

Slice-to-volume medical image registration: a survey

Enzo Ferrante^{a,b}, Nikos Paragios^{a,c}

^a*Center for Visual Computing, CentraleSupélec, INRIA, Université Paris-Saclay,
Grande Voie des Vignes, Chatenay-Malabry, 92295, France*

^b*Biomedical Image Analysis (BioMedIA) Group, Department of Computing, Imperial College London,
South Kensington Campus, 180 Queen's Gate, London, SW7 2AZ, UK*

^c*TheraPanacea, 24 Rue du Faubourg Saint-Jacques, 75014 Paris*

Abstract

During the last decades, the research community of medical imaging has witnessed continuous advances in image registration methods, which pushed the limits of the state-of-the-art and enabled the development of novel medical procedures. A particular type of image registration problem, known as slice-to-volume registration, played a fundamental role in areas like image guided surgeries and volumetric image reconstruction. However, to date, and despite the extensive literature available on this topic, no survey has been written to discuss this challenging problem. This paper introduces the first comprehensive survey of the literature about slice-to-volume registration, presenting a categorical study of the algorithms according to an ad-hoc taxonomy and analyzing advantages and disadvantages of every category. We draw some general conclusions from this analysis and present our perspectives on the future of the field.

Keywords: Bibliographical review, slice-to-volume registration, medical image registration, medical image analysis.

1. Introduction

Image registration is the process of aligning and combining data coming from more than one image source into a unique coordinate system. This problem has become one of the pillars of computer vision and medical imaging. Slice-to-volume registration, a particular case of image registration problem, has received further attention in the medical imaging community during the last decade. In this case, instead of registering images with same dimension, we seek to determine the slice (corresponding to an arbitrary plane) from a given 3D volume that corresponds to an input 2D image.

Several medical applications requiring slice-to-volume mapping have emerged and pushed the community towards developing more accurate and efficient strategies. Such medical imaging tasks can be classified in two main categories: those related to image fusion for image

Email addresses: `e.ferrante@imperial.ac.uk` (Enzo Ferrante), `nikos.paragios@ecp.fr` (Nikos Paragios)

guided interventions; and those related to motion correction and volume reconstruction. In the first category, pre-operative 3D images and intra-operative 2D images need to be fused to guide surgeons during medical interventions. Slice-to-volume registration plays a key role in this process, allowing the physicians to navigate 3D pre-operative high-resolution annotated data using low-resolution 2D images acquired in real-time during surgery. In the second category, the goal is to correct for misaligned slices when reconstructing a volume of a certain modality. A typical approach to solve this task consists in mapping individual slices onto a reference volume in order to correct for inter-slice misalignment. Again, the development of accurate slice-to-volume registration algorithms is crucial to successfully tackle this problem.

Slice-to-volume registration is also known as 2D/3D registration, due to the dimension of the images involved in the registration process. However, this term is ambiguous since it describes two different problems depending on the technology used to capture the 2D image: it may be a projective (e.g. x-ray) or sliced/tomographic (e.g. US) image. Even if both problems share similarities in terms of image dimensionality, every formulation requires a different strategy to estimate the solution. The lack of perspective and different image geometry (Birkfellner et al., 2007) inherent to both modalities, make it necessary to come up with distinct strategies to solve these registration problems. Moreover, a pixel in any 2D projective image does not correspond to only one voxel from the volume (this is the case for slice-to-volume), but to a projection of a set of them in certain perspective. Therefore, the type of functions used to measure similarities between the images is necessarily different in every case. While most of the projective 2D/3D image registration methods require to bring the images into dimensional correspondence (by different strategies like projection, back-projection or reconstruction (Markelj et al., 2010)), in case of slice-to-volume registration, pixels from the 2D image can be directly compared with voxels from the volume. In this review we focus on slice-to-volume registration, while a more comprehensive overview about projective 2D to 3D image registration is presented in the survey by (Markelj et al., 2010).

Slice-to-volume registration could be considered as an extreme case of 3D-3D registration, where one of the 3D images contains only one slice. Even if theoretically true, 3D-3D registration methods can not be extrapolated in a straightforward way to the slice-to-volume scenario. This holds, particularly, for registration methods based on image information, since the descriptors used to quantify similarities between images, normally assume that the amount of information available from both images is balanced. The fact that a single slice (or even a few sparse slices) provides less information than an entire volume, should be explicitly considered in the problem formulation. Moreover, specific geometrical constraints like planarity satisfaction and in-plane deformation restrictions, arise in the case of slice-to-volume registration, which are not applicable in the setting of dimensional correspondence.

Despite the extensive literature available about slice-to-volume registration, to date, no survey has been written to discuss this emerging field. In this review, we discuss the literature about slice-to-volume registration, offering a comprehensive survey of the articles tackling this challenging problem, proposing a categorical study of the related algorithms according to a taxonomy and analyzing the advantages and disadvantages of each category. We propose a taxonomy based on previous surveys on image registration (Markelj et al., 2010; Sotiras et al.,

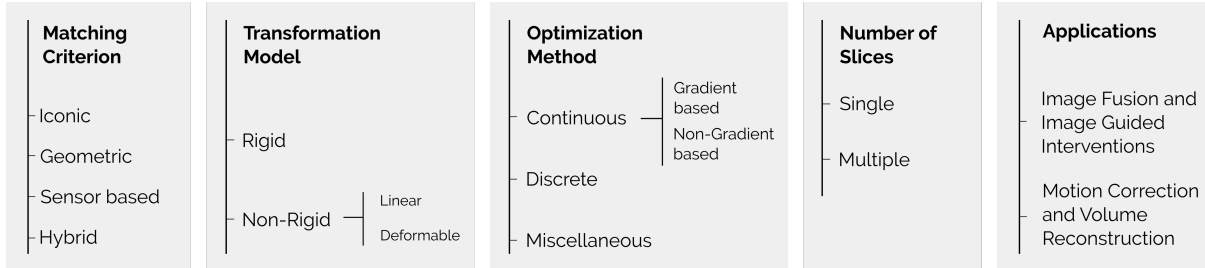


Figure 1: Taxonomy used to classify the slice-to-volume registration publications analyzed in this survey.

2013) and adapted to the particular slice-to-volume case (see Figure 1). We classified the related work according to several principles: (i) matching criterion (section 2), which specifies the strategies to define similarities between the images; (ii) transformation model (section 3), indicating the nature of the models used to bring images into spatial correspondence; (iii) optimization (section 4), differentiating the approaches according to the strategy used to infer the optimal transformation model; (iv) number of slices (section 5), which splits the methods in two groups, according to whether they require a single or multiple source image slices; (v) applications (section 6), where we identify the main clinical scenarios that have motivated research in the field; and finally (vi) validation strategy (section 7) where we summarize the alternative approaches adopted to validate the proposed slice-to-volume registration methods.

Papers included in this review were systematically selected, searching for papers containing the term 'slice to volume' and 'registration' (and their variations), and carefully choosing those that either present a methodological contribution or apply slice-to-volume registration to a particular problem introducing some degree of novelty.

