

Structure Propagation for Image Registration

Mehmet Yigitsoy, *Member, IEEE*, Nassir Navab, *Member, IEEE*

Abstract—Mosaicing is a commonly used technique in many medical imaging applications where subimages are stitched together in order to obtain a larger field of view. However, stitching, which involves alignment or registration in overlapping regions, is often challenging when the information shared by subimages is absent or small. While it is not possible to perform an alignment without overlap using existing techniques, imaging artifacts such as distortions towards image boundaries present further complications during registration by decreasing the reliability of available information. Without taking these into consideration, a registration approach might violate the continuity and the smoothness of structures across subimages. In this paper, we propose a novel registration approach for the stitching of subimages in such challenging scenarios. By using a perceptual grouping approach, we extend subimages beyond their boundaries by propagating available structures in order to obtain structural maps in the extended regions. These maps are then used to establish correspondences between subimages when the shared information is absent, small or unreliable. Using our approach ensures the continuity and the smoothness of structures across subimage boundaries. Furthermore, since only structures are used, the proposed method can also be used for the stitching of multi-modal images. Our approach is unique in that it also enables contactless stitching. We demonstrate the effectiveness of the proposed method by performing several experiments on synthetic and medical images. Moreover, we show how stitching is possible in the presence of a physical gap between subimages.

Index Terms—Image Registration, Stitching, Structure Propagation, Tensor Voting.

I. INTRODUCTION

IMAGE mosaicing is the process of stitching two or more images together in order to obtain an extended field of view (FOV). It has many application areas ranging from computational photography, to computer vision, to medical imaging. Especially in the field of medical imaging, there has been an increasing interest in having larger FOV images. 3D histology reconstruction [1]–[4], extended FOV microscopic [5]–[9], laparoscopic [10] or endoscopic [11]–[13] images, 3D ultrasound mosaicing [14], [15], and whole-body MRI [16], [17] are exemplary applications where an extended view is created by stitching subimages or subvolumes. Having a large FOV is especially important in microscopic images since it for example enables studying several questions that arise in neuroscience [8]. A mosaicing is often needed since most of the existing microscopes have a limited view either in lateral (e.g. confocal, multi-photon) or axial (e.g. histology) dimensions.

M. Yigitsoy, and N. Navab are with the chair for Computer Aided Medical Procedures (CAMP), Department of Informatics, Technische Universität München, Germany (e-mail: yigitsoy@in.tum.de, navab@in.tum.de).

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

The main step of mosaicing is the stitching of subimages by estimating transformations that will bring them into a spatial alignment which is usually done by using an image registration algorithm. Classical image registration techniques usually assume minor differences in the viewpoints of images to be registered. However, this is not the case in microscopic imaging such as confocal microscopy where a large FOV is obtained by combining subimages [8] or in 3D digital pathology where single histology slices are stacked along the axial direction to create a 3D histology volume [3], [4]. Therefore, there are specific challenges of creating wide field of view images/volumes in these cases. Special registration techniques are needed which can handle 1) very small or distorted overlap between subimages or subvolumes¹ (as in the confocal microscopy case), 2) adjacent images having no overlap at all (maybe just touching each other as in 3D histology case).

In the conventional stitching, the amount of overlap between subimages to be stitched is important for finding correspondences for the estimation of necessary transformation parameters through image registration. Usually, a predefined value for the amount of overlap is used to decide whether a stitching between two subimages is feasible [2], [8], [18], [19]. Besides not being able to extract landmarks (if a landmark-based registration is used) or intensity correspondences, the “continuity and smoothness of structures” across subimages cannot be guaranteed unless a sufficient amount of overlap is present. These are the properties which are essential especially in microscopic imaging since the subsequent measurements on the structures will be influenced by the stitching performance.

To motivate the problem, let us assume that one of the initially aligned subimages in Figure 1, which are simply two pieces obtained by cutting a whole image into two, undergoes some misalignments such as translational or rotational transformations. It is very challenging to bring these two pieces into spatial alignment again even in the case of a simple translational misalignment without having an overlapping region. The problem becomes much more challenging and complicated to solve when there is also a physical gap between the pieces, in other words, when there is an occluded region between the pieces as illustrated in the last column of Figure 1.

Since an overlap cannot always be guaranteed, there have been approaches in the literature to address the overlap dependence. One approach is, assuming that there is a shared border between subimages, to perform contour matching after extracting contours at the image boundaries [20], [21]. Although this approach works on puzzle solving tasks, it fails when applied on medical images where boundaries usually have low signal

¹Terms “image” (2D) and “volume” (3D) will be used interchangeably throughout this paper.

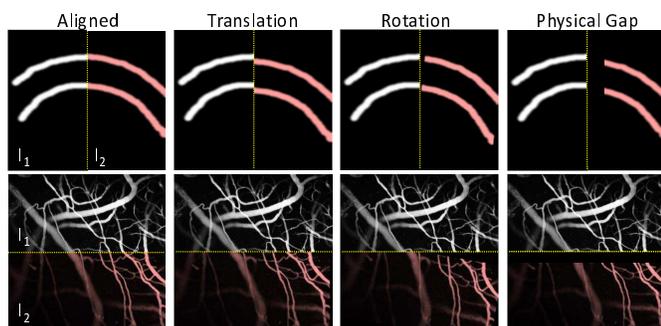


Fig. 1. A simple illustration of the possible challenging scenarios. Initially aligned subimage pairs on the left can be misaligned by translational or rotational transformations applied to one of the subimages. In this case, it is very difficult to restore the applied transformations without having an overlapping region to be used as a base for the alignment. It is much more challenging when there is an occluded region between the subimages.

or distortions. Another approach, which is often used for 3D histology reconstruction, is the use of slices at subvolume boundaries. Using boundary slices directly as in [2], [22] may lead to an incorrect reconstruction due to possible distortions in boundary slices [1]. Based on this observation, Bagci et al. [3] proposed to select “best reference slices” from each subvolume, which are then used for aligning subvolumes. Knowing that there may be large variations between even the neighboring histology slices, reference slices from subvolumes might be quite different from each other leading to possible misalignments.

All of the above mentioned techniques rely on the information contained either at the boundary or in a specific slice of a subvolume. *However, structures in medical data often have certain morphological smoothness and continuity properties which can be exploited to ensure that a final stitching is consistent with respect to the global continuity of structures across boundaries.* Following this idea, Lee & Bajcsy [1] proposed a feature-based approach for the volumetric reconstruction from the confocal laser scanning microscopic subvolumes. Feature trajectories extracted from vessel centerlines in subvolumes are used for fitting polynomial curves that are fused with a corresponding trajectory in a neighboring subvolume. However, they look for an affine transformation only in the lateral plane ignoring transformations in the stitching direction. Their method heavily depends on the extraction of centerlines via segmentation. Moreover, it is usually hard to find feature correspondences and segmentation is often an error prone task.

Based on these observations, we believe that there is still no promising solution to the problem of registering/stitching images/volumes in the absence of sufficient overlap. *Ensuring morphological continuity of structures has remained a challenging task in various applications in medical imaging including digital pathology and wide-field microscopy.* Therefore, in this work, we aim at proposing novel registration techniques addressing this challenging issue.

The main difference between our approach and the other existing solutions for the stitching of subimages is that a sufficient overlap region is not assumed to be *readily available* in the beginning. It is enough for our method if the subimages

are “sufficiently close” to each other. Furthermore, we assume that their relative positions are roughly known. That is to say, we know the adjacency relationship between the subimages in advance. These kind of challenges in medical image mosaicing such as the initial positioning of images by identifying adjacency relationships or error accumulations during mosaicing have been addressed in several recent works [7]–[9], [23], [24]. *Therefore, in this work, we will focus on developing strategies for the registration of subimages having limited or no overlap while ensuring the global continuity and the smoothness of structures crossing their boundaries.* It should be noted that, for a successful registration, it is important that image structures have inherent continuity, which is a property of the most natural images.

We are inspired mostly by the ability of the human visual system to perceive the “good continuity” of structures by integrating information from pieces. It has been reported in the literature that the human visual system can construct the whole from the pieces by integrating the continuity information of apparent contours within each piece, i.e. by perceptually grouping them [25] based on the Gestalt law of good continuation. Recently, there has been extensive research in computer vision to understand and model the perceptual grouping mechanism of the human visual system [25]–[28].

Several methods for modeling the perceptual grouping has been developed for curve and surface inference in computer vision [29]–[33]. Among these methods, “tensor voting” [32] has been applied for curve and surface inference in several applications [34]–[38]. It has been also used to solve many other computer vision problems such as image repairing [39], color correction [40] and terrain extraction [36]. More interestingly, several recent works in the field of medical imaging have successfully used the tensor voting in applications such as gap filling for vascular structures [41], catheter detection in fluoroscopic images [35] and detection of curvilinear structures in microscopic images [38].

In this work, we propose a novel technique which uses tensor voting for the inference of structures beyond image boundaries in order to establish a region shared by the subimages to be stitched as demonstrated in Figure2(a). To this end, we create structure maps by propagating salient structures from image regions into non-image regions. Then, these structure maps are used for the subsequent stitching of subimages. Finally, resulting transformations are transferred to the original subimages for the optimal alignment with respect to the global smoothness and continuity of structures across subimage boundaries. So far, we are not aware of the use of a conceptual grouping technique for the mosaicing of medical images.

