The result is a set of contrast arrival times at discrete locations in 3D. Finally, model based flow quantification is also found in [4].

Manuscript received October 05, 2009; revised December 14, 2009; accepted December 14, 2009. First published February 17, 2010; current version published June 03, 2010. This work was supported in part by the C. S. Draper Laboratory Internal Research and Development funds, in part by the Neely foundation to the Department of Neurosurgery at Tufts Medical Center, in part by NIH/NIBIB Grant 1 R01 EB006161-01A2, and in part NIH/NHLBI Grant 1 R21 HL102685-01. *Asterisk indicates corresponding author.*

A. D. Copeland is with Draper Laboratory, Cambridge, MA 02139 USA, and also with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: acopeland@alum.mit.edu).

*R. S. Mangoubi is with the C. S. Draper Laboratory, Cambridge, MA 02139 USA (e-mail: rmangoubi@draper.com).

M. N. Desai is with the C. S. Draper Laboratory, Cambridge, MA 02139 USA.

S. K. Mitter is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

A. M. Malek is with the Department of Neurosurgery at Tufts Medical Center and Tufts University School of Medicine, Boston, MA 02111 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TMI.2009.2039645

0278-0062/$26.00 © 2010 IEEE

C. Approach

In this paper, we adopt a spatio-temporal data fusion approach that provides a 3D time series (3D+T) reconstruction of contrasted blood flow based on both the 2D time series and a 3D map. The reconstructed 3D time series provides far greater information than the 3D contrast arrival times extracted in [1], [3]. In addition to the 3D map, this approach constrains the reconstruction to be smoothly varying off of a set of edges in time and space and to contain sparse vasculature structures. The fusion based reconstruction process is a variational approach that can be extended by incorporating other constraints in a natural way. This comprehensive approach provides a framework that attempts to address the ill posed nature and associated ambiguities of 3D time series reconstruction from sparse 2D time series projections.

Specifically, these constraints are enforced by developing a variational energy formulation similar to the one used in [5], [6]. Our formulation fuses the static 3D volume and the 2D X-ray time series to produce a reconstructed time series of 3D volumes. The data is prepared by first registering images within each 2D time series, then segmenting the 3D volume to produce a map of the vasculature where the blood flow occurs, and lastly by finding the relative poses of each of the 2D time series with respect to the 3D volume. The solution of the variational problem provides a 3D time series that can be used to estimate additional blood flow characteristics including 3D arrival times.

We obtain results based on two time series projections from two different poses. Furthermore, a real time implementation of the algorithm is also a realistic goal as the algorithm can be parallelized or implemented on a multi-GPU system.

D. Contributions

The work in this paper makes several contributions. The methodological contributions include the following.

1) *Variational Formulation*. The reconstruction is made possible through an innovative variational formulation that includes a back projection model, a sparsity term, a three-dimensional constraint map, and a simultaneous smoothing and segmentation framework that also provides an implicit fluid flow model.
2) *Efficient Algorithm*. The spatio-temporal dynamics of the blood flow allow for the use of a more efficient Gauss-Seidel type algorithm for solving the Euler Lagrange equations. This algorithm provides a significant reduction in memory usage and computation. It can also be easily parallelized.

New data products are made available

1) *Reconstructed 3D time (3D+T) histories*. A time series of 3D volumes of the estimated contrasted blood flow.
2) *Contrast arrival times in 3D*. Thresholding of the 3D time history in time provides arrival times for the contrast front at each voxel location.
3) *Computational synthetic angiographic projections*. The 3D time series allows the computational generation of 2D synthetic projections from arbitrary poses.

These results combine to provide potential *clinical benefits*:

1) *Synthetic angiographic projections*. Currently, if an additional 2D+T projection perspective or pose is needed, contrast must be injected and new time instances captured. The 2D synthetic projections can reduce the need for additional X-ray exposure.
2) *Blood flow time of arrival 3D+T visualization*. The 3D spatio-temporal visualization provides the operator with a 3D view of the blood flow from all poses.
3) *Flow-anomaly detection*. The 3D+T reconstruction enables the quick detection and 3D localization of life threatening blood flow anomalies such as clots that can compromise perfusion and health of downstream brain tissue as demonstrated in Fig. 7

E. Organization

This paper is organized as follows. First, in Section II, we present the core notation and a simple mathematical model for the formation of angiographic images. In Section III, we describe the data preparation algorithms used to process the raw data consisting of one or more 2D time series and a 3D volume. Specifically, we provide a sketch of the algorithms used to do the following: find the groupwise alignment of the raw 2D time series, segment the 3D volume, and find the best 2D-3D registration consistent with the segmented 3D volume and the registered 2D time series. In Section IV, we present the mathematical formulation for the spatio-temporal fusion that includes constraints for simultaneous smoothing and segmentation, tomographic projections, sparseness, temporal smoothness, and the 3D map. In Section V, we discretize the formulation and provide a fully implementable algorithm that efficiently exploits the spatio-temporal structure of the problem. In Section VI, we show the results of applying the spatio-temporal data fusion algorithm to both patient-derived angiographic data and computational phantom data. The clinical results are used to motivate the application, while the phantom data are used to provide proof of accuracy. Section VII provides additional discussion. Finally, in Section VIII, we provide concluding remarks.

Notation and Image Model

The evolution of the three dimensional region of interest that contains the contrasted blood flow is denoted by $F(x,y,z,t)$ in the continuous domain and $F[i,j,k,n]$ in the discrete domain. The continuous and discrete versions of F are defined on a bounded region in $\mathbb{R}^4$ and $\mathbb{Z}^4$ respectively. The overall 3D volume taken while the flow of contrast is in steady state is denoted as $V(x,y,z)$ and $V[i,j,k]$, where V is similarly defined on a bounded region in either $\mathbb{R}^3$ or $\mathbb{Z}^3$, respectively. Similarly, the registered 2D time series are denoted as either $G(u,v,t)$, a function defined on a bounded region in $\mathbb{R}^3$, or $G[l,m,n]$, a function defined on bounded region in $\mathbb{Z}^3$.

Using a vector notation for the spatial arguments, $F(x,y,z,t)$ is written as $F(\mathbf{x},t)$ where $\mathbf{x}$ is a vector with components x, y, and z. Similarly the vector notation is used for $F[\mathbf{x}_p,n]$, $V(\mathbf{x})$, and $V[\mathbf{x}_p]$. The use of $[\mathbf{x}_p]$ is an abuse of notation that is interpreted as the bilinear interpolated value [10] of the corresponding function at the continuous spatial coordinate $\mathbf{x}_p = (x_p,y_p,z_p)^{\mathrm{T}}$ in 3D or $\mathbf{x}_p = (x_p,y_p)^{\mathrm{T}}$ for a function in 2D.

In the continuous case, the constraint map $\mathfrak{M}$ is defined to be the set of points where the blood vessel is thought to be present in the 3D function V, while in the discrete case, $\mathfrak{M}$ is an indicator function that is defined on the same domain as V and is equal to one where blood vessel is thought to be and zero otherwise. The multiple registered 2D time series are also written with the vector notation as $G[\mathbf{x}_q, n]$, where $\mathbf{x}_q$ is instead a 2D coordinate vector. The time series are defined so that X-rays from multiple views are represented as a single function. For example, in the continuous case, if the X-rays captured from two separate angles at a given time each have the same compact support of $[0, N] \times [0, M]$, they can be written as a single function with compact support of $[0, 2N] \times [0, M]$. This defines a single function on a larger domain in place of the two functions. An analogous approach is used in the discrete case.

The continuous tomographic projection operator $\boldsymbol{\mathcal{A}}$ transforms a function defined on $\mathbb{R}^4$ (3D+T) into a function defined on $\mathbb{R}^3$ (2D+T). Following [11], the discretization of $\boldsymbol{\mathcal{A}}$ can be described as a vector computation

$$(\mathbf{A}F)[\mathbf{x}_q, n] = \sum_p \mathbf{A}_{p,q} F[\mathbf{x}_p, n] \tag{1}$$

where p denotes the particular voxel and q a particular pixel and ray through the volume The operators $\boldsymbol{\mathcal{B}}$ and $\mathbf{B}$ are defined to be inverse operators of the continuous projection operator $\boldsymbol{\mathcal{A}}$ and discrete projection operator $\mathbf{A}$, respectively. The inverse $\mathbf{B}$ is a transformation of a function defined on $\mathbb{R}^3$ (2D+T) into a function defined on $\mathbb{R}^4$ (3D+T). For the continuous operator $\boldsymbol{\mathcal{B}}$, one choice is the right Moore-Penrose pseudoinverse

$$\boldsymbol{\mathcal{B}} = \boldsymbol{\mathcal{A}}^*(\boldsymbol{\mathcal{A}}\boldsymbol{\mathcal{A}}^*)^{-1}, \tag{2}$$

where the superscript * denotes the conjugate of the respective operator. The norm of order p of a vector x is defined as $\|x\|_p := \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}$ for $p \geq 1$.

II. Data Preparation

The raw angiographic data from the medical procedure is converted into a form that the spatio-temporal data fusion algorithm can use. In Section III-A, we describe the data acquisition process. The data preparation begins with the groupwise alignment of the raw time series described in Section III-B. In a parallel step, the 3D volume is segmented using the algorithm briefly presented in Section III-C. Finally, the relative pose of the segmented volume with respect to the registered time series is found using the fast 2D-3D registration algorithm in Section III-D.

A. Data Acquisition

During a procedure, physician operators are able to observe the flow of blood in the brain by releasing an X-ray opaque contrast agent into an artery in the neck via a femoral approach. The effect of the contrast is then observed using a Siemens Axiom Artis, a commercially available fluroscopy system with 3 dimensional rotational angiography capability. The methodology is general and can be applied on data from other similar systems. The measurements provide data in each of two forms: 1) a 3D volume $V[\mathbf{x}_p]$ taken in steady state by reconstructing X-ray exposures taken from more than a hundred different perspectives, 2) one or more time series of 2D X-rays $I_n[\mathbf{x}_p]$ taken from a fixed angle during the injection of contrast. The 2D image sequences used in this work are captured from two separate angles at rates of up to 8 frames per second.