1.1. Definition of Slice-to-Volume Registration

Let us start by giving a formal definition of slice-to-volume registration. Given a 2D image I and a 3D volume J , we seek a mapping function $\hat{\Theta}$ which optimally aligns the tomographic slice I with the volumetric image J , through the minimization of the following objective function:

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \mathcal{M}(I, J; \Theta) + \mathcal{R}(\Theta), \quad (1)$$

where \mathcal{M} represents the image similarity term (the so-called matching criterion) and \mathcal{R} the regularization term. Note that this mapping may be rigid or non-rigid, depending on whether we allow image I (or its corresponding reformatted slice from J) to be deformed or not. If we estimate only a rigid mapping (i.e. we calculate a 6 degrees of freedom rigid transformation or even a more restrictive one), we name the problem rigid slice-to-volume registration. In case we also infer some sort of deformation model or we consider more expressive linear transformations (such as affine transformations), we call it non-rigid registration. We refer the reader to section 3 for more information about different transformation models.

The matching criterion \mathcal{M} measures the similarity between the 2D image and its corresponding mapping (slice) to the 3D volume. Usually, it is defined using intensity information or salient structures from I and J . A complete discussion about matching criteria in the context of slice-to-volume registration (where we also consider methods that do not use any image information but rely on external sensors) is presented in section 2.

The regularization term \mathcal{R} imposes constraints on the solution that can be used to render the problem well posed. It also may encode geometric properties on the extended (plane selection and plane deformation in case of non-rigid registration) transformation model. The choice of regularizer depends on the transformation model. While simple models like rigid body transformations can be explicitly estimated even without regularizer, the term \mathcal{R} becomes crucial in more complex non-rigid scenarios to ensure realistic results. In the context of slice-to-volume registration, the regularizer can be used to impose planarity constraints to the solution (when out-of-plane deformations are not allowed) or to limit the out-of-plane deformation magnitude guaranteeing realistic and plausible transformations. When available, prior knowledge about tissue elasticity can also be encoded through the regularizer. Moreover, as we will see in Section 5, when dealing with multiple slices at the same time, regularization can be used to impose consistency among them.

We aim at optimizing the energy defined in equation 1, by choosing the best $\hat{\Theta}$ that aligns the 2D and 3D images. Depending on the variables we are trying to infer, and their optimality guarantees, they can be classified in different categories. A full study of this topic is included in section 4.

The general definition given in equation 1 considers a single slice as the input to the registration process. However, for the sake of completeness, in this survey we also consider the so-called multi slice-to-volume registration approaches, for which several (but sparse) slices are registered to a full 3D volume. It is worth noting that we do not include methods that perform a prior 3D reconstruction from the input 2D slices before registration (see, for example, the work by Arbel et al. (2001)), since these methods reduce the problem to the classic 3D-3D scenario, which is not within the scope of our work. On the contrary, we consider those methods that directly register the subset of 2D slices (which can be orthogonal, parallel, arbitrary or even without an a-priori known spatial relation), rather than reconstructing a volume. Discussion and classification of the approaches according to this criterion is presented in section 5.

Table 1: Comparative table: It includes a selection of those papers cited in this survey that were published in the most important journals and conferences related to medical image analysis. They are classified according to the taxonomy used in this work, namely Matching Criterion (Section 2), Transformation Model (Section 3), Optimization Method (Section 4) and Number of Slices (Section 5).

Reference	Matching Criterion				Transformation Model		Optimization Method			Number of slices	
	Iconic	Geometric	Sensor Based	Hybrid	Rigid	Non-Rigid	Continuous	Discrete	Heuristic	Single	Multi
Gourdon and Ayache (1994)		X			X		X			X	
Kim et al. (1999)	X				X		X			X	
Fei et al. (2002)	X				X		X			X	
Fei et al. (2003a)	X				X		X			X	
Fei et al. (2004)	X				X		X			X	
Fei et al. (2003c)	X				X		X			X	
Park et al. (2004)	X				X		X				X
Penney et al. (2004a)	X		X		X		X				X
Yeo et al. (2004)	X				X		X			X	
Zarow et al. (2004)	X	X		X		X	X				X
Bao et al. (2005)			X		X		-	-	-	X	
Noble et al. (2005)	X				X		X			X	
Rousseau and Glenn (2005)	X				X		X				X
Smolíková Wachowiak et al. (2005)	X				X		X			X	
Rousseau et al. (2006)	X				X		X				X
Chandler et al. (2006)	X				X		X				X
Penney et al. (2006)	X		X		X		X				X
Micu et al. (2006)	X				X	X	X			X	
Yeo et al. (2006)	X				X		X			X	
Birkfellner et al. (2007)	X				X		X		X	X	
Boer et al. (2007)		X			X			X		X	
Jiang et al. (2007a)	X				X		X				X
Jiang et al. (2007b)	X				X		X				X
Sun et al. (2007)	X		X		X		X				X
Chandler et al. (2008)	X				X		X				X
Dalvi and Abugharbieh (2008)		X			X		X				X
Brooks et al. (2008)	X					X	X				X
Gefen et al. (2008)	X					X	X		X	X	
Hummel et al. (2008)	X		X	X	X		X			X	
Kim et al. (2008b)	X				X		X				X
Kim et al. (2008a)	X				X		X			X	
Leung et al. (2008)	X					X	X				X
Wein et al. (2008a)	X		X	X	X	X	X				X
Wein et al. (2008b)	X					X	X				X
Xu et al. (2008)	X		X	X	X		X				X
Yeo et al. (2008)	X				X		X			X	
Frühwald et al. (2009)	X				X		X			X	
Gholipour and Warfield (2009)	X				X		X				X
Huang et al. (2009)	X		X	X	X		X				X
Jiang et al. (2009)	X				X		X				X
Osechinskiy and Kruggel (2009)	X					X	X			X	

Table 2: Continuation of Table 1

Reference	Matching Criterion				Transformation Model		Optimization Method			Number of slices	
	Iconic	Geometric	Sensor Based	Hybrid	Rigid	Non-Rigid	Continuous	Discrete	Heuristic	Single	Multi
San José Estépar et al. (2009)	X				X		X			X	
Elen et al. (2010)	X				X		X				X
Gholipour et al. (2010)	X				X		X				X
Honal et al. (2010)	X					X	X			X	
Kim et al. (2010a)	X				X		X				X
Kio Kim et al. (2010)	X				X		X				X
Tadayyon et al. (2010)	X				X	X	X		X		X
Gholipour et al. (2011)	X		X		X		X			X	
Kim et al. (2011)	X				X		X				X
Marami et al. (2011)	X					X	X				X
Osechinskiy and Kruggel (2011a)	X				X	X	X			X	
Tadayyon et al. (2011)	X				X	X	X		X		X
Xiao et al. (2011)	X					X	X				X
Yu et al. (2011)	X				X		X				X
Kuklisova-Murgasova et al. (2012)	X				X		X				X
Mitra et al. (2012)	X	X			X		X			X	
Yan et al. (2012)	X		X	X	X		X				X
Zakkaroff et al. (2012)	X				X		X				X
Cifor et al. (2013)	X				X				X		X
Ferrante and Paragios (2013)	X					X		X		X	
Lin et al. (2013)	X				X				X	X	
Seshamani et al. (2013)	X				X		X				X
Po Su et al. (2013)	X					X	X			X	
Chicherova et al. (2014)		X			X		X			X	
Eresen et al. (2014)	X		X	X	X			X		X	
Fogtmann et al. (2014)	X				X		X				X
Fuerst et al. (2014)	X		X	X	X	X	X				X
Museyko et al. (2014)	X					X	X			X	
Nir et al. (2014)	X	X				X	X				X
Rivaz et al. (2014c)	X					X	X				X
Rivaz et al. (2014b)	X					X	X				X
Rivaz et al. (2014a)	X					X	X				X
Schulz et al. (2014)			X		X		-	-	-	X	
Xu et al. (2014a)	X					X	X				X
Xu et al. (2014b)	X				X		X			X	
Ferrante et al. (2015a)	X					X		X		X	
Ferrante et al. (2015b)	X					X		X		X	
Hallack et al. (2015)				X		X	X				X
Kainz et al. (2015)	X				X		X			X	
Yavariabdi et al. (2015)		X				X	X			X	
Chen et al. (2016)	X				X	X	X				X
Guzman et al. (2016)	X				X	X	X				X
Marami et al. (2016b,a)	X				X	X	X				X
Murgasova et al. (2016)	X				X		X				X
Porchetto et al. (2016)	X				X			X		X	
Xiao et al. (2016)	X				X		X				X