There are several advantages of using such an approach;

- First of all, in the absence of an overlap region, these structure maps will serve as a basis for the estimation of necessary transformation that brings subimages into a spatial alignment.
- Secondly, if an existing overlap region is small or has severe geometrical distortions or deformations arising from slide preparations etc., then, this region can be supported by structural maps created by extrapolating the

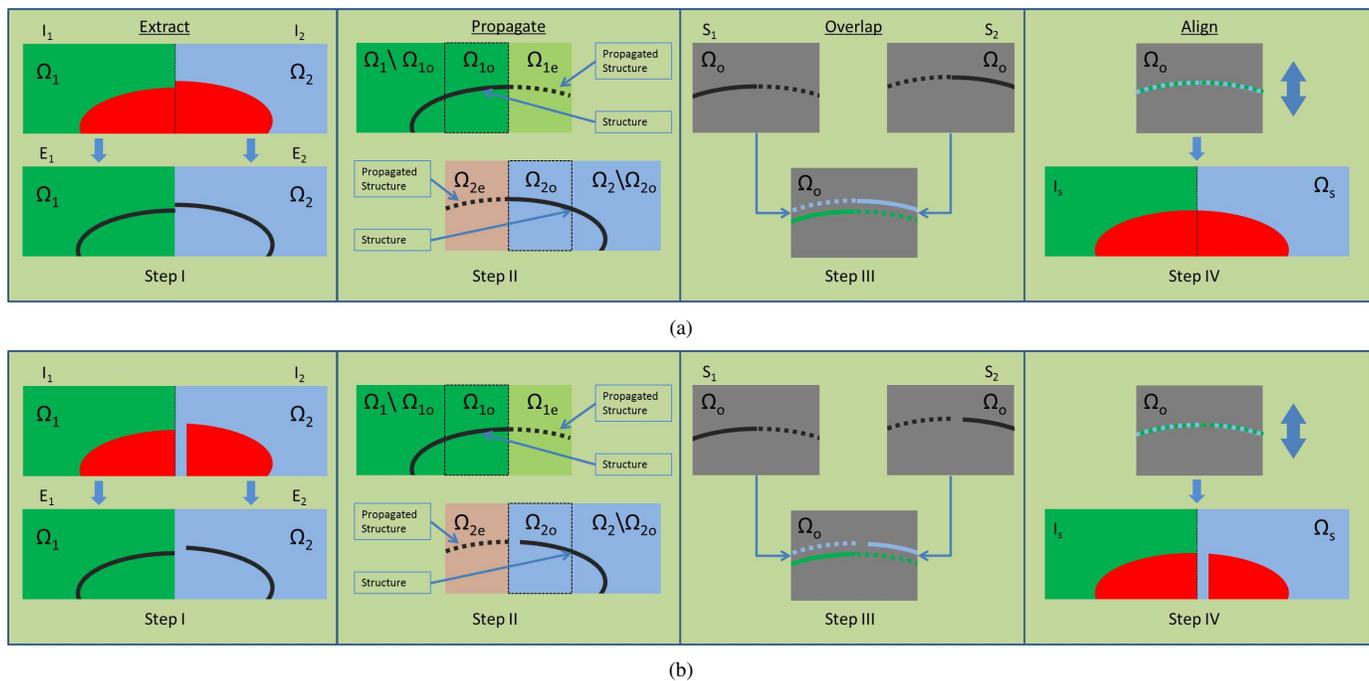


Fig. 2. (a) Overview of the proposed four-step stitching method using structure propagation. Here, we use a simple example where two pieces of a partial ellipse are to be stitched together. In Step I, the edge images, E_i are obtained by extracting structures, i.e. the edges, from the input images, I_i . This is followed by the propagation of structures into a predefined region Ω_{ie} for E_i in Step II. We create structure images denoted by S_i in Ω_o , which is constructed by combining the extended, Ω_{ie} , and the overlapped, Ω_{io} , regions of the edge image E_i in Step III. Finally, these newly created images are registered to each other followed by the application of resulting transformations to the original images, I_i in Step IV. (b) The same technique can be directly applied to the stitching of images having a physical gap (see missing part in I_2). The proposed technique is capable of aligning image even in this case thanks to the fact that the structures are extrapolated by making sure that the global continuity of structures is not violated.

structures that are salient or that present with no or less distortion into this region.

- Finally, an image stitched together from subimages using structure maps will have smooth connections between the corresponding structures of subimages.

Once structure maps are created, then, they can be registered to each other in order to find the necessary transformation. Transformations can be linear or non-linear depending on the needs of applications. Here, we will only consider linear cases for which we employ a registration method based on Markov Random Fields (MRF) as proposed in [42].

The organization of this paper is as follows. In Section II we explain our method where we give a brief theory of tensor voting in Section II-B and describe how it is used for structure propagation in Section II-C. Alignment of the propagated structures is discussed in Section II-D. Experimental results demonstrating the performance and the effectiveness of the proposed method on synthetic and real medical images are presented in Section III. Sections IV and V conclude the paper after some discussions.

II. METHOD

The main idea behind our approach is the extension of subimages in order to allow a stitching based on an artificial overlap created in the extended regions. The extension is in terms of image structures such as strong edges in the subimages. The structures are extrapolated beyond subimage boundaries. By posing this as an inference problem, we

employ tensor voting method to infer structures beyond image boundaries based on the orientation and saliency properties of structures within subimage regions. Inferred structures are then used in estimating transformations at different scales for the optimal alignment between subimages.

In this section, we first give an overview of our approach followed by a background on tensor voting and then we explain how structures are propagated using tensor voting. The employed stitching strategy is discussed at the end of this section.

A. Overview

Let $I_1 : \Omega_1 \subset \mathbb{R}^N \rightarrow \mathbb{R}$ and $I_2 : \Omega_2 \subset \mathbb{R}^N \rightarrow \mathbb{R}$ be two images to be stitched together. In this paper, we use the strong edges as structural information to be propagated. Therefore, we further define $E_1 = \zeta(I_1) : \Omega_1 \subset \mathbb{R}^N \rightarrow \mathbb{R}$ and $E_2 = \zeta(I_2) : \Omega_2 \subset \mathbb{R}^N \rightarrow \mathbb{R}$ as the structural representations where $\zeta(\cdot)$ is an edge operator. Edges can be detected by using an efficient implementation of edge detection with recursive filtering [43]².

In order to do a stitching between two subimages, one has to have a rough initial idea about their positions relative to each other. This prior information can be obtained in several ways such as using a step motor in microscopic imaging or using other existing methods as mentioned in Section I. In this work, we will assume that such prior information is available.

²Implementation available from <ftp://ftp-sop.inria.fr/epidaur/Softs/Malandain/>

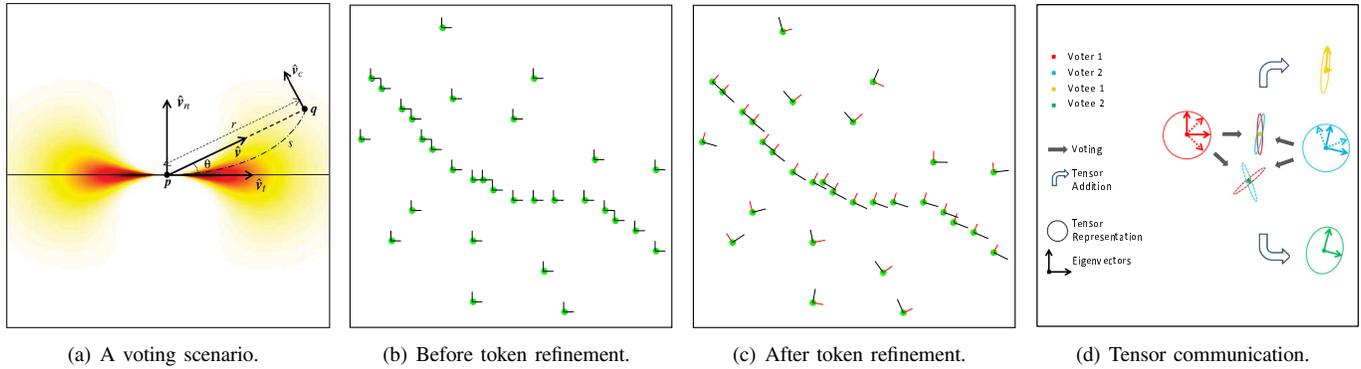


Fig. 3. (a) A voting scenario in 2D from a voter at p to a votee at q along an hypothesized curve s . A weighting function (color) determines the strength of a vote based on the distance between the voter and the votee. Here only a vote for a 1-structure, i.e. curve, which has information regarding only its normal space, \mathcal{N}_1 , is sent. (b)-(c) A simple illustration of the token refinement using a set of unorganized input tokens. Green dots represent the token positions while sticks stand for the eigenvectors of associated tensors. (b) Tokens are assigned ball tensors initially. They have no preferred orientations. (c) After a refinement step, they agreed on a certain orientation by exchanging information with each other in a local neighborhood. Note how tokens lying on a (possible) structure have similar orientation preferences. (d) An illustration of communication between ball tensors. Note how the normal space of a voter (dashed vectors) is oriented for each votee and how different accumulated tensor votes depend on the location of votees. Votee 1 received more stick votes in a certain orientation than votee 2 since it is more aligned with the voters than votee 2.

In other words, we will assume that we know the adjacency relationship between the subimages. Furthermore, without loss of generality, we will assume that two subimages to be stitched together are touching each other as illustrated in Figure 2(a). Other cases where there is a gap between subimages will be discussed later. Note that the goal here is to find the necessary transformation parameters that will bring these two subimages into spatial alignment such that the structures in the subimages are smooth and continuous across the boundary between them.

To this end, we define regions Ω_{1e} and Ω_{2e} corresponding to the extensions of image regions Ω_1 and Ω_2 , respectively. A scalar structure saliency map in the extended region Ω_{ie} will be created using tensor voting from the structures present in the edge image E_i . Ω_{ie} is now assumed to be overlapping with a small, if not entire, portion of region of the other subimage, Ω_j ($j \neq i$ and $i, j \in \{1, 2\}$). This region of overlap will be denoted by Ω_{jo} , i.e. the overlapping portion of Ω_j . Further we define

$$\Omega_o = \Omega_{io} \cup \Omega_{ie}, \quad i \in \{1, 2\} \quad (1)$$

which is simply the union of the overlapped and extended regions for the edge image E_i . We denote the corresponding images with $S_i : \Omega_o \subset \mathbb{R}^N \rightarrow \mathbb{R}$, i.e. the structure images that are constructed by combining the extended, Ω_{ie} , and the overlapped, Ω_{io} , regions of the edge image E_i . This is illustrated graphically in Figure 2(a). In this way, we create an overlap region between subimages by means of a two-sided structure propagation. For the optimization of transformation parameters to align the original images I_1 and I_2 , only the information present in the established overlap region Ω_o will be used. Finally, the stitched image will be denoted by $I_s : \Omega_s \subset \mathbb{R}^N \rightarrow \mathbb{R}$ where Ω_s is the union of the domains of two subimages as shown in the fourth step of Figure 2(a).