B. Groupwise Registration of 2D Time Series

In cerebrovascular procedure where a catheter is used, the patient is conscious, so head motion is inevitable. The first step of data preparation removes the effect of patient motions from the time sequences using a groupwise registration algorithm based on the least squared error metric described in [12].

C. 3D Segmentation

To fuse the two data sets, the 3D volume and the registered 2D time series, the location of the blood vessels need to be extracted from the 3D volume. Because the volumes in the study may be as large as 512^3 voxels, we use a simple threshold on the static volume V to produce the constraint map $\mathfrak{M}$. The map is later used in 2D-3D registration and in the variational fusion algorithm discussed in Section IV. Clearly better segmentation will yield better results. Potential gains from improved segmentation are discussed in Section VII-D.

D. 2D-3D Registration

The final hurdle in preparing the data is to determine the relative pose of the 2D time series in relation to the 3D volumes. This pose, which is often not accurately given *a priori*, can be found by finding the projection of the volume, known as a digital reconstructed radiograph or DRR, that is most similar under correlation to a representative image of the registered time series. The search is initialized with the angiographic system's measured pose information. We find the representative image by subtracting the minimum intensity from the maximum intensity at each pixel location within the time series. This process isolates the vasculature within the time series. The rate limiting step of the registration is the determination of the DRRs from the desired geometry. The term approximate DRR is used to refer to a projection of a symbolic 3D volume instead of the actual 3D intensities themselves. This greatly reduces the complexity and in many cases provides similar results to the standard DRR. Several works in the literature provide efficient algorithms for determining DRR's or approximate DRR's [13]–[20]. Each of these, with the exception of [20], requires a significant amount of computation to transform the data into a more convenient form for the generation of the DRR or approximate DRR.

Our proposed algorithm approximates the DRR by projecting a triangular mesh representation of the blood vessel structure contained within the 3D volume. It begins by first producing a mesh using the Marching Cubes algorithm [21] on the segmented data. This algorithm creates a collection of points that combine to form a surface that is sparse in size compared to the 3D volume. Each triangle in the mesh can then be transformed and projected to form the approximate DRR. This reduces the computation time by a factor of one hundred over the full calculation of the DRR using sums across rays in the original volume as in [22]. After the mesh is projected into the imaging plane,

the pixels location within each triangle are set to one. Mesh reduction algorithms such as [23] and the Matlab function reducepatch [24], used here, can further speed up the algorithm by reducing the number of triangles without affecting the qualitative shape of the mesh. By reducing the number of triangles by 90% we achieved a factor of 6 speedup. A more detailed description of this process along with results can be found in [12]. On a 2.0 GHz Intel Core 2 Duo machine the generation of a single DRR took 0.039 s.

III. Spatio-Temporal Data Fusion Formulation

Spatio-temporal data fusion reconstructs a 3D time series that shows the flow of blood during a medical procedure. The input data consists of the two registered time series $I_k[\mathbf{T}_k(x_i)]$, that we now denote as G, the constraint map $\mathfrak{M}$, and the accompanying 3D pose information. This work builds on the variational formulation developed in [5], [25] for smoothing images while simultaneously finding the set Γ that best segments the image. In [26], [27] the edge set Γ was replaced with the spatial edge indicator function w. Specifically, spatio-temporal data fusion generalizes to the angiography application the stable numerical implementation of the [26], [27] algorithm found in [28], [29].

Spatio-temporal data fusion constrains the reconstruction by taking advantage of the structure and attributes of blood vessels in the following five ways: 1) the reconstruction is consistent with the time projection model, 2) the reconstruction is smoothly varying off of edges in space [25], 3) blood flow over time follows physical laws, 4) blood generally appears in sparse convex structures, and 5) the reconstruction is nonzero only within the 3D constraint map. Using these five attributes helps resolve ambiguities and minimize the presence of artifacts in the reconstructed 3D time series. The constraints are imposed using energy functionals with terms that penalize the violation of constraints.

A. Variational Formulation

The variational minimization problem used to impose constraints on the reconstructed 3D time series $F(\mathbf{x}, t)$ is shown in (3) at the bottom of the page where w is the spatial edge function; v is the temporal edge function; $\boldsymbol{\mathcal{A}}$ and $\boldsymbol{\mathcal{B}}$ are the projection and back projection operators respectively; G is the measured two dimensional data following registration; $\mathfrak{M}$ is the constraint map that specifies the location of the blood vessels in 3D; t_{end} is the duration of the time sequences; and α, β, γ, and ρ are weights. The first term in (3) ensures that the 2D projections of the reconstructed volume are consistent with the measured time series of 2D projections. Note that this approach is different from using the inverse of the projection data with the reconstructed image in the data fidelity term, i.e., $F - \boldsymbol{\mathcal{B}}G$. The goal of the data fidelity term in (3) is to match the projection of the reconstructed data, not to match the final image to the output of any particular inverse tomographic algorithm. The second and third terms ensure that the reconstruction is smoothly varying off of the edges in space and time, respectively. The fourth term ensures that the reconstruction is sparse. The last four terms ensure that the spatial and temporal edge terms are both smoothly varying and sparse.

The back projection operator $\boldsymbol{\mathcal{B}}$ of the data fidelity term is different from [29] where Lagrange multipliers were used to weight differences in the projections. A different approach that instead uses level sets to solve a variational formulation of the tomography problem is presented in [31]. Note that the functional is defined on the domain $\mathfrak{M} \times [0, t_{\text{end}}]$ and integrated with respect to $d\mathbf{x}dt$. Because the blood vessels are spatially sparse, the knowledge of their locations can further regularize the fusion process. Specifically, the reconstruction is confined to the constraint map $\mathfrak{M}$. The map takes on a value of one at the locations in V where the vasculature is thought to be and zero otherwise.

The under determined set of measurements result in commonly known back projection artifacts [22]. These artifacts can be minimized by restricting the contrast to the expected tight convex small regions. Reconstructions of this type can be encouraged by penalizing the $\|F\|_0$ or zero-norm constraint, which represents the desired non-vanishing support of F. Under this sparsity term, the minimization is not convex and is thus difficult to solve. To create an optimization process that is convex with respect to F while w fixed and with respect to w with F fixed, we follow recent work in [32]–[35] and instead use the 1-norm constraint $|F|$. Under certain conditions on the overall sparseness, number of nonzero terms measured by the 0-norm, of the underlying data, the 0-norm and 1-norm optimization problems achieve identical results [32], [36]. Unlike the work in [37], where linear programming is used to minimize the sparseness under a separate constraint on the projections, the formulation in (3) minimizes the two in a single term. For appropriate choices of γ the results should be similar because both optimization techniques minimize the same two terms.

Note that the back projection operator in the first term of the variational formulation brings all terms into the same spatio-temporal domain, thereby facilitating the utilization of Calculus

$$\min_{F,w,v} E(F,w,v) = \int_{\mathfrak{M}\times[0,t_{\text{end}}]} \left(\alpha\boldsymbol{\mathcal{B}}(\boldsymbol{\mathcal{A}}F - G)^2 + \beta(1-w)^2\|\nabla_x F\|^2 + \beta(1-v)^2\|\nabla_t F\|^2 + \gamma|F| + \frac{\rho}{2}\|\nabla_x w\|^2 + \frac{w^2}{2\rho} + \frac{\rho}{2}\|\nabla v\|^2 + \frac{v^2}{2\rho} \right) d\mathbf{x}dt \tag{3}$$

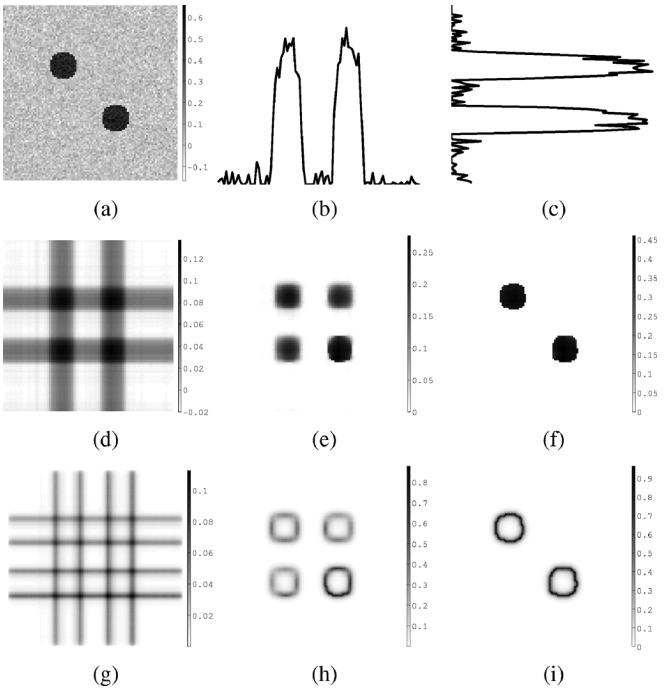


Fig. 1. The first row shows (a) a sample image with added noise, along with projections in (b) the row and (c) column directions. The middle row shows reconstruction using (d) neither map nor sparsity, (e) the L_1 norm based sparsity term, and (f) both the sparsity term and a map to resolve ambiguity. In the last row, (g), (h), and (i) are the edge maps corresponding to, respectively, (d), (e), and (f).

of Variations. Using the back projection operator defined in (2), the resultant Euler Lagrange equations associated with the minimization of functional E found in (3) are

$$\partial_F E = \alpha \mathcal{B}(\mathcal{A}F - G) - \nabla_x \cdot \beta(1-w)^2 \nabla_x F - \nabla_t \cdot \beta(1-v)^2 \nabla_t F + \frac{\gamma}{2}\frac{F}{|F|} = 0 \quad (4)$$

for F,

$$\partial_w E = 2\beta(w-1)\|\nabla_x F\|^2 + \frac{w}{\rho} - \rho\nabla_x^2 w = 0 \quad (5)$$

for w, and

$$\partial_v E = 2\beta(v-1)\|\nabla_t F\|^2 + \frac{v}{\rho} - \rho\nabla^2 v = 0 \quad (6)$$

for v. These coupled equations are used to find the desired minimum of (3) that occurs when the three gradients are equal to zero and the second derivatives are positive. Unless the function in (3) is globally convex with respect to F, w, and v, the minimum is not necessarily global.