2. Matching Criterion

The matching criterion (also known as (dis)similarity measure, merit function or distance function) quantifies the level of alignment between the images, and it is typically used to guide the optimization process of the transformation model. Depending on the nature of information exploited in the matching process, registration methods can be classified as iconic (we use voxel intensities to quantify similarity), geometric (we use a sparse set of salient image locations to guide the registration) or hybrid methods (we combine both strategies). In the particular case of slice-to-volume registration, there are also some approaches which instead of using image information, rely on other non-image technologies; we will refer to them as sensor based methods.

Slice-to-volume registration excluding image based methods (i.e. sensor based slice-to-volume registration) is mainly performed using two different technologies: optical (OTS) and electromagnetic (EMTS) tracking systems (Birkfellner et al., 2008). Optical systems employ different types of cameras (such as infrared (IR), standard RGB or laser cameras) to track markers which help to identify the current position of the objects. Electromagnetic positioning systems perform tracking based on a system of transmitter, sensors and processing unit that localizes the position and orientation of a target object by measuring electromagnetic field properties. Both methods inherit constraints on the nature of transformation that they can estimate, since only rigid transformations can be obtained through these technologies. Moreover, accurate initial calibration is usually required at the beginning, due to the fact that positioning is done through the estimation of the relative displacement from the previous one and therefore errors can be propagated and accumulated through all the process. Instead, registration algorithms exploiting image based information -iconic or geometric- can deal with elastic anatomical changes and they are less sensitive to initial errors (since it is simple to correct them iteratively during the registration process). On the negative side, such methods can be more sensitive to image noise, which is frequently present in intra-operative, real-time and low-quality modalities, normally corresponding to the input 2D image. In addition, the amount of information of an image slice is sparse when compared to a volumetric image, resulting on ambiguities in terms of image matching that render the registration problem ill-posed. Different strategies have been developed to deal with these problems depending on the choice of the matching criterion.

Image registration can be monomodal (when the slice and the volume are captured with the same type of image technology) or multimodal (when slice and volume refer to different modalities, e.g. US slice and MRI or CT volume). In the former case, the task of measuring the similarity between the images is simpler, since pixel/voxel intensity values corresponding to the same anatomical structure are highly correlated, or even identical, in both images. Therefore, iconic methods may perform better since the main issue associated to them -the difficulty to explain image similarities using pixel/voxel correspondences- is already solved. In case of multimodality, where the relation between pixel intensities is not obvious, there are two major alternatives: to continue using the iconic matching criterion but defining more complex similarity measures, or to adopt a geometric or sensor based strategy which appears to be more robust when dealing with different image modalities.

2.1. Iconic

Iconic matching criteria are defined using image intensity information. Similarity between images is measured using functions that act on the pixel/voxel intensity level. Standard signal processing tools, information theoretical approaches or even similarity measures defined for particular image modalities can be considered. The challenge lies in describing both images on a common space where they can be easily compared. In particular, in the context of multimodal registration, where voxel intensities corresponding to the same anatomical or functional structures are dissimilar. There are two desirable properties sought at the definition of any iconic similarity measure: (i) to be convex, since it simplifies the optimization process and (ii) to be discriminative, in the sense that it assigns distinct values to different tissues or anatomical structures. Convexity is also associated

Typically, in the context of slice-to-volume registration, an iconic matching criterion $\delta(I_1, I_2) : \Omega_1 \times \Omega_2 \rightarrow \mathfrak{R}$ is defined to quantify the similarity between two slices (or patches, according to whether we specify a global or local function). Such a similarity measure varies depending on the modalities we are trying to register. In a monomodal scenario, simple similarity measures such as the sum of absolute differences (SAD) (Ferrante and Paragios, 2013; Ferrante et al., 2015a,b; Leung et al., 2008), sum of square differences (SSD) (Fogtman et al., 2014; Heldmann and Papenberg, 2009; Leung et al., 2008; Marami et al., 2011; Miao et al., 2014; Osechinskiy and Kruggel, 2011b; Po Su et al., 2013; Seshamani et al., 2013; Xu et al., 2008; Yu et al., 2008, 2011) or even mean of square differences (Fogtman et al., 2014; Gholipour et al., 2010; Gholipour and Warfield, 2009; Honal et al., 2010; Kim et al., 2008b, 2010a; Tadayyon et al., 2010; Tadayyon, 2010) of the intensity values can be used. In vector notation, SAD can also be seen as the L^1 norm of vectorized image, whereas SSD would correspond to the L^2 norm. These metrics assume a straightforward correspondence between the intensity values in both images, which is not necessarily the case. In real clinical cases like image guided surgery, images captured during the procedures tend to be noisy, low resolution and sometimes even different modalities have to be fused. In these cases, SSD and SAD will present poor performance.

More complex metrics, exploiting statistical properties of the observed intensity values in both images, have also been proposed. These methods, on top of handling identity transformations, can cope with piece-wise linear relationships between the intensities in the images to be registered. In these cases, image pixels in both images are seen as entries of two random vectors X and Y . Cross-correlation (CC) is a well known function widely used in the fields of signal processing and statistics, also applied in several slice-to-volume registration studies (Birkfellner et al., 2007; Elen et al., 2010; Fei et al., 2002, 2003a,b, 2004; Frühwald et al., 2009; Hummel et al., 2008; Jiang et al., 2007a,b, 2009; Kainz et al., 2015; Kim, 2005; Miao et al., 2014; Noble et al., 2005; Osechinskiy and Kruggel, 2009, 2011b; Xiao et al., 2016; Xu et al., 2008; Yan et al., 2012; Zarow et al., 2004). CC measures the correlation between the entries of X and Y . It is simple to compute and, more importantly, invariant to shifts and scaling in the intensity domain. Another metric is the correlation ratio (CR), which has shown promising results even in multimodal image registration (Roche et al., 1998). It measures functional dependencies between X and Y , taking values between 0 (no functional

dependence) and 1 (purely deterministic dependence). CR is intrinsically asymmetrical, since the two variables (images) do not play the same role in the functional relationship. In other words, unlike CC, CR offers different values depending on the order that images were considered. CR has been used as an iconic criterion for slice-to-volume registration in (Marami et al., 2011; Osechinskiy and Kruggel, 2011b; Smolíková Wachowiak et al., 2005).