B. Tensor Voting

The Tensor Voting is a conceptual grouping method designed for the inference of salient structures from a set of incoherent input points [44], [45]. Inference is based on a

communication scheme where every point, *voter*, casts its information encoded as a second order symmetric tensor, \mathbf{T} , to other points, *votee*, over a hypothesized smooth curve with a low total curvature. The strength of the vote cast depends on the voter's perceptual saliency, the voter-to-votee distance and the curvature of the assumed curve connecting them. In the following two subsections, an overview of structure representation and tensor communication will be given. For more details about the tensor voting, we refer the reader to [36].

1) *Structure Representation*: Every structure type is identified by the dimensionality d of its normal space \mathcal{N}_d and its associated saliency s_d [36]. For instance, in 2D a point has a ball-shaped tensor ($d = 2$, i.e. the normal space has the dimensionality of 2), whereas a curve has a stick-shaped tensor ($d = 1$, i.e. the normal space has the dimensionality of 1) where the dominant component represents the normal direction. The strength of a d -structure is related to the magnitude of its saliency s_d .

A tensor is represented as $\mathbf{T} = \sum_{d=1}^N \lambda_d \hat{e}_d \hat{e}_d^T$ and can be decomposed as

$$\mathbf{T} = \sum_{d=1}^{N-1} (\lambda_d - \lambda_{d+1}) \sum_{k=1}^d \hat{e}_k \hat{e}_k^T + \lambda_N \sum_{k=1}^N \hat{e}_k \hat{e}_k^T. \quad (2)$$

It can be further written in terms of normal spaces as

$$\mathbf{T} = \sum_{d=1}^N s_d \mathcal{N}_d \quad \text{with} \quad s_d = \begin{cases} \lambda_d - \lambda_{d+1} & , d < N \\ \lambda_N & , d = N \end{cases} \quad (3)$$

where $\lambda_1 \geq \dots \geq \lambda_N \geq 0$ and $\hat{e}_1 \dots \hat{e}_N$ are eigenvalues and eigenvectors of \mathbf{T} , respectively, s_d is the saliency, $\mathcal{N}_d = \sum_{k=1}^d \hat{e}_k \hat{e}_k^T$ is the d -dimensional normal space, and N is the dimensionality of the input space. This representation can be interpreted as the decomposition of the information to different structure types. Every possible structure type has a normal space and an associated saliency.

2) *Tensor Communication*: For each possible structure, a vote is sent by the voter, which are accumulated at the votee. The transferred information per d -structure is simply a vote tensor, \mathbf{A}_d , weighted by its saliency, s_d . A tensor vote cast by a voter, \mathbf{p} , to a votee, \mathbf{q} , denoted by $\mathbf{A}^{\mathbf{P}}(\mathbf{q})$, can be written as the sum of vote components for every d -structure as

$$\mathbf{A}^{\mathbf{P}}(\mathbf{q}) = \sum_{d=1}^N s_d^{\mathbf{P}} \mathbf{A}_d^{\mathbf{P}}(\mathbf{q}) \quad (4)$$

where $s_d^{\mathbf{P}}$ is the saliency of the d -structure type at \mathbf{p} and $\mathbf{A}_d^{\mathbf{P}}(\mathbf{q})$ is the vote component for d -structure at \mathbf{p} . Here, the voter simply sends a vote component for each d -structure weighted by its saliency. If the voter contains only a single structure type, then, only its saliency will be non-zero and only that information will have an influence at the votee.

A vote component for a d -structure is defined as

$$\mathbf{A}_d^{\mathbf{P}}(\mathbf{q}) = \sum_{j=1}^d \mathbf{S}_{d,j}^{\mathbf{P}}(\mathbf{q}) \quad \text{with} \quad \mathbf{S}_{d,j}^{\mathbf{P}}(\mathbf{q}) = w(r, \theta) \hat{\mathbf{v}}_{c,j} \hat{\mathbf{v}}_{c,j}^T \quad (5)$$

where $\mathbf{S}_{d,j}^{\mathbf{P}}(\mathbf{q})$ is a *stick voting field* for each basis vector of $\mathcal{N}_d^{\mathbf{P}}$. In plain words, a vote component for a certain d -structure is composed of contributions from the basis vectors of its normal space.

To give an intuition, if there is only a curve, then, there is only a 1-structure having \mathcal{N}_1 , i.e. only a normal vector. Then, $\mathbf{A}^{\mathbf{P}}(\mathbf{q})$, i.e. the information cast by the voter, is

$$\mathbf{A}^{\mathbf{P}}(\mathbf{q}) = s_1^{\mathbf{P}} \mathbf{A}_1^{\mathbf{P}}(\mathbf{q}) \quad \text{with} \quad \mathbf{A}_1^{\mathbf{P}}(\mathbf{q}) = \mathbf{S}_{1,1}^{\mathbf{P}}(\mathbf{q}) \quad (6)$$

where s_1 is just 1. This means only one stick voting field for a single basis vector, i.e. the normal vector, was used. If there would be only a 2-structure type, then, the vote component for this structure will be the sum of stick voting fields for each basis vector of the normal space \mathcal{N}_2 . The idea is straightforward, for each structure type, if the saliency is non-zero, then, we create a vote component by adding stick voting fields for each basis vector of the normal space. Now, we explain what a stick voting field means.

The stick voting field for the j^{th} basis vector, $\hat{\mathbf{v}}_{n,j}$, of \mathcal{N}_d is the product of a weighting term $w(r, \theta)$ with the outer product of the implied vector

$$\hat{\mathbf{v}}_{c,j} = \hat{\mathbf{v}}_{n,j} \cos(2\theta) - \hat{\mathbf{v}}_{t,j} \sin(2\theta) \quad (7)$$

where $\hat{\mathbf{v}}_{t,j}$ is a vector in the associated tangent space and θ is the vote angle defined as the angle between $\hat{\mathbf{v}}_{t,j}$ and the voter-to-votee vector, $\hat{\mathbf{v}}$, see Figure 3(a) for an illustration. The implied vector is simply the vector that we would obtain if the normal vector, $\hat{\mathbf{v}}_{n,j}$, is slid on an assumed arc as shown in Figure 3(a). The derivation of Equation 7 from the figure is straightforward. This simply gives the normal at \mathbf{q} if there would be an arc connecting \mathbf{p} and \mathbf{q} . Since a low total curvature of the assumed curve connecting the points is desired [44], the magnitude of the implied vector on arcs having larger curvature should be punished. Furthermore, the requirement that the strength of the vote should decrease as the distance between the points increases justifies the usage of

a weighting term based on the curvature of the assumed curve and the distance between the points.

The weighting term $w(r, \theta)$ essentially controls the strength of a vote depending on the angle, θ , and the distance between the points, $r = \|\mathbf{p} - \mathbf{q}\|$. This is done by setting

$$w(r, \theta) = e^{-\left(\frac{s^2 + c\kappa^2}{\sigma^2}\right)} \quad (8)$$

$$s = r\theta / \sin(\theta) \quad (9)$$

$$\kappa = 2 \sin(\theta) / r \quad (10)$$

as proposed in [32] where s is the arc-length, κ is the curvature of the curve, σ is the scale parameter controlling the voting distance and c is a parameter that can be tuned to change the compactness of the voting field. $w(r, \theta)$, as shown in color in Figure 3(a), is simply a scaling function for the implied vectors. The strength of the stick vote is larger if the curvature is zero and if the voter-to-votee distance is minimal.

Votes cast by a set of voters, \mathcal{P} , to a certain votee at \mathbf{q} are accumulated in the following way

$$\mathbf{T}(\mathbf{q}) = \sum_{\mathbf{P} \in \mathcal{P}} \mathbf{A}^{\mathbf{P}}(\mathbf{q}). \quad (11)$$

Then, a subsequent decomposition is applied to $\mathbf{T}(\mathbf{q})$ as described in Equations 2-3 to extract the saliences of different structure types.

C. Structure Propagation

In this section, we explain how the structures are propagated into the extended regions. The propagation is based on the orientation preferences of voters (see Section II-B). The voters in our case are basically the pixels/voxels on the structures to be propagated. Therefore, since the structures do not explicitly have their orientation preferences, we have to determine them before performing a propagation. A nice feature of tensor voting is that; given an unorganized set of tokens, one can obtain orientation preferences for each token. This is possible thanks to the mutual agreement of tokens on a preferred orientation via local communication. This process is explained further in the following section.

1) *Token Refinement* : We define a set of tokens, i.e. edge pixels, $\mathcal{P}_i \in \Omega_i$, for each edge image, E_i . Assuming that these tokens do not have any preferred orientation in the beginning, each of them is assigned a unit ball tensor as illustrated in Figure 3(b). Making this assumption helps eliminating the dependence on the orientations and edges as detected by the edge detector. Finally, tokens refine their information regarding their preferred direction of orientations by casting tensor votes to each other.