B. Impact of Map Constraint and Sparsity Terms

The 4D spatio-temporal process is difficult to visualize. Fortunately, many of the salient features of the 4D reconstruction can be observed in a 2D spatial reconstruction problem. In fact, the example we present in Fig. 1 can be thought of as a slice from 3 dimensional columns extending into and out of the page. The impact of the map constraint and the sparsity term for the equivalent 4D reconstruction problem are similar in nature.

The first row of Fig. 1 shows a sample cylinder slice with added noise along with projections in the row and column directions. The middle row shows the reconstructions using neither the map nor the sparsity term, [Fig. 1(a)], using the sparsity term and no map [Fig. 1(b)], and lastly using both the sparsity term and the map [Fig. 1(c)]. The third row shows the edge functions for the corresponding smoothed reconstructions in the second row. The left and middle reconstruction show the presence of two "ghost" circles where none appeared in the original. The "ghost" circles are however consistent with the two projections and can be interpreted as giving an equal weighting to each of the circles when insufficient information is available to decide between them. The left reconstruction contains artifacts of back projection that are eliminated by using the sparsity term shown in the center reconstruction. The right figures in the second and third rows show the result of introducing the map to resolve the "ghost circle" ambiguity. In this example, the shape of the map is not important; it must only reject the regions containing the two "ghost" circles. We examine the impact of other terms in the energy functional and provide quantitative results in Section VI-A.

IV. Implementation of the Spatio-Temporal Algorithm

In Section V-A, we present the discretization of the continuous Euler Lagrange Equations shown in (4)–(6). The notation and choice of edge function draw heavily from the 2D work done in [29], which provides both an accurate discretization and one of first few implementable algorithms in the literature [6]. Our works extends the 2D work in [29] into four dimensions, 3 space and 1 time, and includes different terms used in the angiography application. In Section V-C the discretization is extended into an algorithm that exploits the spatio-temporal structure of the problem. Lastly, in Sections V-B and V-D we present a discussion of a few details of the implementation and a few potential enhancements.

A. Discretization

The reconstructed volumes $F(\mathbf{x}_i, n)$ are defined on a rectangular lattice that only includes points where the map function is equal to one. The spatial edge function $w(\mathbf{x}_i, n)$ is defined on a similar rectangular lattice produced by placing a point between each pair of adjacent points in space of the lattice that supports $F(\mathbf{x}_i, n)$. Similarly, the temporal edge function $v(\mathbf{x}_i, n)$ is defined by placing a point between each pair of adjacent points in time of the lattice that supports $F(\mathbf{x}_i, n)$. This discretization provides both an accurate and simple implementation. The lattice for $F(\mathbf{x}_i, n)$ is defined on the subset of $\mathbb{Z}^4$ where the map $\mathfrak{M} = 1$ as

$$\mathfrak{L}_F = \{(i,j,k,n) : i \in \{1,\ldots,M\}, j \in \{1,\ldots,P\}, k \in \{1,\ldots,Q\}, n \in \{1,\ldots,N\}, \mathfrak{M}[i,j,k] = 1\}. \quad (7)$$

The nearest neighbors of the image at point $(\mathbf{x}_i, n) \in \mathfrak{L}_F$ are the points in the set

$$\mathfrak{N}_F(\mathbf{x}_i, n) = \{(\mathbf{x}_j, m) \in \mathbb{Z}^4 : |\mathbf{x}_j - \mathbf{x}_i| + |n - m| = 1\}. \quad (8)$$

The lattices for the spatial edge function $w(\mathbf{x}_i, n)$ and the temporal edge function $v(\mathbf{x}_i, n)$ are respectively,

$$\mathfrak{L}_w = \Big\{ \frac{\mathbf{x}_l + \mathbf{x}_m}{2}, q : (\mathbf{x}_l, q) \in \mathfrak{L}_F, \\ (\mathbf{x}_m, r) \in \mathfrak{N}_F(\mathbf{x}_l, q) \cap \mathfrak{L}_F, q = r \Big\} \quad (9)$$

and

$$\mathfrak{L}_v = \Big\{ \mathbf{x}_l, \frac{q+r}{2} : (\mathbf{x}_l, q) \in \mathfrak{L}_F, (\mathbf{x}_m, r) \in \mathfrak{N}_F(\mathbf{x}_l, q), \\ \mathbf{x}_l = \mathbf{x}_m, |q - r| = 1 \Big\}. \quad (10)$$

The resultant nearest neighbors for w are

$$\mathfrak{N}_w(\mathbf{y}_i, n) = \Big\{ \mathbf{y}_j = \frac{\mathbf{x}_l + \mathbf{x}_m}{2}, n : (\mathbf{x}_l, n) \in \mathfrak{L}_F, \\ (\mathbf{x}_m, n) \in \mathfrak{N}_F(\mathbf{x}_l, n), |\mathbf{y}_j - \mathbf{y}_i| = \frac{1}{\sqrt{2}} \Big\} \quad (11)$$

and for v are

$$\mathfrak{N}_v(\mathbf{z}_i, n) = \Big\{ \mathbf{z}_j \in \mathbb{Z}^3, p + \frac{1}{2} \in \mathbb{Z}, |\mathbf{z}_j - \mathbf{z}_i| + |n - p| = 1 \Big\}. \quad (12)$$

For a simpler implementation, the terms w and v are replaced with $1-w$ and $1-v$, respectively. The interpretation of the value of 1 and 0 are simply reversed. The resultant discrete functional is shown in (13) at the bottom of the page. Without loss of generality, the term γ is replaced with 2γ in the discretization.

B. Discrete Implementation of Back Projection

The discrete back projection operator $\mathbf{B}$ in (13) is difficult to calculate directly, so we instead solve for F using the convergent iterative method, known as the Simultaneous Algebraic Reconstruction Technique (SART) [38]

$$\mathbf{B}_{q,p} = \frac{\mathbf{A}_{p,q}}{\sum_{\hat{q}} \mathbf{A}_{p,\hat{q}} \sum_{\hat{p}} \mathbf{A}_{\hat{p},q}} \quad (14)$$

where, as before, q denotes the particular pixel and ray and p the particular voxel. A detailed discussion on determining $\mathbf{A}_{q,p}$ and $\mathbf{B}_{q,p}$ is provided in [12]. The terms $\hat{q}$ and $\hat{p}$ are used for the summation over all rays (pixels) and all voxels respectively.

C. Exploiting Structure of Space-Time Constraints

This section develops an iterative numerical algorithm based on the Euler Lagrange equations for (13) that takes advantage of the structure of the constraints. Instead of using a Jacobi-type descent, a Gauss-Seidel [39] type descent is used that alternates between solving for reconstructions at an individual times and solving for the time coupling term. This breaks the total 3D time series reconstruction into successive 3D reconstructions. The Gauss-Seidel-type iteration alternates between solving for the individual 3D volumes and solving for the temporal edge term v. This decoupling allows the algorithm to use less memory by writing variables to disk that are complete after the final iteration on a volume and by loading needed variables from disk before iterations are run on the next volume. By using the Gauss-Seidel based method the memory requirements were reduced from 10 GB to 600 MB for a series of 10 volumes.

In order to represent the iterative nature of this algorithm, the variables can be modified to denote the value at a particular iteration k by adding the additional argument to the respective function such as in $F[\mathbf{x}_i, n, k]$, $w[\mathbf{y}_i, n, k]$, and $v[\mathbf{z}_i, n, k]$. In a typical Jacobi-type algorithm, all the variables at each spatial location and time at iteration $(k-1)$ are saved for determining the variables at the next iteration (k). This type of algorithm requires keeping two copies of the entire state F, w, and v, one at iteration (k) and another at iteration $(k-1)$. This doubles the memory requirements of this already demanding application.

The Gauss-Seidel type algorithm provides an in place algorithm for finding the reconstruction of F, w, and v at each time. This process can be extended to further decouple the determination of v at each time. Specifically, given v for each spatial location at time $n + 1/2$ and $n - 1/2$, $w[\mathbf{y}_i, n, k+p]$ and $F[\mathbf{x}_i, n, k+p]$ can be determined for several gradient descent

$$\begin{aligned} E(F,w,v) = & \sum_{(\mathbf{x}_i,n)\in\mathfrak{L}_F} \left(2\gamma |F[\mathbf{x}_i,n]| + \sum_{(\mathbf{x}_i,m)\in\mathfrak{N}_F(\mathbf{x}_i,n)} \alpha \Big(\sum_l \boldsymbol{B}_{l,i} \Big(\sum_j \boldsymbol{A}_{j,l} F[\mathbf{x}_j,n] - G[\mathbf{x}_l,n] \Big) \Big)^2 \right. \\ & + \sum_{(\mathbf{x}_j,n)\in\mathfrak{N}_F(\mathbf{x}_i,n)} \frac{\beta (F[\mathbf{x}_i,n] - F[\mathbf{x}_j,n])^2}{2} \left(w \left[\frac{\mathbf{x}_i + \mathbf{x}_j}{2}, n \right] \right)^2 \\ & \left. + \sum_{(\mathbf{x}_i,m)\in\mathfrak{N}_F(\mathbf{x}_i,n)} \frac{\beta (F[\mathbf{x}_i,n] - F[\mathbf{x}_i,m])^2}{2} \left(v \left[\mathbf{x}_i, \frac{n+m}{2} \right] \right)^2 \right) \\ & + \sum_{(\mathbf{y}_i,n)\in\mathfrak{L}_w} \left(\frac{(1 - w[\mathbf{y}_i,n])^2}{2\rho} + \sum_{(\mathbf{y}_j,n)\in\mathfrak{N}_w(\mathbf{y}_i,n)} \rho (w[\mathbf{y}_i,n] - w[\mathbf{y}_j,n])^2 \right) \\ & + \sum_{(\mathbf{z}_i,n)\in\mathfrak{L}_v} \left(\frac{(1 - v[\mathbf{z}_i,n])^2}{2\rho} + \sum_{(\mathbf{z}_j,m)\in\mathfrak{N}_v(\mathbf{z}_i,n)} \frac{\rho (v[\mathbf{z}_i,n] - v[\mathbf{z}_j,m])^2}{4} \right). \end{aligned} \quad (13)$$