Information theoretic similarity measures are usually the choice of preference for multi-modal registration, and slice-to-volume registration is not an exception. The most popular is mutual information (MI), which measures the statistical dependence or information redundancy between the image intensities of corresponding distributions in both images, that is assumed to be maximal if the images are geometrically aligned (Maes et al., 1997). It requires an estimation of joint and marginal probability density functions (PDFs) of the intensities in every image. Given the information sparse nature of slice-to-volume registration when compared to the volume-to-volume scenario, the estimation of these joint PDFs for every slice -especially in slices of low image resolution/number of samples- is a hard task and may redound to poor MI-based registration results. An alternative approach to improve MI-based slice-to-volume registration was developed by incorporating informative PDF priors in the context of fMRI time-series registration Bhagalia et al. (2009). First, it was shown that slices located near the middle of the head scans give more reliable PDFs and MI estimations because they refer to a richer information space than the end slices (top or bottom). End-slices registration is then guided by a joint PDF prior based on intensity counts from registered center-slices. Alternatively, a better MI calculation in slice-to-volume registration can be achieved by using a PDF estimate that retains as much information about voxel intensities as possible from the higher resolution anatomical data set, when registering 2D MR scout scan to a complete 3D MR brain volume (Chandler et al., 2004). MI is a widely used similarity measure for slice-to-volume registration, adopted in an important number of methods in the last decades, e.g. (Birkfellner et al., 2007; Brooks et al., 2008; Eresen et al., 2014; Fei et al., 2002; Ferrante et al., 2015b; Fogtman et al., 2014; Gill et al., 2008; Huang et al., 2009; Kim et al., 1999; Museyko et al., 2014; Park et al., 2004; Rousseau et al., 2006; Seshamani et al., 2013; Smolíková Wachowiak et al., 2005; Tadayyon et al., 2010; Xiao et al., 2011; Xu et al., 2014a; Yeo et al., 2004; Zakkaroff et al., 2012; Zarow et al., 2004; Chen et al., 2016). One of the main drawbacks of MI is that it varies when the overlapping area between the images changes, i.e. it is not invariant to changes in the overlap region throughout registration. It could happen that while estimating the transformation model, some potential solutions lie out of the volume. In such cases, an overlap invariant function would be of choice. To this end, a modified version of MI, the normalized mutual information (NMI), can be applied, which is simply the ratio of the sum of the marginal entropies and the joint entropy (Studholme et al., 1999). Another advantage of NMI with respect to MI is its range: it conveniently takes values between 0 and 1. NMI has been used as well for slice-to-volume registration (Chandler et al., 2006; Elen et al., 2010; Gefen et al., 2008; Hummel et al., 2008; Jiang et al., 2007a; Kainz et al., 2015; Kuklisova-Murgasova et al., 2012; Leung et al., 2008; Marami et al., 2011; Miao et al., 2014; Rousseau and Glenn, 2005).

Using prior knowledge like segmentation masks during the iconic registration process can

be useful (Shakeri et al., 2016). These approaches, also known as region-based methods, employ intensities information or statistics to describe a pre-segmented region. Chan-Vese metric (Chan and Vese, 2001) for instance, aims to minimize the intensity variances on the regions inside and outside a given segmentation contour. Nir and coworkers (Nir and Tannenbaum, 2011; Nir et al., 2014) applied this matching criterion to the problem of aligning multiple slices of histological images to in vivo MR images of the prostate.

Border information is another low level visual cue that was exploited by iconic methods. It is usually determined from the intensity gradient of the images, which gives an idea of the image structure defined by intensity changes, independently of their actual value. MI as well as SAD or SSD can be applied on top of the gradient magnitudes of both images. In (Brooks et al., 2008), for example, MI between the gradient magnitudes of an US slice and MRI volume was used. Po Su et al. (2013) applied SSD on both image intensities and gradients, combining them in a unique similarity measure. In (Xu et al., 2014b), the matching criterion was defined using the normalized gradient field of the images, while CC of both intensity and gradient magnitude was adopted in (Xu et al., 2008). In (Lin et al., 2013), a merit function called Weighted Edge-Matching Score (WEMS) was used to evaluate the similarity of a 2D real-time image and the corresponding reformatted slice to search for the 3D pose of a bone model.

More robust similarity measures have been introduced more recently to deal with problems in medical image registration, and in the slice-to-volume case in particular. Remarkable contributions have been made by (Wein et al., 2008a; Fuerst et al., 2014) to this field. They propose different similarity measures based on the simulation of US images from MRI and from CT, which can deal with these challenging multimodal registration problems. In (Wein et al., 2007, 2008a), novel methods for simulation of ultrasonic effects from CT data are presented, together with a new similarity measure entitled Linear Correlation of Linear Combination (LC^2), which is invariant to missing simulation details, yielding smooth properties and a global optimal corresponding to the correct alignment. Since they simulate US imaging effects with respect to the probe geometry, the original B-mode scan planes of the sweep are used instead of 3D reconstruction, making it suitable for multi slice-to-volume registration. LC^2 was used by Fuerst et al. (2014), although without the simulation process: the similarity measure was defined by locally matching US intensities to both MRI intensity and gradient magnitude. More recently, a variation on the LC^2 entitled $BOXLC^2$ which produces a measurable speed improvement over LC^2 when running on GPU (while retaining similar performance in the context of single slice-to-volume registration) was presented in (Pardasani et al., 2016).

Another robust similarity measure, the modality independent neighborhood descriptor (MIND) was proposed by Heinrich et al. (2012) for multimodal rigid and deformable registration. It is based on the concept of local self-similarity at the patch level. They create a multi-dimensional descriptor through ranking of the local intensity distribution of the two images, therefore providing a very good representation of the local shape of an image feature. It can be computed in a dense fashion across all the pixels (or voxels) of the images; once it is computed, the SSD of the MIND representations can be used as a similarity measure.

Thanks to its point-wise (pixel or voxel-wise) calculation nature, it can be adapted to almost any registration algorithm. An extension to MIND, named Self Similarity Context (SSC) is also estimated using patch-based self-similarities (Heinrich et al., 2013). In (Cifor et al., 2013), MIND was successfully used in a multi slice-to-volume registration framework to align untracked freehand 2D US sweeps to CT volumes.