For instance, in 2D, without any preferred orientations, a voter will always vote in the favor of lines, i.e. stick votes, passing through a certain votee and itself. This practically means that the stick voting field in Figure 3(a) is always aligned with the voter-to-votee vector $\hat{\mathbf{v}}$. This is possible due to the fact that the basis vectors of the normal space of a unit ball tensor can always be chosen such that one of the basis vectors is along $\hat{\mathbf{v}}$ and others are orthogonal to it. At the end of this procedure, each token accumulates votes coming from

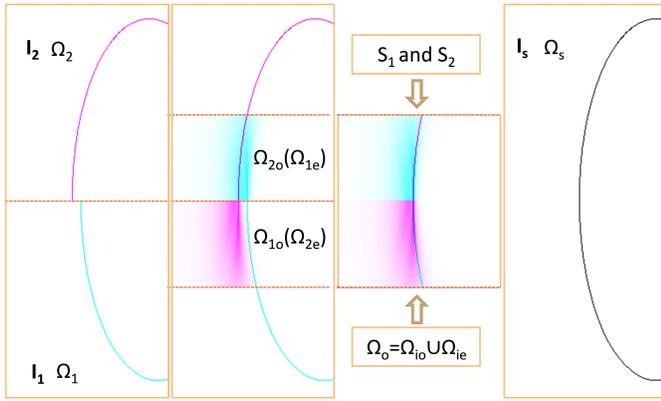


Fig. 4. An illustration of the stitching approach on a synthetic image pair. Only the structure images are shown here. Structures in one of the subimages are propagated towards the other one. Then, the subimages in the established overlapping region are used for estimating the necessary transformation parameters. Finally, resulting parameters are applied to the original images yielding a smooth and continuous stitching of structures at the stitching boundary.

other tokens in its neighborhood into a tensor which is then decomposed according to Equation 3. Those tokens which are close to each other or lie on a structure mutually agree on a preferred orientation and have higher saliences for a certain structure type as illustrated in Figures 3(c) and 3(d). In our 2D example, if a votee receives more line votes from voters along a certain orientation, then, after decomposition, this will be its preferred orientation, which is an indication of the presence of a line passing through this token with an associated saliency. At the end of this step, edge pixels will be assigned tensors encoding possible structure types and their saliences.

2) *Structure Propagation*: Having assigned tensors to edge pixels, i.e. tokens, encoding structure types and their associated saliences, every token can now propagate its information via tensor voting. In other words, based on their self information about the underlying structures, tokens vote for possible structures in other regions. We use this for the inference of structures in the extended regions, Ω_{ie} . Practically, each token in Ω_i sends a vote to every pixel in Ω_{ie} .

A dense voting is performed by

$$\mathbf{T}(\mathbf{q}_i) = \sum_{\mathbf{p}_i \in \mathcal{P}_i} \mathbf{A}^{\mathbf{p}_i}(\mathbf{q}_i), \quad i \in \{1, 2\}, \mathbf{p}_i \in \Omega_i, \mathbf{q}_i \in \Omega_{ie}. \quad (12)$$

where every pixel position \mathbf{q}_i in the extended region Ω_{ie} is considered as a votee. Once dense voting is finished, scalar structural (i.e. curve) saliences are extracted from the accumulated tensors as follows

$$S_i(\mathbf{q}_i) = \lambda_1(\mathbf{q}_i) - \lambda_2(\mathbf{q}_i) \quad i \in \{1, 2\}, \mathbf{q}_i \in \Omega_{ie} \quad (13)$$

where $\lambda_1(\mathbf{q}_i)$ and $\lambda_2(\mathbf{q}_i)$ are calculated using the eigendecomposition of $\mathbf{T}(\mathbf{q}_i)$ according to Equations 2 and 3. Note that only these extracted scalar saliences are used in the following steps including registration.

Next, we create the structure images S_i by setting intensity values from the overlapping region $\Omega_o = \Omega_{io} \cup \Omega_{ie}$. This is done simply by cropping the portion of the extended edge image E_i falling in Ω_o as shown in Figure 2(a). Now, we

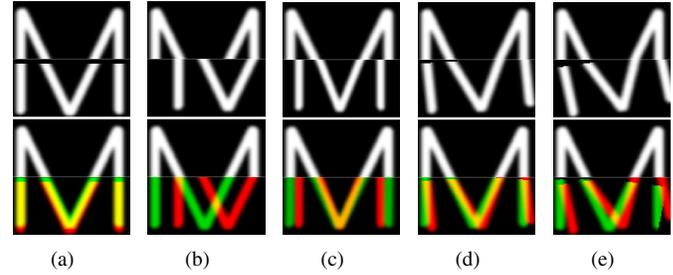


Fig. 5. Demonstration of image stitching without overlap. A synthetic image of “M” letter is cut into two and the lower piece is translated, scaled or rotated which is then restored using the proposed approach. In the first row, the initial stitchings before the alignment is shown. The gray horizontal line delineates the boundary between the pieces. In the second row, the final stitching is overlaid onto the initial one where the red and green colors show the lower piece of the letter before and after the alignment, respectively. In (a)-(b), translations in horizontal and vertical directions, respectively, are applied to the lower piece which are then successfully restored using the proposed stitching method. Note the physical gap in (a) created by applying a vertical translation. (c) demonstrates the ability to restore a change in scaling. In (d)-(e), a small and a relatively large rotation applied is recovered. Again, there is a gap as a result of the applied rotation which makes the stitching much more challenging than the case where there is no gap between the pieces to be stitched together.

have a pair of completely overlapping structure images, S_1 and S_2 . The observed structures in S_1 corresponds to the estimated structures in S_2 , and vice versa. Alternatively, one could also obtain a saliency map for the observed structures if the observed structures are too noisy or distorted. Here, for the sake of simplicity, we will assume that our observed structures are free of noise. However, in the experiments, we will show that the proposed method is already robust against noise.

In order to make sure that structures at different scales are treated appropriately, we employ a multiscale approach for both structure propagation and registration.

D. Stitching

In order to have a wide field of view, subimages have to be stitched together in a consistent way. This consistency is usually ensured by allowing an overlap region between subimages, which is then used for estimating the transformation parameters that will bring the subimages into a consistent alignment. In this work, we claim that for a successful alignment of subimages such an overlap does not need to be secured during acquisition. In the previous sections, we have demonstrated how such an overlap can be established in the absence of an existing one. In order to achieve a consistent alignment, a registration method is needed to recover the transformation parameters using the information in the established overlap region. We have turned the problem of stitching two non-overlapping images into the problem of registering two newly created images S_1 and S_2 . Transformations obtained from the successful registration of these images can be transferred to the original images I_1 and I_2 . See Figure 4 for an illustration.

Registration of S_1 and S_2 , which are assumed to be the source and the target images respectively, is posed as an optimization problem such that an optimal transformation \mathcal{T} optimizes an energy \mathcal{E} . Transformation \mathcal{T} is parametrized in terms of a set of n parameters $p = p_1, \dots, p_n$. Optimal

transformation parameters \hat{p} can be estimated via the following equation

$$\hat{p} = \arg \min_p \mathcal{E}(p). \quad (14)$$

where \mathcal{E} is defined in terms of a similarity measure as

$$\mathcal{E}(p) = \xi(S_1 \circ \mathcal{T}_p, S_2) \quad (15)$$

with \mathcal{T}_p being the parametrization of \mathcal{T} by p .

Misalignment between subimages can be global, which can be corrected by using an affine registration, but it can also be local where a non-rigid or deformable registration method is needed. In this work, we consider and address only the global misalignments. We employ a recently proposed intensity-based affine registration method [42] which is based on a discrete Markov Random Field (MRF) formulation. For the details about the used affine registration and discrete optimization methods, we refer the reader to [42], [46], [47].

Registration is performed in a multi-resolution setting to avoid local minima, where the optimal transformation is communicated between levels. Note that both structure propagation and stitching take place in a multi-resolution setting. Such an approach during structure propagation ensures that structures present at larger scales are considered first for the optimization whereas structures at a smaller scale are used for refining the transformation parameters. Although not implemented here, alternatively, one could also combine different scales of structure propagation and registration to make use of larger scale propagations also in finer scale registrations for regularization purposes.

In this work, after experimenting with several standard similarity measures for intensity based image registration, we ended up using the normalized cross-correlation (NCC) measure as the similarity measure $\xi(\cdot, \cdot)$ in Equation 15.

The implementation was done in C++. We integrated the tensor voting framework formulated in [36] with its implementation as a part of VXL³ library into the discrete MRF based image registration framework [42], [48].

III. EXPERIMENTS AND RESULTS

To evaluate the proposed technique, several experiments have been conducted on synthetic and real medical images. Synthetic images were used to demonstrate the capabilities of the proposed approach especially in cases, where not only there is no overlap but also there is a physical gap between the subimages to be stitched together as well as in cases where the images to be stitched are noisy. In order to evaluate the performance of the stitching of subimages in such cases without violating the smoothness and the continuity of structures across subimage borders, we have created several synthetic image pairs. In Section III-A, we show and interpret the results of synthetic experiments.

The proposed method has also been applied to the stitching of real medical images. To do this, we have extracted 2D slices from a two-photon microscopic image data set acquired from a rat brain. Again, as in the synthetic experiment case, slices were cut into two and several transformations constructed

in the same way were applied to one of the subimages. Furthermore, four microscopic image pairs with overlapping subimages to be stitched are used without applying any initial transformations. These experiments were important in that they showed the feasibility of the proposed method for being used for the stitching of real medical images.

For the evaluation of results, in addition to the visual assessment, we have compared the stitched images to the ground truth data using a correlation method. Furthermore, we evaluated the performance of the method in recovering the applied individual transformation parameters. Applied transformations are designed by varying one of the parameters at a time in a predefined range while keeping the others fixed to their initial values. In this way, we assess the sensitivity of our method to the variations in transformation parameters. In an ideal recovery of the transformation parameters, the sum of the applied and the recovered translational or rotational parameters should add up to zero. For the scaling parameters, the multiplication of the applied and the recovered parameter values should be equal to one. Shearing angle was not tested individually as it makes no difference without a non-identity scaling. Moreover, since a variation in a single parameter can also be compensated by variation in other parameters during optimization, we have also evaluated the recovery error for transformation matrices as a whole where all parameters (including shearing) have non-trivial values.