iterations $p = 1, \ldots, P$ instead of just one. The number of iterations P can be set to a number greater than one to avoid the cost of switching from the determination of both w and F to the determination of $v[\mathbf{z}_i, n-1/2, k+p]$. For P iterations, this process decouples the individual 3D reconstructions from the smoothing of each spatial location across time. The reconstruction still begins at time one, but now P iterations are performed at each time instead of just one. Following [29], the iterative scheme for simultaneously determining $F[\mathbf{x}_i, n, k+p]$ and $w[\mathbf{y}_i, n, k+p]$ is given by

$$\begin{aligned}
&F[\mathbf{x}_i, n, k+p] \\
&= F[\mathbf{x}_i, n, k+p-1] \\
&\quad - c_F \Bigg(\alpha \sum_l \boldsymbol{B}_{l,i} \Big(\sum_j \boldsymbol{A}_{j,l} F[\mathbf{x}_j, n, k+p-1] - G[\mathbf{x}_l, n] \Big) \\
&\quad + \gamma \frac{F[\mathbf{x}_i, n, k+p-1]}{|F[\mathbf{x}_i, n, k+p-1]|} \\
&\quad + \beta (F[\mathbf{x}_i, n, k+p-1] - F[\mathbf{x}_i, n+1, k]) \Big(v\Big[\mathbf{x}_i, n+\frac{1}{2}, k\Big] \Big)^2 \\
&\quad + \beta (F[\mathbf{x}_i, n, k+p-1] - F[\mathbf{x}_i, n-1, k+P]) \\
&\quad \cdot \Big(v\Big[\mathbf{x}_i, n-\frac{1}{2}, k\Big] \Big)^2 \\
&\quad + \sum_{(\mathbf{x}_j, n) \in \mathfrak{N}_F(\mathbf{x}_i, n)} \beta (F[\mathbf{x}_i, n, k+p-1] - F[\mathbf{x}_j, n, k+p-1]) \\
&\quad \cdot \Big(w\Big[\frac{\mathbf{x}_i + \mathbf{x}_j}{2}, n, k+p-1\Big] \Big)^2 \Bigg)
\end{aligned} \tag{15}$$

where l denotes the index of a particular ray or pixel and j denotes the index of a particular voxel, and

$$\begin{aligned}
w[\mathbf{y}_i, n, k+p] &= \frac{1}{2} w[\mathbf{y}_i, n, k+p-1] \\
&\quad + c_w \left(\frac{1}{2\rho} + \sum_{(\mathbf{y}_j, n) \in \mathfrak{N}_w(\mathbf{y}_i, n)} \rho w[\mathbf{y}_j, n, k+p-1] \right)
\end{aligned} \tag{16}$$

using the weighting terms

$$c_F = \frac{1}{2} \left(\alpha + 8\beta + \frac{\gamma}{|F[\mathbf{x}_i, n, k+p-1]| + \epsilon} \right)^{-1} \tag{17}$$

and

$$\begin{aligned}
c_w = \frac{1}{2} \Bigg(&\beta (F[\mathbf{x}_l, n, k+p-1] - F[\mathbf{x}_m, n, k+p-1])^2 \\
&+ \frac{1}{2\rho} + 6\rho \Bigg)^{-1}.
\end{aligned} \tag{18}$$

To arrive at the gradient descent equation found in (15), we approximated the matrix product $\mathbf{AB}$ with the identity matrix in the first variation with respect to F of (13). After the variables $F[\mathbf{z}_i, n-1/2, k+P]$ and $F[\mathbf{z}_i, n+1/2, k+P]$ are obtained using P iterations of (15) and (16),

$$\begin{aligned}
v[\mathbf{z}_i, n, k+p] &= \frac{1}{2} v[\mathbf{z}_i, n, k+p-1] \\
&\quad + \frac{c_v}{2} \Bigg(\rho v[\mathbf{z}_i, n-1, k+P] \\
&\quad + \frac{1}{\rho} + \rho v[\mathbf{z}_i, n+1, k] \\
&\quad + \sum_{\substack{(\mathbf{z}_j, m) \in \mathfrak{N}_v(\mathbf{z}_i, n) \\ n=m}} \rho v[\mathbf{z}_j, n, k+p-1] \Bigg),
\end{aligned} \tag{19}$$

is evaluated for P iterations using the weight

$$\begin{aligned}
c_v = \frac{1}{2} \Bigg(&\beta \Big(F\Big[\mathbf{z}_i, n+\frac{1}{2}, k+P\Big] - F\Big[\mathbf{z}_i, n-\frac{1}{2}, k+P\Big] \Big)^2 \\
&+ \frac{1}{2\rho} + 8\rho \Bigg)^{-1}
\end{aligned} \tag{20}$$

D. Implementation Details

The algorithm is started by initializing F to one inside the constraint map $\mathfrak{M}$ and zero outside, w to one over the whole domain, and v to zero over the whole domain. The use of other initial conditions change the algorithms speed and point convergence. For example, if F is set to zero instead of one, the convergence is very slow because the $1/(F+\epsilon)$ term in c_F, defined in (17), makes the gradient step very small. A corollary to this observation is that larger values of F yield bigger gradient steps. As desired, this difference in speed causes the functional to favor regions with smaller support. Finally, the 2D time series are scaled so that the numerical values of F are less than one and greater than or equal to zero.

V. Spatio-Temporal Results and Validation

This section shows results for the spatio-temporal data fusion algorithm on both real angiographic data and computational phantom data. Section VI-A provides a quantitative analysis of the algorithm using computational phantom vasculature blood flow data. In Section VI-B, results are shown for data sets taken of two different patients: set A contains a time series acquired at a high frame rate and set B contains time sequences acquired at a lower frame rate taken both with a clot is present and after the clot is removed.

A. Performance Evaluation on Computational Phantom Data

We use computational phantom data to evaluate our algorithm. We obtain quantitative performance results based on blood intensity residual comparison and the time of arrival for each voxel of the contrasted blood flow front. The quality of the 3D+T reconstruction in terms of signal to noise ratio can be attributed to the formulation's ability to provide simultaneous spatio-temporal smoothing and segmentation during reconstruction. We continue the discussion in Section IV-B and distill our experience in examining the effects of various terms and weights in the energy functional (3).

1) Quantitative Results: Contrasted Blood Flow Time of Arrival and Residual Comparison: Fig. 2(a) shows a view from a time series of a particular blood flow in a 3D computational phantom vasculature. Gaussian noise is added to the 3D phantom time series data. The resultant values are then clipped to be nonnegative. From the noisy phantom data, 2D

TABLE I
QUANTITATIVE RESULTS FROM COMPUTATIONAL PHANTOM

Time	Time of arrival % accuracy	Intensity (RSME) / *SNR Improvement (dB)* Pose 1	Pose 2	Pose 3	Pose 4
3	98.70	12.47 / *18.08*	12.27 / *18.22*	11.39 / *18.87*	10.52 / *19.56*
6	97.16	10.24 / *19.80*	9.66 / *20.30*	8.93 / *20.98*	8.78 / *21.13*
9	95.75	10.92 / *19.23*	10.24 / *19.79*	9.52 / *20.43*	9.60 / *20.35*
10	95.35	11.26 / *18.97*	10.66 / *19.46*	9.93 / *20.06*	9.91 / *20.08*

Numerical results based on comparison of the reconstructed data and the original computer generated phantom. The first column indicates the time instances, while the second the percentage of voxels that are assigned a correct time of arrival for each of four instances of the time series. The 3rd to 6th column gives the percentage error based on root square mean error (RSME) in the residual of the projection from 4 different poses, as well as the improvement in dB (italic), for the four time instances. A residual error visualization is shown in Fig. 2(e).

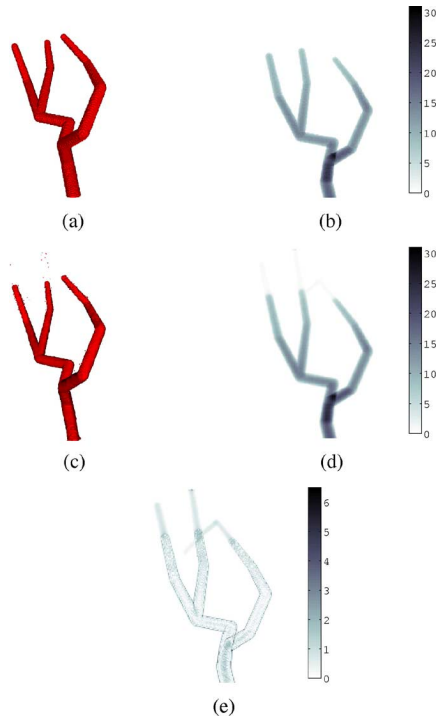


Fig. 2. Original and reconstructed computational phantom time series. The top row shows (a) the truth computational 3D+T phantom vasculature flow at a particular time in the time series and (b) a 2D synthetic X-ray projection of a noisy version of (a) from a particular pose. The middle row shows (c) the algorithm's 3D reconstruction of (a) from a time series of synthetic projections taken from two different poses and (d) a 2D synthetic X-ray projection from the reconstruction in (c) using the same pose as in (b). The bottom row shows the magnitude of the 2D residual (e) = |(b)–(d)|, with the gray scale adjusted for visibility. The residual (e) would be invisible if shown at the scale of (b) or (d). Table I presents numerical results for additional poses and time instances.