A last family of robust similarity measures was introduced by (Rivaz et al., 2014c,b,a) in the context of mutlimodal US/MRI image registration. In (Rivaz et al., 2014c), a metric called Contextual Conditioned Mutual Information (CoCoMI) was proposed. The metric aims at tackling one of the main drawbacks of classic MI based methods, that is taking into account the intensity values of corresponding pixels and not of neighbor. Images are treated as “bag of words” and consequently contextual information is ignored. CoCoMI overcomes this limitation by conditioning the MI estimation on contextual information. In (Rivaz et al., 2014b), Self Similarity α -MI (SeSaMI) -another MI based matching criterion- is proposed. α -MI is usually calculated on multiple features like intensities and their gradients, as opposed to standard MI which is usually calculated on intensities only. SeSaMI combines this multi-feature α -MI formulation with self-similarities in a kNN α -MI registration framework by penalizing clusters (i.e. the nearest neighbors) that are not self-similar. Finally, in (Rivaz et al., 2014a), a CR similarity measure is introduced. Robust Patch Based Correlation Ratio (RaPTOR) computes local CR values on small patches and adds them to form a global cost function. Authors claim a property that makes suitable such methods: their metric is invariant to important spatial intensity inhomogeneity. This is especially useful when dealing with US images due to wave attenuation, shadowing and enhancement artifacts. One of the main advantages of these Rivaz’s metrics (CoCoMI, SeSaMI and RaPTOR) is that their gradient can be derived analytically, and therefore the cost function can be efficiently optimized using stochastic gradient descent methods.

2.2. Geometric

Geometric registration finds correspondences between meaningful anatomical locations or salient landmarks (Joshi and Miller, 2000; Glocker and Sotiras, 2011). These methods aim at minimizing an energy function that, for a given transformation, measures the discrepancy between the key-points detected in both, the 2D slice and the volumetric image. Simplicity of the registration process once the landmarks are appropriately determined, no sensitivity to initializations and a wider capture range in terms of deformation are the main strengths of such approaches. On the other hand, the landmark detection and matching processes are not that trivial, and errors on their position compromise the accuracy of the registration process. Moreover, due to the sparsity of the key-points, the quality of the deformations may become insufficient (due to the limited support on the interpolation).

The early work by Gourdon and Ayache (1994) presents a geometric method to perform slice-to-volume registration between a curve and a surface. In this work, Gourdon and Ayache exploited the knowledge about the differential properties computed on both, the curve and the surface, to constrain the rigid matching problem. The most relevant contribution of this work is their discussion about how differential constraints can be used to rigidly register a curve on a surface. However, they used the basic Marching Cubes algorithm (Lorenson

and Cline, 1987) to extract them from simulated and medical data. The extraction of these structures remains an arduous task for images with low resolution, a scenario often valid in slice-to-volume registration. In these cases, alternative methods (such as super-resolution reconstruction (Oktay et al., 2016) or registration-based interpolation methods (Penney et al., 2004b)) could be used to improve on the quality of the original images before extracting the structures of interest, leading to more accurate results after the geometric registration process is applied.

Extracting distinctive features becomes even more complicated for medical images (opposed to natural ones), since the first ones are usually not as discriminative (lack of texture) as the last ones. Invariance to scaling, rotation and changes in illumination or brightness constitute useful properties for methods seeking to extract salient points. Highly distinctive features (in both spatial and frequency domains) simplify the matching task since it is likely that they are correctly matched. Classical examples of such descriptors successfully applied in different computer vision tasks are the Scale-Invariant Feature Transform (SIFT) (Lowe, 1999), the Harris detector (Harris and Stephens, 1988), the Histogram of Oriented Gradients (HOG) (Dalal and Triggs, 2005), Speeded Up Robust Features (SURF) (Bay et al., 2008), etc. More recently, features learned using deep learning have been successfully applied for scenarios involving massive amount of annotated data for training (Long et al., 2014). These features are used to extract the salient landmarks that will guide the geometric registration process. Once the landmarks are established in both slice and volumetric images, there are two choices. First, we can perform a matching process which results in establishing correspondences between pairs of points from the two sets. Once the matching process is finished, correspondences can be used to estimate the desired -rigid or non-rigid- transformation models. Alternatively, methods directly estimating the transformation model without inferring any correspondence could also be applied. Sometimes, both correspondences and transformation model can be inferred at the same time, like in case of Iterative Closest Point (ICP) (Besl and McKay, 1992) algorithm. ICP is an algorithm subsequently improving the matching of point pairs. It minimizes the sum of geometric distances between the transformed set of source image landmarks and the closest detected landmarks in the target one. It is a simple and fast method that follows the closest neighbor principle to perform the matching task, which often converges to local minimum though.

Dalvi and Abugharbieh (2008) proposed a slice-to-volume registration approach which uses ICP. First, it extracts phase congruency information from the slices/volume using oriented 2D Gabor wavelets. Then, using local non maximum suppression, a robust and accurate set of feature points is automatically obtained, which are subsequently matched by ICP inferring a rigid body transformation. Nir et al. (2014) presented a particle filtering framework using a geometric matching criteria where the closest point principle is applied (they use the SSD between pairs of closest points in the two sets). In their proposal, the inferred transformation model involves 9-DOF (affine) matrices instead of 6-DOF transformations (rigid body). Nir's method was applied by Moradi and Mahdavi (2013) as a pre-processing step in a tumor detection learning based framework, to register US to whole-mount histopathology references of prostate images, mapping the location of tumors to the US image domain.

Another method that uses ICP was proposed by Yavariabdi et al. (2015) for deformable registration. They register 2D transvaginal US (TVUS) images with 3D MRI volumes to localize endometrial implants. They use contour to surface correspondences through a novel variational one-step deformable ICP method, finding a smooth deformation field while establishing point correspondences automatically. The main drawback of this approach is that it relies on the user to segment the organs. Manual pelvic organ segmentation is a laborious and time consuming task, and consequently the applicability of the approach in real scenarios is limited. Note that in the papers discussed in this paragraph, different types of transformation models were estimated based on the closest point principle, showing the flexibility of such approach.

Signed distance functions (or signed distance maps) can be used to avoid the landmark matching step when performing non-rigid registration of shapes and points. In this case, the landmarks or shapes are assigned to zero distance, while the rest of the pixels of the image are labeled with the distance to the nearest geometrical primitive (landmark or boundary). Once the distance map is created, the optimal transformation model can be estimated by means of standard iconic registration (e.g. using SSD on the distance maps)(Paragios et al., 2003; Xiaolei Huang et al., 2006; Taron et al., 2009). In (Boer et al., 2007), signed distance functions are used to register slices of histological images with a pre-reconstructed 3D model. Authors proposed to use a naive brute force approach for optimization, which incurs in extremely high computational cost.

2.3. Sensor Based Methods

Sensor based systems are an alternative to image-based methods (even if images are often used for their initialization), where slice location and orientation are continuously updated using sensors. Such information can be used to extract the corresponding slice from a given volume assuming an initial correspondence. Optical (OTS) and electromagnetic (EMTS) tracking systems are the most widely used technologies to perform this task. OTS systems determine in real-time the position of an object by tracking the positions of either active or passive markers attached to the object. It requires a line-of-sight to be maintained between the tracking device and the instrument to be tracked. This might cause an inconvenience to the physician -especially in image guided interventions- during his work, resulting highly inconvenient. EMTS systems usually consist of three basic components: the electromagnetic field generator, a system control unit that interfaces with a PC, and tracked sensor coils together with their interfaces to the system control unit. By measuring the behavior of each coil, the position and orientation of the object can be determined. EMTS do not inherit the same constraints as the line-of-sight requirements but it is sensitive to distortion from nearby metal sources and exhibits limited accuracy compared to optical tracking (Birkfellner et al., 2008). Moreover, neither OTS nor EMTS can deal with elastic deformations between volume and slice. In general, one can claim that sensor based methods are more reliable (when they are calibrated correctly) than image based methods since they are not affected by ambiguities that arise during the image interpretation process.