As an additional quantitative evaluation of the proposed method, we have developed a continuity index called *structure continuity index* (SCI), which is designed to evaluate the smoothness of structures across stitched subimages. We compute SCI again using tensor voting where structures extracted from one side of the stitching boundary *after* the alignment vote for the possible locations of structures on the other side. The real structures are then correlated with the estimated ones via tensor voting using the same notation in Section II. SCI can be formulated as

$$SCI = \sum_{\substack{i,j \in \{1,2\} \\ i \neq j}} \sum_{\mathbf{p} \in \mathcal{P}_{io}} \frac{S_i(\mathbf{p})S_j(\mathbf{p})}{|\mathcal{P}_{io}|} \quad (16)$$

where $\mathcal{P}_{io} \subset \mathcal{P}_i$ is the set of structure points in Ω_{io} . Note that this measure is in fact a cross correlation where only the detected structure points are used in the computation. This measure can also be used as a cost function during registration. Here, we use it only for evaluation in order to have an objective testing.

The robustness of the proposed approach against noise was also tested where varying degrees of white noise was added to a pair of synthetic images before applying any initial transformation. The noise added ranges between 0%-50% of image intensities.

There are several parameters to be set both for the tensor voting and the registration method. First of all, the size of the overlapping region Ω_{oi} to be established on each subimage in the stitching direction was set to be 20mm. This was considered to be sufficient for performing a registration. This can also be fine tuned depending on the image data. If the images to be stitched are well structured, then, a smaller value is also fine.

³VXL library is available for download at <http://vxl.sourceforge.net/>.

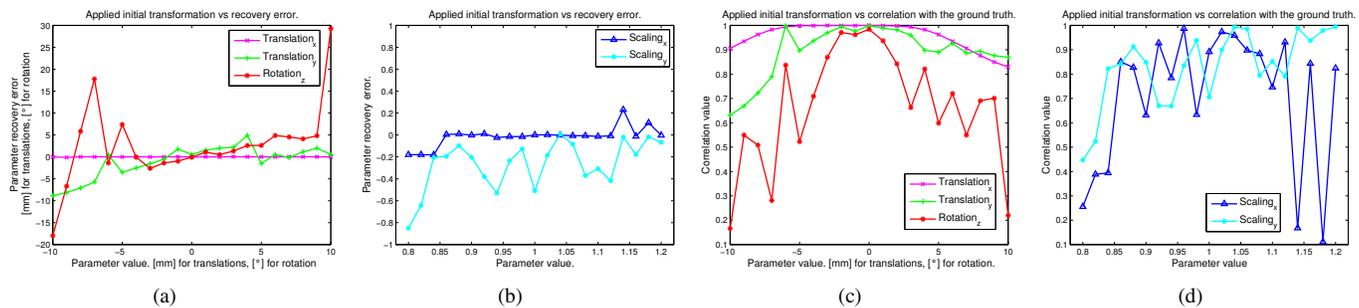


Fig. 6. Parameter recovery and ground truth correlation experiments. Several transformations, which differ from the identity only by the variation of one of the parameters in p , have been applied to the lower piece of the letter image shown in Figure 5. In each case, only one of the parameters changed its value in specific interval which was chosen to be ± 10 for translations (mm) and rotations (degrees) and 1 ± 0.20 for scalings. 21 equally spaced values for each parameter were sampled from the given interval. In each case, the error between the applied and the recovered transformation parameter was computed for (a) translation and rotation and (b) scaling. (c)-(d) Correlations of the aligned subimage with the ground truth subimage for (c) translation and rotation and (d) scaling.

In turn, a value of 40 was used for the scale parameter σ in Equation 8. This value guaranteed that a structure point 20mm inside a subimage can also vote for an image point 20mm inside of the other image although the influence decreases by the increasing distance. Also for the registration method, there are some parameters to be set. For defining the search interval for the affine transformation parameters, we set the search range to ± 40 mm for the translational and rotational components and to 1 ± 0.20 for the scaling components. The size of the discrete label space for each parameter is set to be 7 which is refined by a factor of 0.6 as suggested in [42] in every pyramid level. For all experiments, a PC equipped with Intel[®] Core[™] i7 2820QM, 230 GHz and 12GB RAM was used. It took about 2s to stitch an image of size 512x512 pixels with 1.6x1.6mm pixel size.

A. Synthetic Images

Synthetic experiments were conducted to demonstrate the capabilities of the proposed method for the stitching of subimages without overlap in different scenarios. In order to make sure that the method successfully copes with the lack of overlap region, we simply cut a full image into two pieces and created simple scenarios by applying affine transformations to one of the pieces. Given a full 2D image with smooth structures, it was cut into two subimages I_1 and I_2 , which are stitched again to obtain the original image I . Note that by simply dividing into two, we do not allow any overlap between the subimages. This is important because this is exactly the problem which we are trying to address in this work.

In Figure 5, we demonstrate how a stitching can be done in different scenarios. A synthetic image of “M” letter of size 102x124 pixels with 1mmx1mm pixel size was divided into two and the lower piece was translated, scaled or rotated which was then restored using the proposed approach. In the first row, the initial stitching before the alignment is shown. The gray horizontal line stands for the boundary between the subimages. In the second row, the final stitching is overlaid onto the initial one, where the red and green colors show the lower piece of the letter before and after the alignment, respectively. In Figure 5(a)-(b) translations in horizontal and vertical directions, respectively, are applied to the lower piece

which is successfully restored using the proposed stitching method. Note the physical gap in Figure 5(a) which is a result of the applied vertical translation. Figure 5(c) demonstrates the ability to restore a change in scaling. In Figure 5(d)-(e), a small and a relatively large rotation, respectively, are applied and recovered. Again, there is a gap as a result of the applied rotation which makes the stitching much more challenging than the case where there is no gap between the pieces to be stitched together. In all cases, it is visually obvious that a stitching is possible in different scenarios, where the common issue is the lack of image overlap.

For the quantitative evaluation of the proposed method, several transformations, in which only one of the parameters in p was varied, were applied to the lower piece of the letter image. In each case, only one of the parameters changed its value in specific interval which was chosen to be ± 10 for translations (mm) and rotations (degrees) and 1 ± 0.20 for scalings. Shearing parameter was not tested here as it makes no difference without a non-identity scaling. Instead, it was tested together with the other parameters while evaluating the transformation recovery performance which is explained later in this section. 21 equally spaced values were sampled from the give interval. For each case, the error between the applied and the recovered parameter was computed. For translation and rotation parameters, absolute value of the summation of the applied and the recovered parameter values are computed whereas for scaling deviation of their ratio from identity is computed. The error plots for each parameter in the specified interval are shown in Figures 6(a) and 6(b). We can make several observations looking at the figure. First of all, it is obvious that the most problematic parameter is the rotation. Although the parameter recovery error is reasonable in the vicinity of 0, i.e. small rotations, it gets quite large towards the marginal values. Another observation is that the recovery error for the translation parameter in the direction orthogonal to the stitching boundary is quite larger than its counterpart in the parallel direction, which is the same for the scaling. The reason for this is obvious since a physical gap is created in both cases, which makes the recovery process more error prone. The recovery error for the other parameters are almost zero, meaning that all the parameters are successfully restored

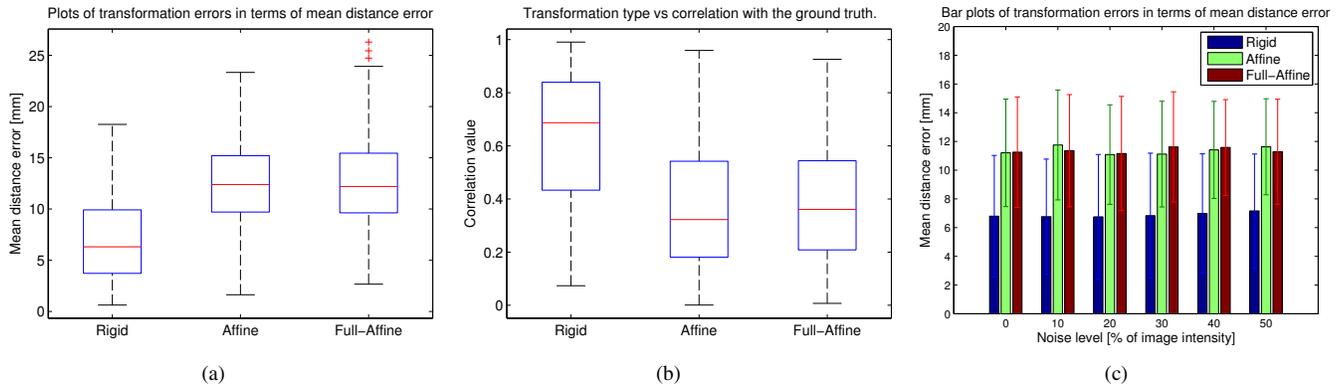


Fig. 7. Stitching experiments testing the sensitivity of the proposed technique to non-trivial transformations. The applied transformations are composed by randomly drawing transformation parameters from the following ranges; translations from ± 10 mm, rotation from $\pm 10^\circ$, scalings range from 1 ± 0.20 , and finally the shearing from $\pm 10^\circ$. Three cases were tested; only rigid parameters (rigid), including anisotropic scaling (affine) and finally including all six parameters (full-affine). In each case, 200 random initial transformations were composed of only relevant subset of parameters and only those parameters were optimized for. (a) Box plots show the transformation recover error statistics in each case. It is clear that the error increases as the number of degrees of freedom increases. Furthermore, adding shearing parameter did not make a significant difference in the overall performance. (b) Mean correlations of the aligned subimage with the ground truth subimage for different transformations. (c) Mean distance error in the presence of white noise. The added noise ranges from 0% to 50% of image intensities. The results present parallelism with (a) and (b) with a good performance in rigid case and a decreased performance for affine and full-affine transformations. Moreover, it is obvious that the noise did not affect the performance of the proposed method.