X-ray synthetic projections are obtained for two perspectives or poses. The effective signal to noise ratio of the synthetic 2D X-ray projections, obtained by taking the ratio of the square root of the power in the noiseless projection, to the square root of the power in the noise, varies between 0.71 and 0.79 for all projection measurements at most instances, or an equivalent of -6.8 to -4.65 dB.

Fig. 2(b) shows a noisy 2D synthetic projection from one pose of a noisy version of the 3D+T vasculature flow in Fig. 2. These time series projections taken from two different poses are in turn used to reconstruct the 3D+T time series of the flow. Fig. 2(c) shows the algorithm's 3D time series reconstruction for the same pose and time instance of Fig. 2(a), with a corresponding 2D synthetic X-ray projection from the reconstruction shown in Fig. 2(d). The projections of Fig. 2(b) and (d) are taken from the same pose at the same instance, and the residual difference is shown in Fig. 2(e), after gray scale adjustment for enhanced visibility; the residual, which shows no significant artifacts, is not visible at the scale of either (b) or (d).

The 3D+T reconstructed image intensities are appropriately thresholded to assign each voxel a contrast arrival time. Table I, second column, gives the percentage of voxels that are assigned the correct contrast time of arrival. In all cases, at least 95% of voxels are given a correct assignment.

Next, for quantitative residual comparison, 2D synthetic projections from four poses at four time instances are obtained from both the original and reconstructed volumes. These projections, whose pose is different from those of the two original projections used to reconstruct the 3D+T time series, are compared at each of the four time series instances.

For each pose's residual at each of the four instances the root mean square error (RMSE) is computed and normalized by the square root of the power in the projection measurement. Table I, third to sixth column, displays the percentage error for each pose at each instant, as well as the equivalent dB gain. The errors vary between 8.78% and 12.27%. The equivalent dB improvment varies between 18.029 and 20.08 dB.

2) Robustness to Weight Parameter Values: We also tested the robustness of these results with respect to the choice of the weigh parameters α, β, γ, and ρ in (3). On the average, results, over about a dozen trials, were for most views at most instances comparable. For instance, a different set of weights, gave a maximum error in Table I of 12.57%, and a minimum one of 2.42%, with most errors below 9%.

3) Effects of Terms in Energy Functional: The effect of the sparsity term and the map were investigated in Section IV-B. The effect of the other terms on the final reconstruction is now examined on the phantom model of Fig. 2 by setting the different weights to zero one at a time. In sum, the edge functionals are needed for simultaneous smoothing and segmentation while reconstruction evolves simultaneously across space and time [5], [6].

The spatial smoothness modulated by the spatial edge terms and the cost terms associated with the edge functional w and its gradient are needed for more accurate estimation of the bolus front inside the vasculature. Although, due to the presence of a prior map, its effect on vasculature segmentation is not as important. In the presence of inaccuracies in the map, however, this term will contribute to more effective vasculature segmentation.

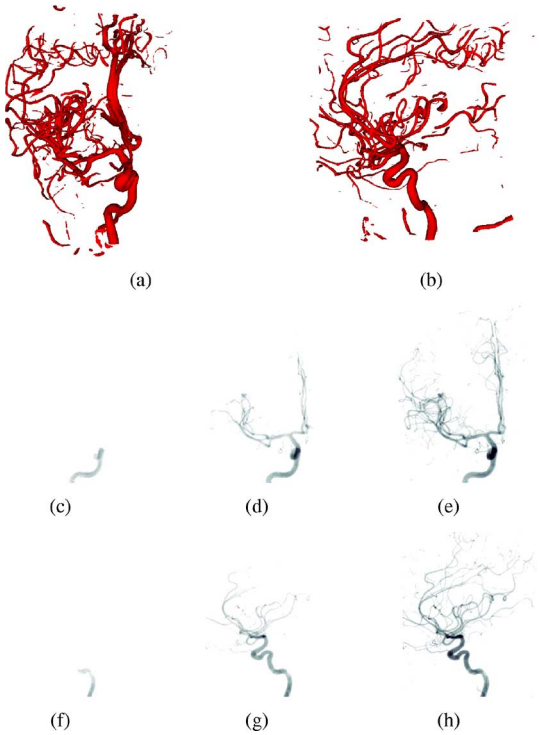


Fig. 3. Raw data for patient A—The first row shows a visualization of the 3D map generated from patient A from two views. The second and third rows show the measured 2D projections at different times for the antero–posterior and lateral views, (a) and (b), respectively. (c) Time 1. (d) Time 5. (e) Time 10. (f) Time 1. (g) Time 5. (h) Time 10.

The temporal smoothness term modulated by the time edge term and the cost terms associated with v and its time derivative are crucial. For instance, about four times as many pixels were wrongly assigned in the absence of these terms.

B. Validation on Angiographic Data

1) 3D+T Time Series Reconstruction: Patient A: The top row of Fig. 3 provides a visualization of the vasculature map of patient A that is used as an input to the spatio-temporal data fusion algorithm. The map is accurate for the larger vessels, but contains noticeable errors in the smaller vessels. In addition, a sampling of the original 2D image sequences of patient A are shown from the antero–posterior view (middle row) and from the lateral view (bottom row). The algorithm fuses all the information contained in Fig. 3 to provide the 3D+T time history of Fig. 4.

The top two rows of Fig. 4 shows a visualization of the reconstruction of a 3D time series of real angiographic data at six of nine different times. The figure shows the result of rendering a mesh at an isosurface based on a particular threshold. These panels show the outward flow of blood as time progresses. This flow is consistent with the type of blood flow that is expected and with what is observed in the 2D time series. Another way of visualizing the 3D flow is using the color coded arrival times, as shown in the third row of Fig. 4. The color coded arrival times show the advancing front of the contrast. This allows the 3D flow of the blood to be accessed in a single manipulatable surface from a multitude of views.

Fig. 5 illustrates the impact of varying the visualized iso-level of the reconstruction at time 8. As the threshold is lowered more of the blood vessel can be seen, but at the same time more artifacts appear. Because of this trade off an appropriate threshold needs to be carefully chosen. The figure also gives a visualization of where the large values are in the reconstruction. This appearance is consistent with the reconstructions within the 2D images in Section IV.

Fig. 6 shows a few of the reconstructed (right panels) 2D image sequences of patient A from the antero–posterior (top row) and from the lateral (middle row). These reconstructions can be compared to Fig. 3. As hoped, the reconstructed image sequences are visually similar to the original measured image sequences. This similarity was determined by Dr. Malek, an experienced Neurosurgeon. The few differences between the original and reconstructed time series are due to vessels not appearing within the 3D map in Fig. 3, the small delay between the acquisition times of the image pair (the antero–posterior view had a small delay from the lateral view), and the regularization terms used in reconstruction.

In addition to the reconstruction of 2D series from observed views, the synthetic projections or DRRs can be produced for the same blood flow instances from unobserved views as seen in the bottom row of Fig. 6. This process allows surgeons to find the best view for observing an area of interest such as one containing a clot or an aneurysm. This capability allows angiographic time sequences taken from different views to be compared directly. It also enables the generation of projections from unobservable views, such as from the top or the bottom of the patients head. In addition, this process reduces the X-ray exposure dose a patient receives by replacing actual X-rays with synthetic X-Rays.

2) Blood Flow Anomaly Detection: Patient B: This case highlights the ability of the proposed algorithm to visualize changes of blood flow that may develop during procedures. Fig. 7 shows the reconstructed data sets for patient B that were visualized from different views. The panels on the left are for flow in the presence of the clot while the panels on the right are for the normal flow after the clot is removed. The arrival times of the contrasted blood flow downstream from the blocked artery are delayed by 1–3 frames from what they were in the no clot case. The visualizations shown in Fig. 7 highlight the presence of the clot and also confirm the successful removal of the clot.

These visualization allow for straightforward comparison and can alert the surgeon to the presence of a problem that can be further assessed and acted on if necessary, without the need to compare the multitude of frames from multiple 2D X-ray time sequences. More importantly, it may help detect subtle changes in regional blood flow that may escape detection by the visual comparison method currently employed by physician operators.

VI. Discussion

A. Variations on Algorithm

In addition to the forward Gauss–Seidel implementation, the algorithm can be run in the reverse direction (starting at the last time and moving backward in time) or with alternating forward