In (Gholipour et al., 2011), a three-dimensional magnetic field sensor is used to track the motion of a subject during MRI scanning. It allows estimating the location of the slices for

volume reconstruction. The rigid 6-DOF transformation of the sensor in three dimensions (6-DOF) is calculated in real-time using the native gradient fields of the MRI scanner. The relative 3D location of each slice is computed through the sensor motion parameters at the time of slice acquisition. Other works exploiting EMTS systems can be found here (Hummel et al., 2008; Olesch et al., 2011; Olesch and Fischer, 2011; Xu et al., 2008).

Optical tracking systems can also replace classical image based slice-to-volume registration algorithms. In (Schulz et al., 2014) a slice-by-slice motion correction for fMRI image reconstruction (see section 6.2 for a complete discussion about this problem) is achieved thanks to an optical tracking system. Authors proposed to use three tracking cameras with embedded image processing, to track the position of the three optical markers attached to the skull using goggles. Tracking information is then used to replace the classical slice-to-volume registration step (Kim et al., 1999) necessary to account for motion correction during fMRI image reconstruction. In (Bao et al., 2005), optical tracking is used to register laparoscopic US to CT images of a phantom liver, based on an infrared camera. The advantage of infrared cameras is that the position sensor does not pick up interference from reflections and ambient light. However, the line-of-sight requirement still holds and therefore sensor must not be blocked. Other papers using OTS to track the position of image slices when registering them with volumetric images can be found here (Huang et al., 2009; Penney et al., 2006, 2004a; Schulz et al., 2014).

More recently (Eresen et al., 2014), smartphone tracking technology was considered as a navigation tool to initialize a slice-to-volume registration process between a histological 2D slice and a MR volume. The Inertial Measurement Unit (IMU) of the smartphone (also available as a standalone component) is used to define the orientation of the slice. Given the orientation and position data of an IMU sensor, one can then interpolate the corresponding slice from the MR volume. This interactive alignment is applied to determine initial orientation for the 2D slice, and is refined using an iconic MI based registration process optimized via brute force. IMU tracking systems use a combination of accelerometers and gyroscopes to measure acceleration and angular velocity, respectively. Since acceleration is the second derivative of position with respect to time, and angular velocity is the first derivative, angular changes integrated over time from a known starting position yield translation and rotation (i.e. a 6-DOF transformation) (Birkfellner et al., 2008). These sensors are cheap and widely available. However, a major issue of using IMUs for tracking refers to accumulation of errors (either systematic or statistical), leading to degradation estimation over time. Kalman filters (Kalman, 1960) can be used to deal with this type of issues and have been often adopted from the scientific community. In Marami et al. (2016b,a) for example, Kalman filters were coupled with slice-to-volume registration to estimate head motion parameters.

Note that, as described in the next section, several papers that propose sensor based methods, are actually a combination of this technology with some image-based registration technique, resulting in what we call a hybrid method.

2.4. Hybrid

Sotiras et al. (2010) states that the hybrid matching criteria take advantage of both, iconic and geometric approaches, in an effort to get the best of both worlds. This category can

be extended by including sensor based technology as it can also be combined with iconic or geometric methods. In this case, sensors are used to initialize the registration process or to provide a real-time update for the slice position which can then be refined using image-based registration.

Mitra et al. (2012) proposed a hybrid slice-to-volume registration approach, combining geometric and iconic matching criteria in a probabilistic framework to register transrectal US (TRUS) with MR images of the prostate. The geometric component is based on shape-context (Belongie et al., 2002) representations of the segmented prostate contours. Bhattacharyya distance (Bhattachayya, 1943) between the shape-context histograms of the two shapes is used to find point to point correspondence in every axial MR image. The Chi-square distances between the TRUS slice and each of the MR slices are calculated and used to determine the matching slice. Once the TRUS-MR slice pair with the minimum Chi-square distance is determined, it is used to retrieve a 2D rigid transformation (in-plane rotation and translation) between them. This transformation is applied to the remaining MR slices to ensure similar 2D in-plane rigid alignment with the 2D TRUS slice. The iconic step is performed by measuring the similarity (NMI and CC are used as metric) between the TRUS image and every rotated axial slice of the MRI. Finally, shape and image similarity measures are transformed into probability density functions and mapped into a statistical similarity framework towards retrieving the MR slice that resembles the TRUS image. Two issues limit the applicability of such an approach in clinical practice. First, the registration is determined by the manual segmentations of the prostate on both images. Second, the method assumes that the TRUS slices are parallel to the MR axial planes; therefore, the result of the registration does not consider any out-of-plane rotation, which could occur in a realistic scenario.

Tracking information (coming from a EMTS or OTS system) can be combined in a hybrid approach with iconic or geometric criteria to perform slice-to-volume registration. In this case, the tracking signal could be used to initialize the process or calculate the additive position of a new 2D slice with respect to the previous one, bringing them to the same coordinate system. Iconic or geometric registration can then be incorporated to complement this process. In Penney et al. (2004a) for example, image-based registration was used to calculate the initial transformation that relates this unique coordinate system with the volumetric image, and sensor-based updates for the relative position of the slice where then used to perform the tracking. This approach was used by Penney et al. (2004a) to register intra-operative US images to a pre-operative MR volume. The algorithm was extended in (Penney et al., 2006) to deal with 2D US to volumetric CT registration. San José Estépar et al. (2009) refines the rigid tracking information provided by an EM system attached to a laparoscopic or endoscopic probe. The refinement is made in terms of translation parameters, using an edge-based iconic registration method combined with a phase correlation technique. In (Xu et al., 2008), EM tracking and intra-operative iconic image registration are used to superimpose MRI data on TRUS images of the prostate. To this end, a three-step algorithm is applied, with the intermediate one corresponding to slice-to-volume registration. At the beginning of the TRUS procedure, the operator performs a 2D tracked axial sweep from

the prostate's base to its apex, which is used to produce a volumetric US reconstruction of the prostate volume. The reconstructed US volume and the MRI are manually aligned. Thus, 2D TRUS slices are registered to the US reconstructed volume (reducing the MRI/US multimodality problem to US/US monomodality) while the estimated transformation can still be used to recover the corresponding MRI slice. During the intervention, out-of-plane motion compensation is achieved using intermittent multi slice-to-volume registrations between nearly real-time 2D US images and the 3D US image. This rigid body registration is based on minimizing SSD with a Simplex algorithm. The final step refers to a 2D-2D registration of US (using CC and gradient based similarity measures) and it seeks to account for in-plane misalignment. In this workflow, EM tracking information is used to simplify the problem from multi to monomodality registration, and combined with iconic slice-to-volume registration to account for motion compensation. Other hybrid methods that combine tracking information with iconic or geometric registration approaches are (Sun et al., 2007; Yan et al., 2012; Xiao et al., 2016).