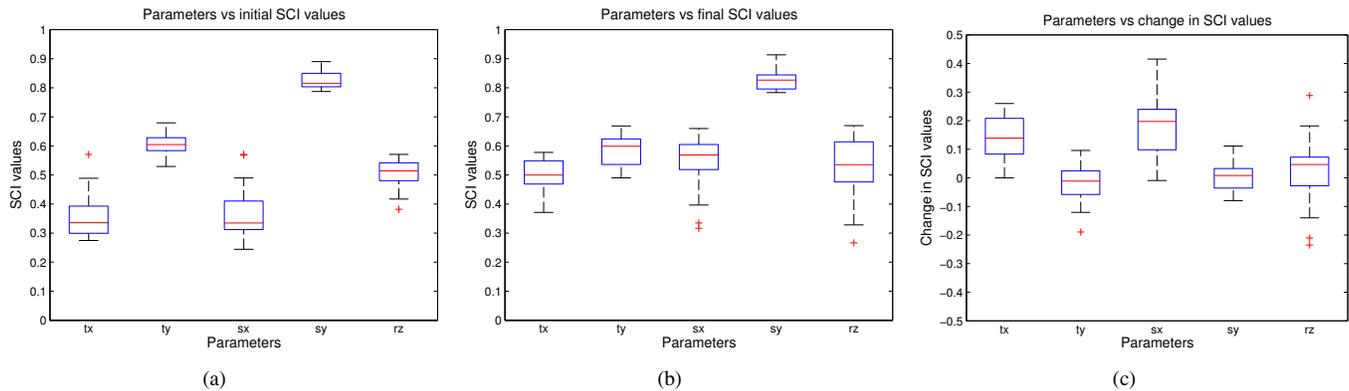


Fig. 8. Evaluation of the stitching performance in terms of the structural continuity index (SCI). SCI has been computed for each stitching task as shown in Fig 6. For each parameter there are 21 experiments where the parameter takes a value from an interval as discussed in Section III-A. Here, we show the box plots of SCI values for each parameter. (a) Initial SCI values before performing a stitching. (b) Final SCI values after stitching. (c) Change in SCI after stitching. Looking at (c), one can say that there has been a positive change for the translation and the scaling parameters in the direction parallel to the stitching boundary. On the other hand, there has been little or no improvement for their counterparts in orthogonal direction. Rotation parameter also presents with an increased SCI value.

even for the marginal parameter values.

We also compute the correlation of the aligned subimages with the ground truth to see whether the alignment was correct in terms of the image intensities. Figures 6(c) and 6(d) show the correlation results for each parameter for the specified value range. What is noteworthy is that although the correlation results for the rotation and the translation parameters are consistent with the parameter recovery error shown in Figure 6(a), this is not the case for the other parameters especially for the scaling parameter in the direction parallel to the stitching boundary, which show a better result in Figure 6(b). This indicates that the misalignment caused by the variation of this parameter has been mostly compensated by the variations in other parameters. This is quite possible as all of the parameters were optimized simultaneously for the best alignment. Furthermore, we take the middle of the stitching border as the origin of the image coordinate space.

For this reason, although a scaling in the direction orthogonal to the stitching boundary does not lead to a padding, a scaling in the parallel direction definitely results in paddings during registration. Therefore, a direct correlation of image intensities might be misleading in this case. Yet another observation that can be made from Figure 6 is that, in almost all cases, the alignment performance decreases towards the marginal values of the varied parameters. This is an expected behavior since the registration performance highly depends on the initial position.

In addition to the evaluation of individual parameter errors, we also compute the recovery error for transformation matrices as a whole. The error is computed as the mean distance of a set of points before applying an initial transformation and after restoring it. The point set consists of 200 points drawn randomly from a 200mm centered square. Ideally, a successful recovery should lead to zero mean distance error. The applied transformation are composed by drawing parameter

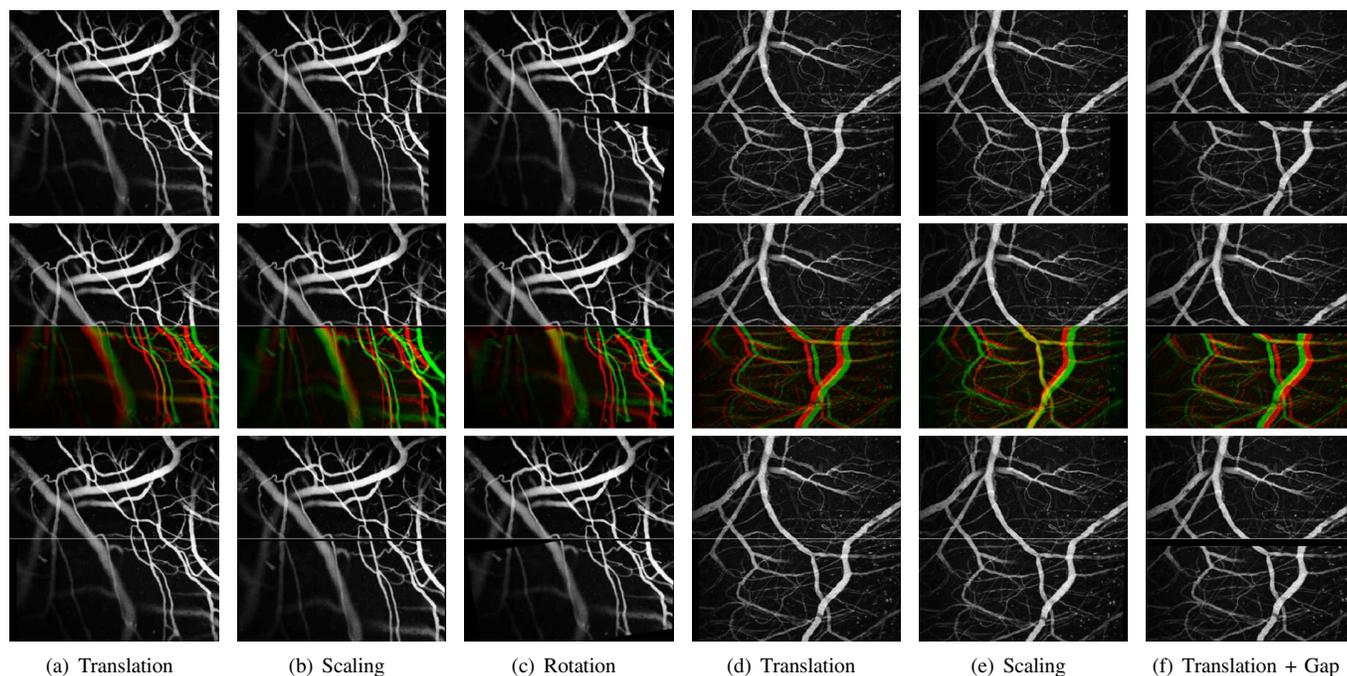


Fig. 9. Stitching experiments on 2D microscopic images. There are two microscopic image pairs P1 [(a)-(c)] and P2 [(d)-(f)] used for the experiments. The first row shows the initial stitchings obtained by applying affine transformations. In the second row, the final stitchings after the alignment is overlaid onto the initial one where red and green colors are used for initial and final versions of the aligned subimage. The last row shows the final stitchings without overlay. In (a)-(b), similarly in (d)-(e), misalignments caused by the variations in the translation and scaling parameters, respectively, are restored. In (c), the correction of a misalignment due to rotation is demonstrated. Finally, a misalignment caused by a translational transform in the presence of a physical gap between the subimages is recovered in (f). See the text for the interpretation of this figure.

values randomly from the following ranges; translations from $\pm 10\text{mm}$, rotation from $\pm 10^\circ$, scalings range from 1 ± 0.20 , and finally the shearing from $\pm 10^\circ$. Three cases were tested; only rigid parameters (rigid), rigid plus anisotropic scaling (affine) and finally including all six parameters (full-affine). In each case, initial transformations were composed of only relevant parameters and only those parameters were optimized for. 200 random registration experiments were performed in each cases. The statistics of transformation recovery errors are shown in Figure 7(a). From the figure, it is obvious that the performance decreases as the number of degrees of freedom increases as expected. The median error for each case are 6.3mm, 12.4mm and 12.2mm, respectively. It is interesting that including shearing in transformations did not introduce a significant change in the performance. Non-trivial random transformations were also used for ground-truth correlation of image intensities. The results shown in Figure 7(b) present parallelism with the results of transformation recovery error as shown in Figure 7(a). The median correlation is about 70% for rigid case dropping to 35% as the complexity increases.

The proposed method was further assessed in terms of its robustness against noise. To this end, varying degrees of white noise was added to the initial pair of synthetic images before applying any transformation. The noise added ranges between 0%-50% of image intensities. For each noise level and for each transformation type (i.e. rigid, affine, full-affine), 100 random transformations were applied to one of the synthetic images which was then recovered by the proposed approach. Since a multi-resolution setting would eliminate the added noise

during image pyramid creation, we performed registration only on the finest level to make sure that the noise is present during structure propagation. The results shown in Figure 7(c) demonstrate the robustness of the proposed approach to the amount of noise contained in subimages.