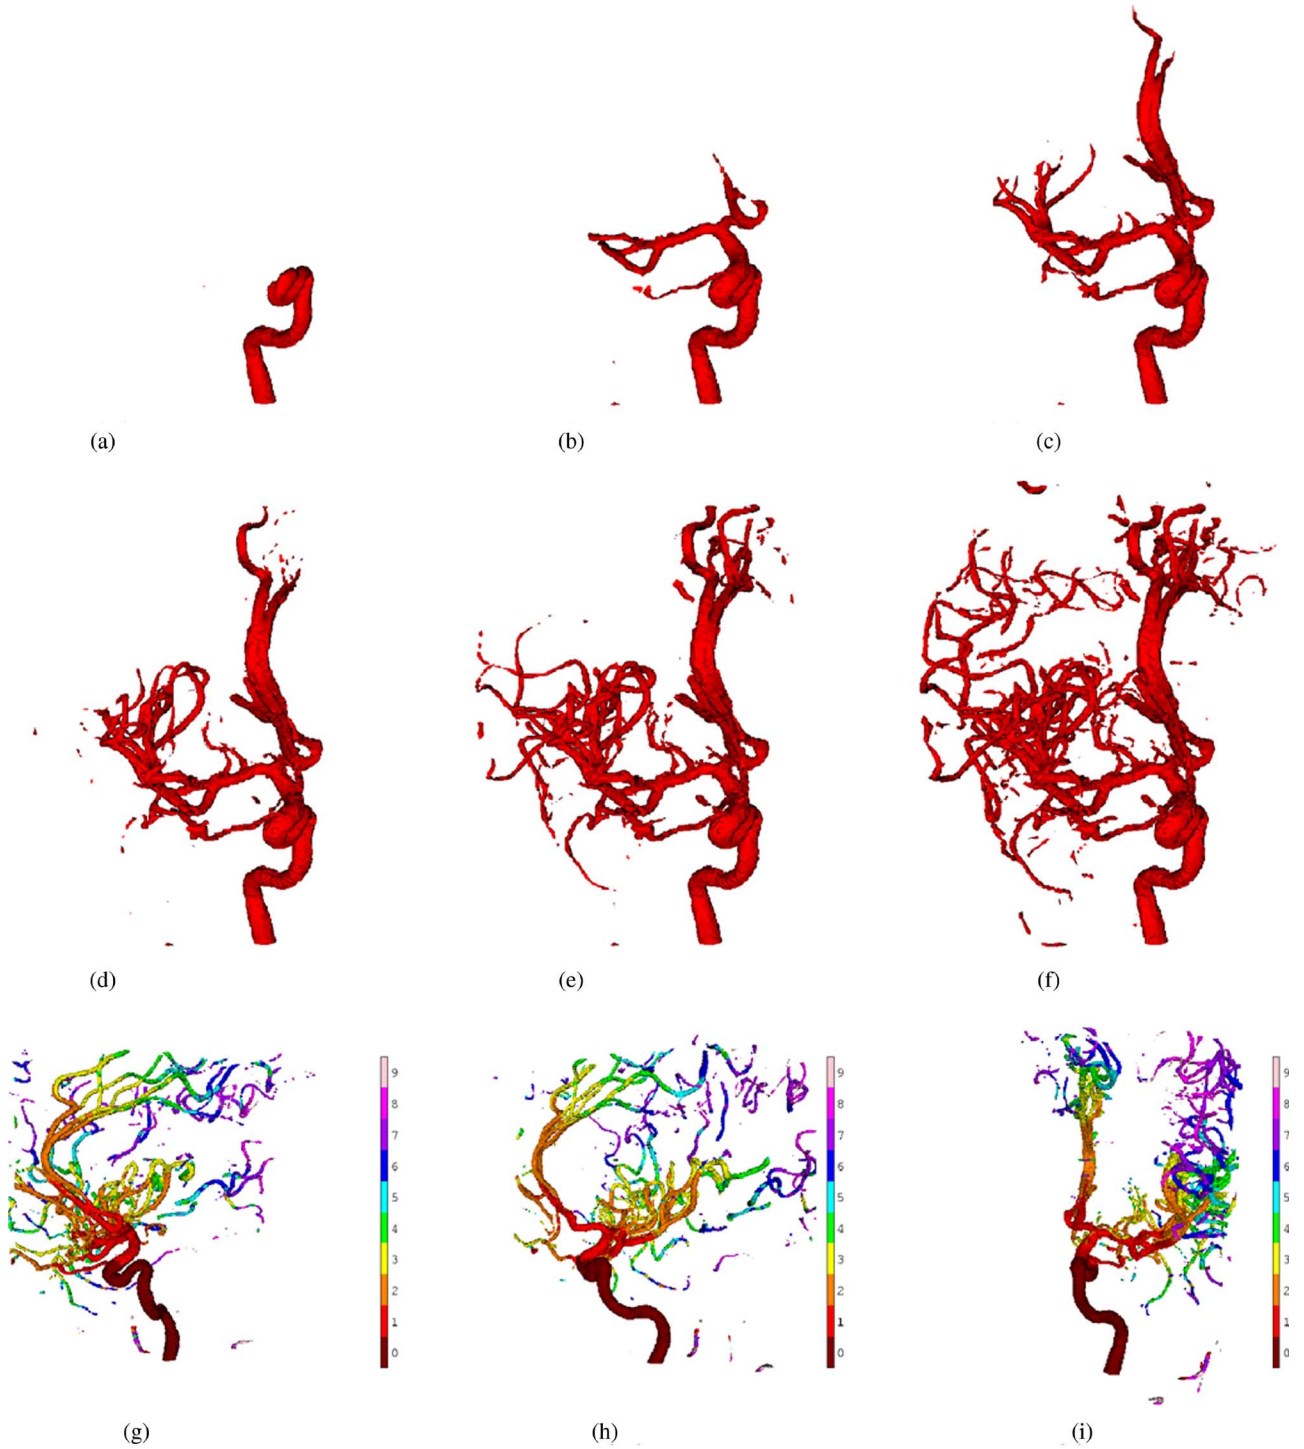


Fig. 4. Visualization of reconstructed 3D time series for patient A. The top two rows in (a)–(f) show visualizations of the reconstructions at six of the nine time steps. The total observed sequence took one second. The threshold of 0.01 was used to generate the isosurfaces. While the bottom row shows in (g)–(i) a 3D time series using color to denote the arrival time of contrast at the different points on the surface of the mesh. This visualization is shown for three separate views and are sampled at eight hertz.

and reverse sweeps. Because each volume is linked to its immediate neighbor, another interesting choice is to calculate every other volume in a single sweep (e.g., the odd times) and then go back to calculate the volumes at the remaining times (e.g., the even times) in a second sweep. This type of algorithm is similar to a red-black Gauss–Seidel [40] algorithm, where every other item is labeled red and each of its neighbors are labeled black. Such an algorithm consists of "red" iterations followed by "black" iterations. Determining each volume of a given color is independent of all other volumes of that same color. An algorithm of this form is parallelizable, leading to significant gains in computational performance. Restriction of the reconstruction to the map allows for further optimizations that yield improved speed and reduced memory constraints. The resultant algorithm provides a memory efficient and stable method for reconstructing the blood flow.

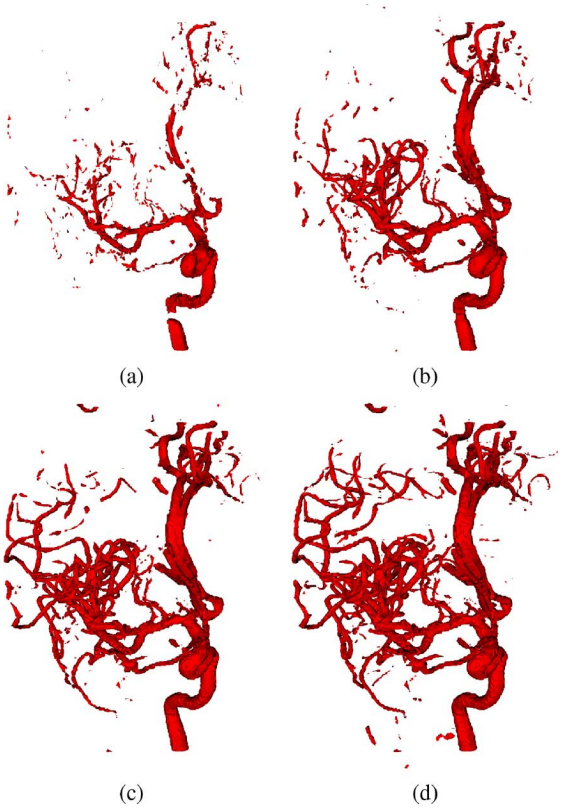


Fig. 5. Effect of thresholds on visualization. Visualization of the reconstruction for patient A at time 8 using four separate thresholds. (a) $Threshold = 0.1$. (b) $Threshold = 0.05$. (c) $Threshold = 0.015$. (d) $Threshold = 0.005$.

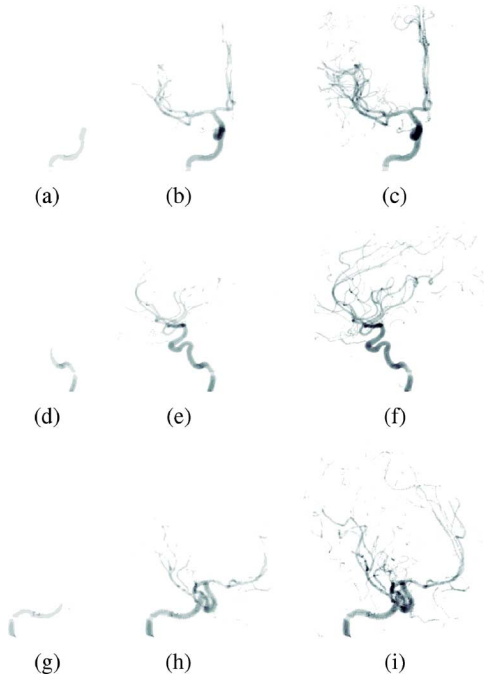


Fig. 6. Reconstructed time series projections of patient A. Digital reconstructed radiographs of the reconstructed 3D time series for patient A from the antero-posterior (top row), from the lateral view (middle row), and from a view that was not observed in the original angiographic sequences (bottom row) at three different times. Compare the first and second rows to the measured data shown in Fig. 3. (a) Time 1. (b) Time 5. (c) Time 10. (d) Time 1. (e) Time 5. (f) Time 10. (g) Time 1. (h) Time 5. (i) Time 10.

B. Implicit Fluid Flow Model

The temporal smoothness constraint in (3) provides an implicit fluid flow model. This model smoothes the intensities through time off of a set of edges or discontinuities in the flow. The set of discontinuities is penalized to be smoothly varying in space and time. A higher fidelity fluid flow model would capture a more precise numerical model for the flow. The temporal flow of the contrast can be captured by the temporal derivative of the intensities or $\nabla_t F$. If $\nabla_t F$ satisfies

$$\frac{\partial}{\partial t}(\nabla_t F) = \Psi(\nabla_t F), \tag{21}$$

where Ψ enforces the Navier Stokes equation, then $\nabla_t F$ is a valid flow. For such formulations to be practical, an appropriately chosen operator Φ that approximates Ψ can instead be used. A possible functional for incorporating models of this form is shown in (22) at the bottom of the next page where ζ is a penalty term that penalizes deviations from the flow model. Note that, for simplicity, we did not include a temporal edge function v as is used in (3). If, however, we wish to capture abrupt changes within our model, the edge function can be included. Finally, usage of such a model would help resolve potential flow ambiguities in more complex settings.