In (Ghanavati et al., 2010) a method that combines the information from a tracked free-hand 2D US transducer with an iconic matching criterion based on US image simulation was proposed, in the context of Total Hip Replacement (THR) surgeries. They use a statistical shape model (SSM) of the pelvis, constructed from several CT images. They developed a multi-slice to volume registration method, to register multiple 2D US slices to a statistical atlas of the pelvis. The mean shape of the atlas is first registered rigidly to the freehand 2D US images, based on an US simulation method (Wein et al., 2008a). Then, the atlas is instantiated by measuring the iconic similarity between the actual US and the US images simulated from the instance of the atlas. This two-step algorithm estimating the rigid transformation and followed by the optimization of the deformable parameters, was turned into a one-shot optimized approach in (Ghanavati et al., 2011).

In (Wein et al., 2008b), a position sensor attached to the patient's skin was used to extract a scalar surrogate measurement, which represents an anterior-posterior translation used to detect and compensate for respiratory movements of the liver. This information is combined with an iconic criterion to estimate an affine 12-DOF transformation model which maps pre-operative plans and imaging into the interventional scenario.

Huang et al. (2009) proposed a hybrid slice-to-volume registration strategy exploiting multiple temporal frames. A multimodality image navigation system was introduced that integrates 2D US images with pre-operative cardiac CT volumes, using electrocardiogram (ECG) information, optical tracking and iconic MI based registration. The ECG is used to synchronize the US images to the corresponding pre-operative dynamic 3D CT, depending on the cardiac phase indicated by the ECG. Spatial information given by the optical tracker is used to produce a near-optimal starting point for every slice, that is refined by maximizing the MI similarity measure. Such an approach works in real-time, making it applicable to real interventional scenarios. A different combination of sensors to solve the registration problem was proposed in (Hummel et al., 2008), where EM and OTS sensors are combined with iconic registration to register endoscopic US to CT data.

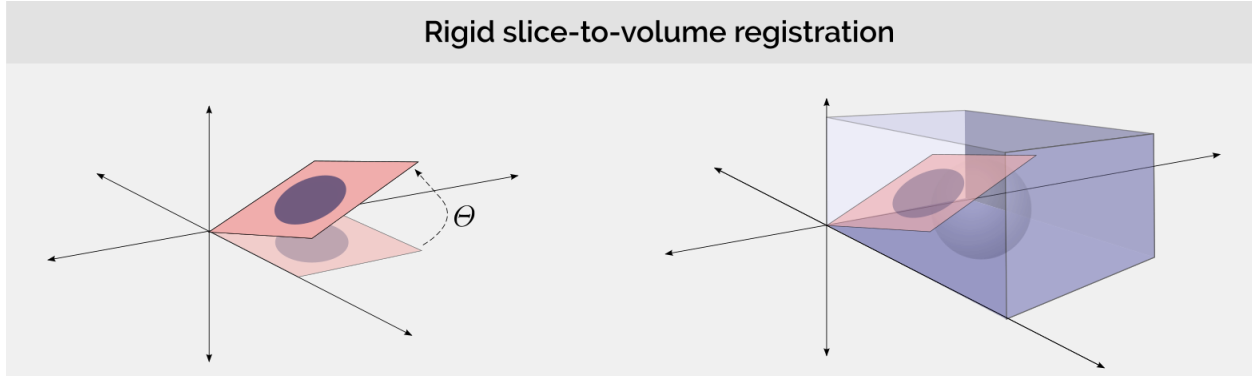


Figure 2: Rigid slice-to-volume registration. In this case, the transformation Θ parameterizes a simple 6-DOF transformation which can only model rotation and translation. Even if simple, this type of transformation has been largely adopted by the works reviewed in this survey (see Table 1)

3. Transformation Model

Transformation models explain the relation between the slice and the volume being registered, and are the outcome of the registration process. They are often classified according to their degrees of freedom. Rigid transformations deal with global rotations and translations, while deformable models -the complex case- can produce local in-plane and out-of-plane deformations (see Figure 3). The richness of the deformation model is proportional to the number of parameters we need to specify it, and therefore the trade-off is to be found between the model complexity and power of expression.

When a transformation is applied to a given image (at the final step of the registration process or during the intermediate optimization steps), an interpolation method must be used to infer the intensity value of every pixel in the space of the transformed image during resampling. In this regard, there are two alternatives to consider when it comes to slice-to-volume registration: to resample the slice or the volume using the estimated transformation. The most commonly adopted strategy consists in resampling the volumetric image, since resampling a slice in 3D is not as well defined as it is for a volume. However, in some special cases where the transformation model only allows in-plane deformations in the space of the 2D image (like (Ferrante and Paragios, 2013; Ferrante et al., 2015a,b)), slice resampling can be performed.

3.1. Rigid

The simplest transformation model accounts for rotation and translation parameters. It is usually expressed as a 6 degrees of freedom (6-DOF) transformation Θ^R composed by 3 rotation and 3 translation parameters. Such a basic model is the most common choice in the literature for slice-to-volume registration. Rigid transformations are expressive enough to explain simple slice-to-volume relations (see Figure 2). They can deal with in-plane and out-of-plane translations and rotations. Clinical scenarios that do not inherit image distortion -like simple inter-slice motion correction (Jiang et al., 2007b; Kim et al., 2008b; Rousseau and Glenn, 2005; Smolíková Wachowiak et al., 2005) or basic nature image guided surgeries

(San José Estépar et al., 2009; Birkfellner et al., 2007; Gill et al., 2008)- can be modeled with rigid transformations.

When out-of-plane motion is avoided, even simpler models can be used. Zakkaroff et al. (2012) proposed to recover in-plane slice rotations in cardiac MR series, using the stack alignment transform. In-plane translation along X and Y and rotation around a user-supplied center of rotation for the individual slices were parameterized independently. It also includes a parameter for global translation along the Z direction. Hence, it combines the individual slice transforms into a unique space of parameters -that contains 3-DOF per slice instead of 6- so that all of them can be optimized simultaneously. This type of models has been largely criticized suggesting that (especially in the context of volume reconstruction) the exact corresponding slice can only be found through a slice-to-volume registration which considers out-of-plane rotation and translations as well. In (Xiao et al., 2011), authors considered a slice-to-slice matching initial step, where every 2D image from a histological volume is matched to one axial MRI slice, via a group-wise MI based comparison. It implies a restricted transformation model involving simple slice correspondences (1-DOF). Subsequently, they correct the out-of-plane misalignment applying 2D-2D affine registration for each pair of matched slices, and a posterior 3D-3D affine registration between the histology pseudo-volume (reconstructed after transforming every 2D slice) and the MRI volume. Restricted rigid body transformations can be a convenient initialization component of a complete slice-to-volume registration pipeline.