Another method that we use for the assessment of the stitching performance is the structure continuity index, SCI, as defined in Section III. Here, for the same experiments described above, we computed the SCI before and after a stitching was performed. Statistics on the SCI values computed from 21 experiments for each parameter are presented in Figure 8. Figure 8(a) shows the box plots for the SCI values before doing a stitching. Parameters t_x , s_x , i.e. translation and scaling in the direction parallel to the stitching boundary, respectively, as well as the rotation parameter, r_z , have a low median SCI value, which is mainly due to the fact that the initially applied transformations resulted in misalignments occurred less in the cases of t_y and s_y , i.e. translation and scaling parameters in the direction orthogonal to the stitching boundary. This is mainly because the structures are either scaled or shifted in the stitching direction which did not influence the structure alignments that much. SCI values were also computed after the stitching as presented in Figure 8(b) and the changes in SCI values is shown in Figure 8(c). We can make several observations looking at Figure 8(c). First of all, one can say that there has been an improvement in the estimation of the translation and the scaling parameters in the direction parallel to the stitching boundary. On the

other hand, there has been little or no improvement for their counterparts in the orthogonal direction. Rotation parameters also present an increased SCI value. One can conclude that stitching performance for correcting the misalignments caused by the variations in translation and scaling parameters in the direction orthogonal to the stitching boundary is not as good as the performance for correcting other types of misalignments. However, one should note that the structure continuity index, SCI, is based on the correlation between the votes cast by structure points in one of the subimages and the actual structure points in the other subimage. A transformation in the direction of stitching usually leads to a loss of structure points, which also influences the value of the SCI. Therefore, SCI should not be taken as a sole performance indicator when evaluating the proposed method.

B. Microscopic Images

Experiments have also been conducted on medical images to show the effectiveness of the proposed technique on real images. To this end, we have extracted 2D slices from a two-photon microscopy data set, which was taken from a mouse brain. The images were taken to examine the vasculature in the brain. There are two slices of size 512×512 pixels with $1.16 \mu\text{m}$ uniform pixel size. Again, as it was done in the previous synthetic experiments, the slices were cut into two pieces where one of the pieces underwent a transformation in order to induce a misalignment of structures across the cutting border. Figure 9 shows a series stitching experiments on the created pairs demonstrating the performance of the proposed method on microscopic images with synthetic misalignments.

In order to test the proposed registration technique using structure propagation on a real data without synthetic transformations and also to compare it to a competitive registration method, four microscopic image pairs, with subimages of 256×256 pixels resolution and $2.32 \mu\text{m}$ uniform pixel size, were used for the experiments. An example pair along with corresponding structure images are shown in Figures 10(a) and 10(b). The overlap region was further removed from one of the subimages for each pair in order to assess the proposed approach in no overlap case. Ground truth rigid transformations were obtained by using corresponding landmarks manually annotated in both subimages which resulted in overlap sizes of $53.47 \mu\text{m}$, $62.76 \mu\text{m}$, $55.79 \mu\text{m}$ and $62.76 \mu\text{m}$, respectively. For the experiments with overlap, three different scenarios were used for registration with and without structure propagation. First, a good initialization of transformation close to the ground truth was provided in the beginning. Secondly, white noise in the range of 50% of the image intensity were added to the images. Finally, subimages were placed only side-by-side where stitching boundaries were touching each other. The last scenario was further used for aligning without overlap. Again, mean distance error was used for the evaluation of the registration performance. For structure propagation, the value of scale parameter was fixed to $50 \mu\text{m}$ for each pair. The errors shown in Figure 11 indicate the comparable performance of the proposed approach in the presence of an overlap when a good initial estimate of the transformation is provided. Our

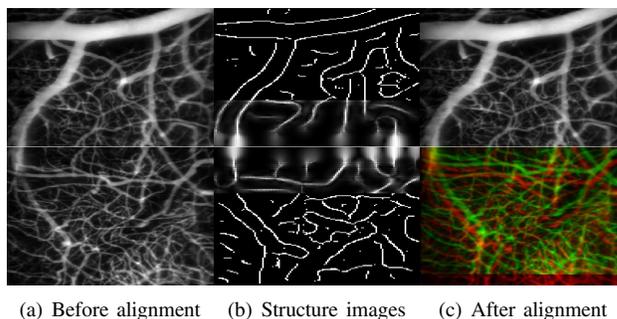


Fig. 10. A pair of microscopic images used for the evaluation of the proposed method in comparison to the classical approach. (a) Before an alignment was performed. (b) Corresponding structure images with propagated structures. (c) After the subimages were aligned using the proposed method. Aligned subimage is overlaid onto the original one.

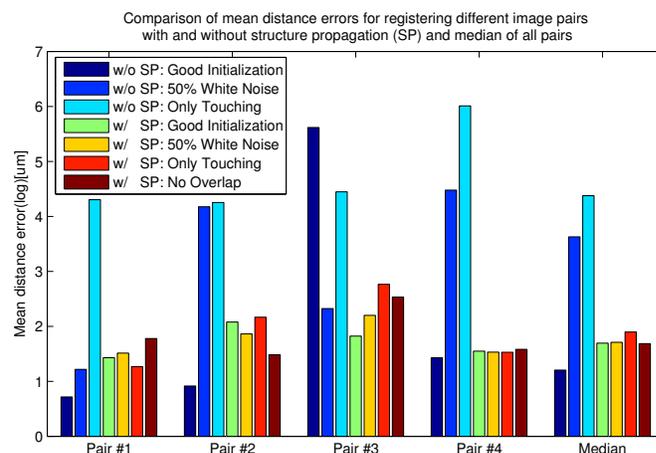


Fig. 11. Comparison of the proposed technique using structure propagation to registration without structure propagation. four different microscopic image pairs are used. For the experiments with overlap, three different scenarios are used. First, a good initialization close to the solution is provided in the beginning. Secondly, white noise in the range of 50% of the image intensity is added to the images. Finally, subimages are placed only side-by-side where stitching boundaries are touching each other. The last scenario is further used for aligning without overlap. Mean distance error of a randomly drawn point set is used for the evaluation. Note the errors are shown in log scale for better visualization. The results indicate the good performance of the proposed approach in various scenarios with and even without overlap whereas it is only possible to do an alignment in the presence of an overlap and when a good initialization is provided.

approach already outperforms the classical approach when the overlap region is degraded with white noise. However, when the alignment is initialized by placing the subimages only side-by-side (i.e. touching), the classical approach fails to find a transformation close to the ground truth whereas the current approach can still successfully align the subimages. Moreover, it was not possible to use the classical approach in no-overlap case where, again, a good performance was obtained by using structure propagation. The results reveal that registration with structure propagation can be successfully used also for correcting misalignments in real images.

Based on the previous experiments, it is clear that the physical gap between subimages has a negative influence on the stitching performance. Since one of the goal of the proposed techniques it to overcome the limitations caused by

the physical gap between the subimages, we have conducted further experiments to test the sensitivity of the technique to the size of physical gap between the subimages. To this end, we have manually introduced a physical gap of varying size to the lower subimage of the microscopic image pair P2 as illustrated in Figure 9(f). There are 21 different gap sizes sampled from the interval $[0\mu\text{m}, 20\mu\text{m}]$. For each gap size, misalignments of varying magnitude were introduced by applying 21 different translational transformations from the interval $\pm 20\mu\text{m}$ in the direction parallel to the stitching boundary. We have calculated the parameter recovery error in order to evaluate the influence of the amount of translation and the size of physical gap, respectively. Figure 12(a) shows that the stitching performance is not influenced too much by the amount of misalignment in the presence of a physical gap. However, when the parameter recovery error is plotted against the size of the introduced physical gap as shown in Figure 12(b), it becomes obvious that the stitching performance decreases with the increasing gap size. This is again an expected behavior since initially it was claimed that a non-overlap stitching is possible if the subimages are sufficiently close to each other. This is a much looser condition than the one that classical stitchings often require. It is still noteworthy that with the proposed technique it is still possible to stitch two subimages even if there is a physical gap on the order of $20\mu\text{m}$ between them. However, looking at Figure 12(b), in order to guarantee a successful stitching, one can allow a physical gap of around $12\mu\text{m}$ as a threshold after which the deviation of the error from the median starts to increase.

IV. DISCUSSION

The proposed stitching method presented here is designed to overcome the limitations of classical techniques by enabling a “perceptually good” alignment of images under difficult conditions like having small or no overlap by means of structure propagation. The method is addressing the state of art. However, there are still several limitations of the technique arising either from the employed method for the structure propagation or from the used image registration technique.

The tensor voting technique is capable of propagating the information regarding the underlying structure. However, with the current technique only the local orientation is taken into account for the propagation of the information. Therefore, currently, structures can be propagated only as linear structures. Although there have been attempts to include the local curvature of the underlying structure into the voting procedure [38], [49], they are usually iterative approaches which are not optimal due to high computational complexity and possibility of drifting. In the ideal case, a point should propagate its information based not only on its local orientation but also on the local curvature of the underlying structure.

The scale parameter of the weighting term is another factor that has an influence on the quality of propagation. A too large scale parameter causes blurry propagation of structures while having it too small yields to a limited propagation of local structural information. The choice of this parameter is strongly related to the size of the region of extension, Ω_{ie} in

Figure 2(a). In our experiments, we use a value of 20mm for the overlap region which proved to be enough. Accordingly, the scale parameter was chosen to be 40, which guaranteed that a structure point located 20mm into Ω_i can send a vote to the farthest point within the region of extension, Ω_{ie} . The choice of this parameter depends also on the size of physical gap. Obviously, for a large gap size, a larger value of scale parameter is needed.

Although promising results were observed during experiments, the obtained registration errors, especially the ones shown in Figure 7 indicate that the proposed registration is still far from being compared to the classical registration techniques. However, it should be noted that these results are despite the absence of an overlap whereas no errors have been reported so far in such challenging cases.

Another important point to be discussed is the similarity measure for the affine registration. Currently, normalized cross-correlation (NCC) is employed as the cost function for evaluating the (dis)similarity between S_1 and S_2 as described in Section II-D. Although, satisfactory results are obtained using NCC, a more specific cost function could be proposed, which accounts for the differences in the representation of the propagated and existing structures.

Although the demonstrations include only two images, the proposed method could also be extended for the alignment of multiple images. If a prior information is available, then, a simultaneous optimization strategy can be employed to obtain the optimal transformation parameters for neighboring images. Otherwise, a method similar to the one employed in [23] can be used to first identify adjacency relationships between subimages. Here, the proposed structure propagation technique can be employed to enable the comparison between subimages when no overlap is available.