C. Algorithm Run Time

The current algorithm runs in approximately six hours on a single core of an Intel Xeon 5160. The structure of the proposed algorithm contains multiple parts that can be computed in parallel and that rely on data that are located in 2D and 3D arrays. These types of algorithms map particularly well to Graphical Processing Units. Typical speedups of such algorithms in practice are on the order of 50 ×, [41]. Furthermore, a red black version of the algorithm can easily be spread across multiple GPUs. Workstations containing 4 GPUs are readily available at the writing of the paper. Using such a system and code adapted to NVIDIA's Compute Unified Device Architecture (CUDA) programing language [42], execution times on the order of a few minutes are likely. An example from a similar application is shown in [43].

$$E(F,w) = \int_{\mathfrak{M}\times[0,t_{\text{end}}]} \left(\alpha \boldsymbol{\mathcal{B}}(\boldsymbol{\mathcal{A}}F - G)^2 + \beta(1-w)^2 \|\nabla_x F\|^2 + \gamma|F| + \zeta \left\| \frac{\partial}{\partial t}(\nabla_t F) - \Phi(\nabla_t F) \right\|^2 + \frac{\rho}{2}\|\nabla_x w\|^2 + \frac{w^2}{2\rho}\right) d\mathbf{x}dt \tag{22}$$

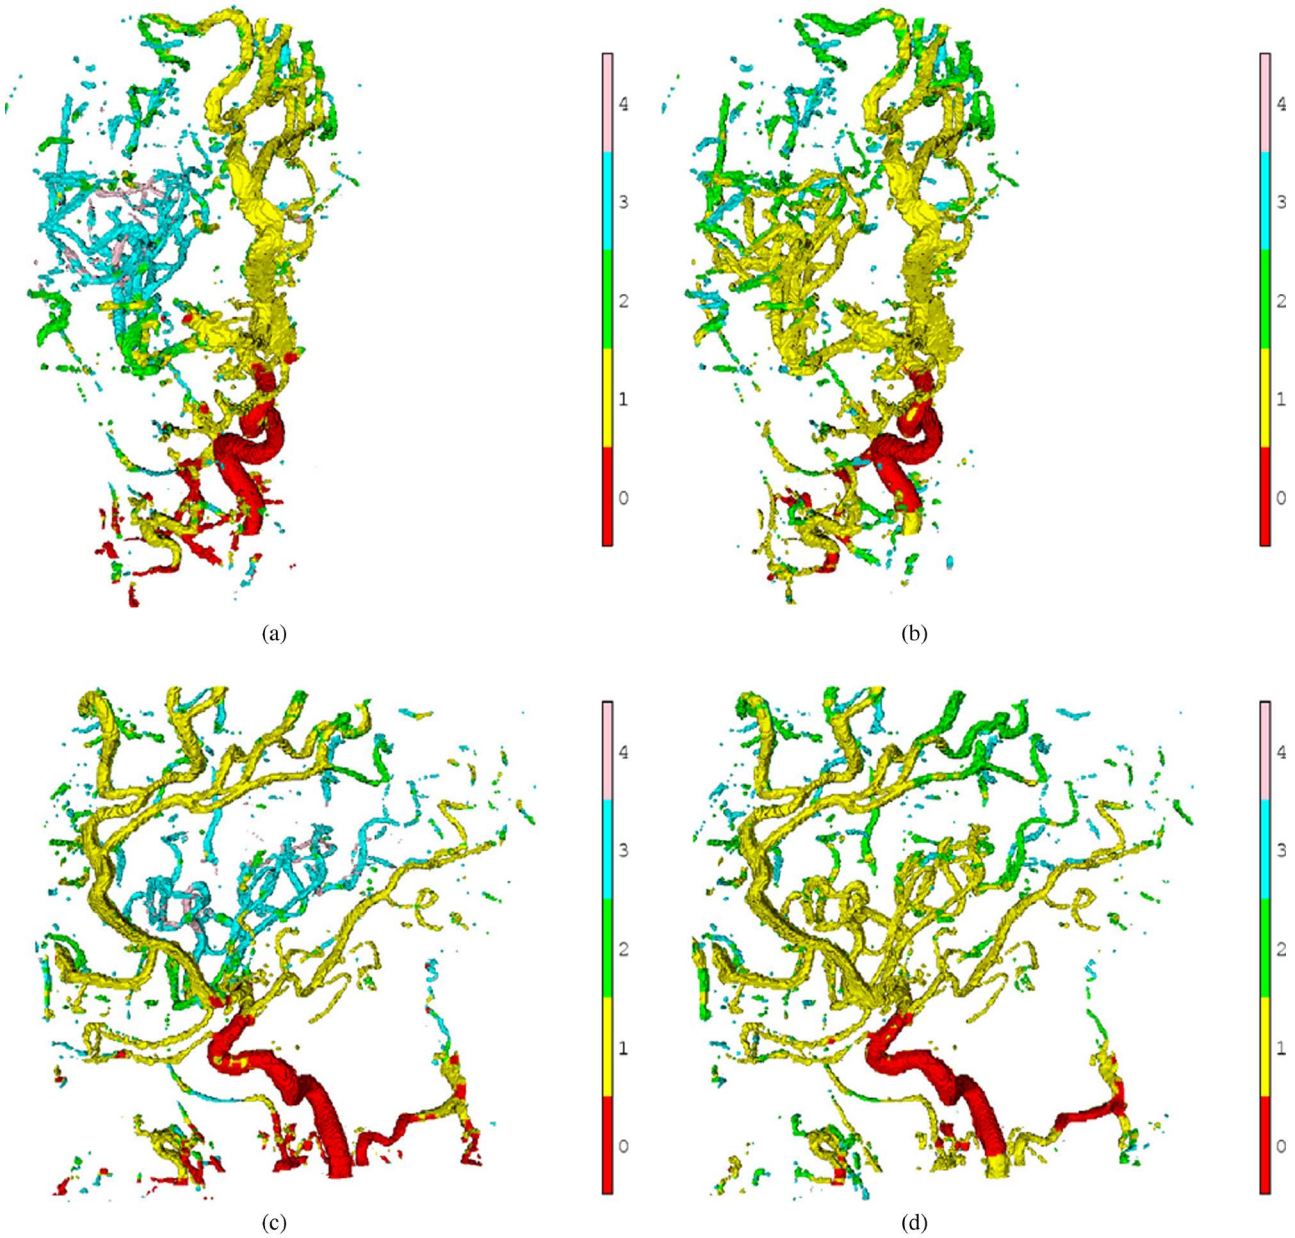


Fig. 7. With clot and post-clot removal 3D arrival times for patient B—Visualization of the reconstructed time series of patient B before (left) and after (right) a clot is removed from the middle cerebral artery in two views. Color is used to show contrast arrival time. The comparison of the arrival times images (left with the right) verified the removal of the clot. The blocked artery delays the down stream contrast arrival times by 1–3 frames (sampled at 3 Hertz) when compared to the baseline case. (a) Clot. (b) No Clot. (c) Clot. (d) No Clot.

D. Generation of Constraint Map

The map that was generated in this work provided good results, but it was clear that if a better map were available the results could be improved. A better segmentation algorithm such as [6] would allow finer vasculature structures to be resolved thereby leading to a more accurate reconstruction. Additionally, care should be taken to make sure that the entire region of contrasted vessels appear in the volume used to generate the constraint map. Other imaging modalities such as MRA or CTA can provide constraint maps. Constraint maps from these modalities can be enhanced through application of algorithms such as those in [44].

The primary purpose of the constraint map is to resolve ambiguities, while the $L1$ term is used to make it sparse. Consequently the constraint map can be enlarged to include all regions that do not reintroduce ambiguities resolved by the original constraint map. This can be found by comparing volume generated from the image pair at the time when the vessels are fully contrasted using the constraint map to those without the constraint map. Such an approach may compensate for deficiencies in the segmentation algorithm.

E. Extension to One or More Time Series

The use of more than two time series would lead to improved performance at the expense of additional exposure to the patient. Angiographic systems typically accommodate only two poses

simultaneously for the same blood flow history. To capture the additional time sequences new equipment and/or imaging techniques would need to be utilized. The temporal derivatives in the formulation can be modified and extended to support the use of image sequences from more than two poses if the relative timing of the different sequences is known. If the timing information is not known, a synchronization algorithm is needed.

In reconstructions from a single time series, the risk of ambiguities goes up, and operators will feel less comfortable using reconstructed information. These ambiguities most likely can be reduced by increasing the temporal frame rate, however, further study would be required to verify this.

VII. Conclusion

Angiographic imaging is essential in clinical management of patient-based lesions, both diagnostic and therapeutic. Extracting maximum information from raw images is paramount to safely performing medical procedures. This research has succeeded in extracting previously unavailable angiographic 3D+T histories from currently available surgical X-ray raw data without requiring additional image collection and consequent patient x-ray exposure. These time series also allow the generation of 2D projections from arbitrary views. The new images can improve the physician operator's procedure monitoring and allow quicker detection of intra-procedural life-threatening anomalies such as stroke-inducing blood clots.

The main contribution of this work is a variational formulation that achieves reconstructions by incorporating a 3D vasculature map and the 2D time sequences along with the use of sparsity and spatio-temporal constraints that approximate an implicit fluid flow model. Our implementation provides an efficient algorithm by exploiting the space-time structure of the problem.

This approach can be generalized to angiographic imaging in other vascular beds such as the heart and periphery and to reconstruction of higher dimensional (N greater than 4) phenomena from lower (less than N) dimensional observations.

Acknowledgment

The authors would like to thank the anonymous reviewers for their patient and thorough work, and for making excellent suggestions that improved the quality of this paper. This work was conducted while A. Copeland was completing the Ph.D. degree in Electrical Engineering and Computer Science at the Massachusetts Institute of Technology as a Draper Fellow.

References

[1] H. Schmitt, M. Grass, R. Suurmond, T. Kohler, V. Rasche, S. Hahnel, and S. Heiland, "Reconstruction of blood propagation in three-dimensional rotational x-ray angiography (3D-RA)," *Computerized Med. Imag. Graphics*, vol. 29, no. 7, pp. 507–520, Oct. 2005.

[2] S.-Y. J. Chen and J. D. Carroll, "Kinematic and deformation analysis of 4-D coronary arterial trees reconstructed from cine angiograms," *IEEE Trans. Med. Imag.*, vol. 22, no. 6, pp. 710–721, Jun. 2003.

[3] H. Schmitt, M. Grass, V. Rasche, O. Schramm, S. Hähnel, and K. Sartor, "An x-ray based method for the determination of the contrast agent propagation in 3D vessel structures," *IEEE Trans. Med. Imag.*, vol. 21, no. 3, pp. 251–262, Mar. 2002.

[4] I. Waechter, J. Bredno, R. Hermans, J. Weese, D. C. Barratt, and D. J. Hawkes, "Model-based blood flow quantification from rotational angiography," *Med. Image Anal.*, vol. 12, no. 5, pp. 586–602, Oct. 2008.

[5] D. Mumford and J. Shah, "Boundary detection by minimizing functionals," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 1985, pp. 22–26.

[6] A. Tsai, A. Yezzi Jr., W. Wells, III, C. Tempany, D. Tucker, A. Fan, W. E. Grimson, and A. Willsky, "Model based curve evolution technique for image segmentation," in *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2001, vol. 1, pp. 463–468.

[7] M. Desai, R. Mangoubi, J. Shah, W. Karl, H. Pien, A. Worth, and D. Kennedy, "Functional MRI activity characterization using response time shift estimates from curve evolution," *IEEE Trans. Med. Imag.*, vol. 21, no. 11, pp. 1402–1412, Nov. 2002.

[8] M. Desai, R. Mangoubi, and P. Sammak, "Noise adaptive matrix edge field analysis of small sized heterogeneous onion layered textures for characterizing human embryonic stem cell nuclei," in *Proc. IEEE Int. Symp. Biomed. Imag*, Boston, MA, Jun. 2009, pp. 1386–1389.

[9] M. Desai, D. Kennedy, R. Mangoubi, J. Shah, W. Karl, A. Worth, N. Makris, and H. Pien, "Model-based variational smoothing and segmentation for diffusion tensor imaging in the brain," *Neuroinformatics*, vol. 4, no. 3, pp. 217–233, Sep. 2006.

[10] J. S. Lim, *Two-Dimensional Signal and Image Processing*. Englewood Cliffs, NJ: Prentice Hall, 1999.

[11] P. Oskoui-Fard and H. Stark, "Tomographic image reconstruction using the theory of convex projections," *IEEE Trans. Med. Imag.*, vol. 7, pp. 45–58, Mar. 1988.

[12] A. D. Copeland, "Spatio-temporal data fusion in cerebral angiography," Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, 2007.

[13] H. Chan, A. C. Chung, S. C. Yu, and W. Wells, III, "2D-3D vascular registration between digital subtraction angiographic (DSA) and magnetic resonance angiographic (MRA) images," in *Proc. IEEE Int. Symp. Biomed. Imag.: Macro Nano*, Apr. 2004, pp. 708–711.

[14] J. Feldmar, G. Malandain, N. Ayache, S. Fernández-Vidal, E. Maurincomme, and Y. Trousset, "Matching 3D MR angiography data and 2D x-ray angiograms," in *Proc. 1st Joint Conf. Comput. Vis., Virtual Reality Robotics Medicine and Medial Robotics Computer-Assisted Surg. (CVRMed-MRCAS '97)*, London, U.K., 1997, pp. 129–138.

[15] M. Groher, F. Bender, R.-T. Hoffmann, and N. Navab, "Segmentation-driven 2D-3D registration for abdominal catheter interventions," in *Proc. MICCAI*, 2007, pp. 527–535.

[16] J. Hipwell, G. Penney, R. McLaughlin, K. Rhode, P. Summers, T. Cox, J. Byrne, J. Noble, and D. Hawkes, "Intensity-based 2-D – 3-D registration of cerebral angiograms," *IEEE Trans. Med. Imag.*, vol. 22, no. 11, pp. 1417–1426, Nov. 2003.

[17] D. Russakoff, T. Rohlfing, K. Mori, D. Rueckert, A. Ho, J. Adler, and C. Maurer, "Fast generation of digitally reconstructed radiographs using attenuation fields with application to 2D-3D image registration," *IEEE Trans. Med. Imag.*, vol. 24, no. 11, pp. 1441–1454, Nov. 2005.

[18] J. Jomier, E. Bullitt, M. H. V. Horn, C. Pathak, and S. R. Aylward, "3D/2D model-to-image registration applied to tips surgery," in *Proc. MICCAI*, 2006, pp. 662–669.

[19] J. Weese, R. Goecke, G. P. Penney, P. Desmedt, T. M. Buzug, and H. Schumann, K. M. Hanson, Ed., "Fast voxel-based 2D/3D registration algorithm using a volume rendering method based on the shear-warp factorization," in *Proc. SPIE, Med. Imag. 1999: Image Process.*, May 1999, vol. 3661, pp. 802–810.

[20] L. Zöllei, E. Grimson, A. Norbash, and W. Wells, III, "2D-3D rigid registration of X-ray fluoroscopy and CT images using mutual information and sparsely sampled histogram estimators," in *Proc. 2001 IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2001, vol. 2, pp. II–696–II–703.

[21] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," in *Proc. 14th Annu. Conf. Comput. Graphics Interactive Techniques (SIGGRAPH '87)*, New York, 1987, pp. 163–169.

[22] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Piscataway, NJ: IEEE Press, 1999.

[23] M. Garland and P. S. Heckbert, "Surface simplification using quadric error metrics," in *Computer Graphics*, 1997, vol. 31, pp. 209–216.

[24] Matlab Function Reference7th ed. Natick, MA, The Mathworks, 2006.

[25] D. Mumford and J. Shah, "Optimal approximation by piecewise smooth functions and associated variational problems," *Comm. Pure Appl. Math.*, vol. 17, pp. 577–685, 1989.

[26] L. Ambrosio and V. Tortorelli, "Approximation of functionals depending on jumps by elliptic functionals via Gamma convergence," *Commun. Pure Appl. Math.*, vol. 43, no. 8, pp. 999–1036, 1990.

[27] L. Ambrosio and V. Tortorelli, "On the approximation of free discontinuity problems," *Bolletino U.M.I.*, vol. 6B, pp. 105–123, 1992.

[28] T. J. Richardson, "Limit Theorems for a Variational Problem Arising in Computer Vision," in *Annali della Scuola Normale Superiore di Pisa*, 1992, vol. XIX, Pisa Serie IV. Classe di Scienze, pp. 1–49, 1.

[29] T. J. Richardson and S. K. Mitter, "A variational formulation-based edge focusing algorithm," *Sadhana*, vol. 22, no. 4, pp. 553–574, 1997.

[30] J. Kybic, T. Blu, and M. Unser, K. M. H. Milan Sonka, Ed., "Variational approach to tomographic reconstruction," in *Proc. SPIE*, San Diego, CA, Feb. 2001, vol. 4322, pp. 30–39.

[31] C. V. Alvino, "Multiscale Active Contour Methods in Computer Vision with Applications in Tomography," Ph.D., Georgia Institute of Technology, , 2005.

[32] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, pp. 489–509, Feb. 2006.

[33] M. Çetin and W. C. Karl, "Feature-enhanced synthetic aperture radar image formation based on nonquadratic regularization," *IEEE Trans. Image Process.*, vol. 10, no. 4, pp. 623–631, Apr. 2001.

[34] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans. Image Process.*, vol. 6, no. 2, pp. 298–311, Feb. 1997.

[35] M. Li, H. Yang, and H. Kudo, "An accurate iterative reconstruction algorithm for sparse objects: Application to 3D blood vessel reconstruction from a limited number of projections," *Phys. Med. Biol.*, vol. 47, no. 15, pp. 2599–2609, 2002.

[36] E. J. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.

[37] E. J. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, pp. 4203–4215, Dec. 2005.

[38] A. H. Andersen and A. C. Kak, "Simultaneous algebraic reconstruction technique (SART): A superior implementation of the art algorithm," *Ultrason. Imag.*, vol. 6, pp. 81–94, 1984.

[39] W. F. Ames, *Numerical Methods for Partial Differential Equations*, 2nd ed. New York: Academic, 1977, Computer Science and Applied Mathematics.

[40] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A Multigrid Tutorial*, 2nd ed. Philadelphia, PA: SIAM, 2000.

[41] P. Muyan-Ozcelik, J. D. Owens, J. Xia, and S. S. Samant, "Fast deformable registration on the GPU: A CUDA implementation of demons," in *Proc. 2008 Int. Conf. Computat. Sci. Its Applicat. (ICCSA '08)*, Washington, DC, 2008, pp. 223–233.

[42] N. Corporation, NVIDIA CUDA Programming Guide 2.1 NVIDIA, 2008.

[43] A. Khamenea, P. Blochb, W. Weina, M. Svatosc, and F. Sauer, "Automatic registration of portal images and volumetric CT for patient positioning in radiation therapy," *Med. Image Anal.*, vol. 10, no. 1, pp. 96–112, Feb. 2006.

[44] A. Gooya, H. Liao, K. Matsumiya, K. Masamune, and T. Dohi, "Effective statistical edge integration using a flux maximizing scheme for volumetric vascular segmentation in MRA," in *Proc. IPMI*, 2007, pp. 86–97.