There are a few challenges associated with the extension of the proposed approach to 3D. First of all, although the tensor voting framework is readily extensible to 3D, the interpretation of the saliences would need further consideration. In 2D, only curve saliences are derived from tensors whereas, in 3D, both surface and curve saliences are present and they should simultaneously be considered during registration. The additional degrees of freedom add further complexity during transformation estimation. Furthermore, additional computational cost should be expected in 3D also during the determination of adjacency relationships if the proposed structure propagation technique is used to establish overlap regions between subimages.

V. CONCLUSION

In many medical imaging applications ranging from microscopy to ultrasound to digital pathology, a wide field of view is usually desired to enable a better analysis at different scales. However, having a wide field of view is often limited by the capabilities of the imaging devices in guaranteeing large field of view at highest resolution. Therefore, it is a common practice to acquire smaller pieces which are then stitched together in order to get a larger field of view. There has been many solutions proposed in the past for the stitching of subimages in various applications. The common requirement

- [12] J. Totz, K. Fujii, P. Mountney, and G.-Z. Yang, "Enhanced visualisation for minimally invasive surgery." *Int J Comput Assist Radiol Surg*, vol. 7, no. 3, pp. 423–432, May 2012. [Online]. Available: <http://dx.doi.org/10.1007/s11548-011-0631-z>
- [13] A. Warren, P. Mountney, D. Noonan, and G.-Z. Yang, "Horizon stabilized–dynamic view expansion for robotic assisted surgery (hs-dve)." *Int J Comput Assist Radiol Surg*, vol. 7, no. 2, pp. 281–288, Mar 2012. [Online]. Available: <http://dx.doi.org/10.1007/s11548-011-0603-3>
- [14] C. Wachinger, W. Wein, and N. Navab, "Three-dimensional ultrasound mosaicing." *Med Image Comput Comput Assist Interv*, vol. 10, no. Pt 2, pp. 327–335, 2007.
- [15] L. Brattain and R. Howe, "Real-time 4d ultrasound mosaicing and visualization." *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2011*, pp. 105–112, 2011.
- [16] C. Wachinger, B. Glocker, J. Zeltner, N. Paragios, N. Komodakis, M. S. Hansen, and N. Navab, "Deformable mosaicing for whole-body mri." *MICCAI*, vol. 11, no. Pt 2, pp. 113–121, 2008.
- [17] O. Dzyubachyk, B. P. F. Lelieveldt, J. Blaas, M. Reijnierse, A. Webb, and R. J. van der Geest, "Automated algorithm for reconstruction of the complete spine from multistation 7t mr data." *Magn Reson Med*, Jul 2012. [Online]. Available: <http://dx.doi.org/10.1002/mrm.24404>
- [18] V. Kaynig, B. Fischer, E. Mueller, and J. Buhmann, "Fully automatic stitching and distortion correction of transmission electron microscope images." *Journal of structural biology*, vol. 171, no. 2, pp. 163–173, 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.jsb.2010.04.012>
- [19] P. Khurd, R. Oketokoun, T. Gajera, J. Frangioni, L. Grady, H. Sundar, S. Gibbs-Strauss, and A. Kamen, "Global error minimization in image mosaicing using graph connectivity and its applications in microscopy." *Journal of Pathology Informatics*, vol. 2, 2011. [Online]. Available: <http://dx.doi.org/10.4103/2153-3539.92039>
- [20] W. Kong and B. B. Kimia, "On solving 2d and 3d puzzles using curve matching," in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition CVPR 2001*, vol. 2, 2001.
- [21] E. Tsamoura and I. Pitas, "Automatic color based reassembly of fragmented images and paintings." *IEEE Transactions on Image Processing*, vol. 19, no. 3, pp. 680–690, 2010.
- [22] M. Berlanga, S. Phan, E. Bushong, S. Wu, O. Kwon, B. Phung, S. Lamont, M. Terada, T. Tasdizen, M. Martone *et al.*, "Three-dimensional reconstruction of serial mouse brain sections: solution for flattening high-resolution large-scale mosaics." *Frontiers in neuroanatomy*, vol. 5, p. 17, 2011. [Online]. Available: <http://dx.doi.org/10.3389/fnana.2011.00017>
- [23] M. Brown and D. Lowe, "Automatic panoramic image stitching using invariant features." *International Journal of Computer Vision*, vol. 74, no. 1, pp. 59–73, 2007.
- [24] M. Emmenlauer, O. Ronneberger, A. Ponti, P. Schwarb, A. Griffla, A. Filippi, R. Nitschke, W. Driever, and H. Burkhardt, "Xuvtools: free, fast and reliable stitching of large 3d datasets." *J Microsc*, vol. 233, no. 1, pp. 42–60, 2009. [Online]. Available: <http://dx.doi.org/10.1111/j.1365-2818.2008.03094.x>
- [25] D. J. Field, A. Hayes, and R. F. Hess, "Contour integration by the human visual system: evidence for a local "association field"." *Vision Res*, vol. 33, no. 2, pp. 173–193, Jan 1993.
- [26] P. Parent and S. W. Zucker, "Trace inference, curvature consistency, and curve detection." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 8, pp. 823–839, 1989.
- [27] J. August and S. Zucker, "Sketches with curvature: The curve indicator random field and markov processes." *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 25, no. 4, pp. 387–400, 2003.
- [28] M. Singh and J. M. Fulvio, "Bayesian contour extrapolation: geometric determinants of good continuation." *Vision Res*, vol. 47, no. 6, pp. 783–798, Mar 2007. [Online]. Available: <http://dx.doi.org/10.1016/j.visres.2006.11.022>
- [29] A. Sashua and S. Ullman, "Structural saliency: The detection of globally salient structures using a locally connected network," in *Proc. Second Int Computer Vision. Conf.*, 1988, pp. 321–327.
- [30] S. Sarkar and K. L. Boyer, "Integration, inference, and management of spatial information using bayesian networks: perceptual organization." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 3, pp. 256–274, 1993.
- [31] K. K. Thornber and L. R. Williams, "Analytic solution of stochastic completion fields," in *Proc. Symp. Int Computer Vision*, 1995, pp. 617–622.
- [32] G. Medioni, M. Lee, and C. Tang, *A computational framework for segmentation and grouping*. Elsevier Science, 2000, vol. 1.
- [33] P. Mordohai and G. Medioni, "Tensor voting: a perceptual organization approach to computer vision and machine learning." *Synthesis Lectures on Image, Video, and Multimedia Processing*, vol. 2, no. 1, pp. 1–136, 2006.
- [34] W. Tong, C. Tang, P. Mordohai, and G. Medioni, "First order augmentation to tensor voting for boundary inference and multiscale analysis in 3d." *IEEE Trans. PAMI*, vol. 26, no. 5, pp. 594–611, 2004.
- [35] E. Franken, M. van Almsick, P. Rongen, L. Florack, and B. ter Haar Romeny, "An efficient method for tensor voting using steerable filters." *ECCV 2006*, pp. 228–240, 2006.
- [36] B. King, *Range data analysis by free-space modeling and tensor voting*. ProQuest, 2008.
- [37] J. Jia and C. Tang, "Image stitching using structure deformation." *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 30, no. 4, pp. 617–631, 2008.
- [38] L. A. Loss, G. Bebis, and B. Parvin, "Iterative tensor voting for perceptual grouping of ill-defined curvilinear structures." *IEEE Trans Med Imaging*, vol. 30, no. 8, pp. 1503–1513, Aug 2011. [Online]. Available: <http://dx.doi.org/10.1109/TMI.2011.2129526>
- [39] J. Jia and C.-K. Tang, "Image repairing: robust image synthesis by adaptive nd tensor voting." in *CVPR*, vol. 1, 2003.
- [40] J. Jia and C. Tang, "Tensor voting for image correction by global and local intensity alignment." *IEEE Trans. PAMI*, vol. 27, no. 1, pp. 36–50, 2005.
- [41] L. Risser, F. Plouraboué, and X. Descombes, "Gap filling of 3-d microvascular networks by tensor voting." *IEEE Transactions on Medical Imaging*, vol. 27, no. 5, pp. 674–687, 2008.
- [42] D. Zikic, B. Glocker, O. Kutter, M. Groher, N. Komodakis, A. Kamen, N. Paragios, and N. Navab, "Linear intensity-based image registration by markov random fields and discrete optimization." *Med Image Anal*, vol. 14, no. 4, pp. 550–562, Aug 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.media.2010.04.003>
- [43] O. Monga, R. Deriche, G. Malandain, and J.-P. Cocquerez, "Recursive filtering and edge tracking: two primary tools for 3-d edge detection." *Image and Vision Computing*, vol. 9, no. 4, pp. 203–214, August 1991.
- [44] G. Medioni, C. Tang, and M. Lee, "Tensor voting: Theory and applications." *Proceedings of RFIA, Paris, France*, 2000.
- [45] G. Medioni and S. Kang, *Emerging topics in computer vision*. Prentice Hall PTR, 2004.
- [46] N. Komodakis and G. Tziritas, "Approximate labeling via graph cuts based on linear programming." *IEEE Trans. PAMI*, vol. 29, no. 8, pp. 1436–1453, 2007.
- [47] N. Komodakis, G. Tziritas, and N. Paragios, "Performance vs computational efficiency for optimizing single and dynamic mrfs: Setting the state of the art with primal-dual strategies." *Computer Vision and Image Understanding*, vol. 112, no. 1, pp. 14–29, 2008.
- [48] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios, "Dense image registration through mrfs and efficient linear programming." *Med Image Anal*, vol. 12, no. 6, pp. 731–741, 2008.
- [49] W. Tong and C. Tang, "Robust estimation of adaptive tensors of curvature by tensor voting." *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 27, no. 3, pp. 434–449, 2005.