6-DOF rigid body transformations are part of nearly all slice-to-volume registration algorithms. Literature seeking deformable registration, often initially employs rigid alignment to account for big range displacements before performing local deformable mapping (examples where this two-steps approach is applied can be found in (Tadayyon et al., 2011; Xu et al., 2014a; Fuerst et al., 2014; Hallack et al., 2015; Murgasova et al., 2016)). The standard way to estimate 6-DOF rigid transformations, consists in minimizing an energy functional (based on an iconic or geometric matching criterion) often with a continuous optimization algorithm (see section 4.1) where the search space is part of the Euclidean group $SE(3)$ of rigid transformations. Sensors as described in section 2.3 can be used as well for rigid alignment as proposed in a number of papers (Chandler et al., 2008; Elen et al., 2010; Fei et al., 2002, 2003a; Frühwald et al., 2009; Fuerst et al., 2014; Gholipour et al., 2010; Huang et al., 2009; Jiang et al., 2007a; Kainz et al., 2015; Kim et al., 1999, 2010a; Nir and Tannenbaum, 2011; Penney et al., 2004a; Rousseau et al., 2006; Smolíková Wachowiak et al., 2005; Xu et al., 2014b; Yeo et al., 2004; Yu et al., 2011; Chen et al., 2016).

3.2. *Non-Rigid*

In this review, we call non-rigid transformations to those models which perform -at some extent- changes in the structure of the images. These changes vary from simple operations -which can still be modeled using linear transformations, such as scale, flip or shear- to more complex models that produce local deformations.




Figure 3: Depending on the application, the slice-to-volume registration algorithm may allow for out-of-plane, in-plane or null deformations.

3.2.1. Linear Models

Linear transformations are the first step towards non-rigid registration. In (Wein et al., 2008a), authors suggested that affine models can eliminate most of the large-scale deformations when registering intra-operative 2D US slices to CT images. These deformations are particularly observed between CT and US exams in different respiratory configurations, and could be handled with affine transformations. However, in a different scenario where histological slices of the prostate are registered to a MRI volume, Nir and coworkers Nir et al. (2014) concluded that the affine transformation model was unable to capture deformations, especially in the rectum area. This is due to the fact that local deformations cannot be expressed from global affine models. In such cases, deformable registration is required (further described in the next section).

Recently, affine models were used for registering 2D slices and 3D images of the bones (Hoerth et al., 2015). Their proposal includes a novel way to initiate the registration between 2D slices and 3D images of the bones. The method uses the Generalized Hough Transform (GHT) (Ballard, 1981) to identify suitable starting positions. They create a template version of the 2D slice by thresholding its gradient vector field, used to explore the 3D volume through a GHT shape-detection process, providing a set of initial positions. Standard affine registration process is then used to deform images according to the initial configurations of the GHT space and updated using a MI based criterion. Such a principle could be considered to handle alternative transformation models. Other papers that estimate affine models to solve slice-to-volume registration are (Gefen et al., 2008; Micu et al., 2006; Museyko et al., 2014; Wein et al., 2008b; Xiao et al., 2011, 2016).

3.2.2. Deformable Models

Elastic deformations are powerful transformations with strong adoption in slice-to-volume mapping. The expressive power of these models depends on the number of parameters used to define them. An extensive description of different types and categories of deformable models used for non-rigid registration is presented in (Sotiras et al., 2013) and (Holden, 2008). Here, we focus on models applied for slice-to-volume mapping.

Thin-Plate Splines (TPS) are frequently used to generate a dense deformation field from a

sparse set of control points. These methods involve a set of control points that can be located in arbitrary positions, which are usually obtained by detecting salient structures. Radial basis function (RBF) -where the value at any interpolation point is given as a function of its distance from the control points- are centered at the control points and combined with affine terms to define an interpolation function. TPS minimize a bending energy based on this interpolation function, which gives a closed-form solution whose uniqueness is guaranteed in most cases. TPS can be decomposed into an affine and a local component. In (Osechinskiy and Kruggel, 2011a,b), TPS are used to parametrize a smooth 3D deformation of a 2D surface (slice). Control points are placed in a regular grid on the 2D image domain, and a 3D warp is defined using three independent TPS functions. Similarly, in (Miao et al., 2014), several regular 2D grids of control points define one TPS deformation model per slice in a multi slice-to-volume registration scenario. Kim et al. (2000) also applied TPS to a single slice-to-volume registration problem, and compared two variants of the elastic mapping (involving different number of control points to support the TPS model) with a rigid body registration algorithm. Authors extracted two main conclusions: (i) TPS based registration outperformed rigid body registration, at least in their multi slice fMRI scenario where local deformations are encountered due to the local field induced deformations or localized out-of-plane motion artifacts; and (ii) the number and location of the control points have a significant impact on the final results.

Free-form deformation (FFD) models are also widely applied for medical image registration. Originated in computer graphics (Sederberg and Parry, 1986), they became popular thanks to the seminal work by Rueckert et al. (1999). In this model, the weighting functions are cubic B-splines. The control points, have limited local support (in contrast to TPS where the control points influence the complete domain during interpolation) and are uniformly distributed over the image domain in a grid-like manner (Glocker and Sotiras, 2011). In (Ferrante and Paragios, 2013; Ferrante et al., 2015a,b), FFDs were used in a graph-based discrete optimization framework to perform slice-to-volume deformable mapping. In this model, a 2D grid-like graph encodes at the same time the plane position (rigid body transformation) and the in-plane deformation of a slice with respect to its corresponding position at the 3D. The in-plane deformation is obtained through a FFD interpolation, where the nodes of the graph play the role of control points. Similarly, but in a continuous setting, Osechinskiy and Kruggel (2011b) optimizes the position of the control points defined on a 2D grid, with unknown variables corresponding to 3D displacements. In a different scenario, Fuerst et al. (2014) proposed the use of a 3D grid based FFD within the bounding box of a tracked US sweep. Therefore, they estimated a single 3D deformation field out of a 3D grid, which can be used to deform all the slices contained in the US sweep. The main advantage of the FFD over TPS, is that it does not require solving a linear system for computing weights from displacements. However, FFD imposes constraints in terms of modeling, since the control points must be placed in a regular grid and must enclose the domain boundaries, whereas for TPS, control points can -in principle- be arbitrarily placed in the domain (Osechinskiy and Kruggel, 2011b). Other interesting works combining B-splines to model a deformable transformation in the context of slice-to-volume registration can be found here (Brooks et al.,

2008; Honal et al., 2010; Museyko et al., 2014; Po Su et al., 2013; Rivaz et al., 2014c,b,a; Tadayyon et al., 2010, 2011; Xu et al., 2014a).

An alternative model based on finite-element (FE) meshes has been used in (Marami et al., 2011) to model a dynamic linear elastic deformation field. It imposes a regularization constraint on permissible volume transformations based on an iconic image similarity criterion. The advantage of FE models is that they allow incorporating the dynamic behavior of tissue deformation into the registration process, through physically meaningful constraints.

4. Optimization Method

Optimization methods aim to determine the instance of the transformation model that minimizes a function based on the aforementioned matching criterion (see section 2). Depending on the nature of the variables being involved, those methods can be classified to continuous or discrete. The continuous approaches exploit the entire space of parameters, while the discrete ones a discretized/quantized version of the admissible solutions.

Numerous problems in computer vision and medical imaging are inherently discrete (like semantic segmentation); however, this is not the case of slice-to-volume image registration, where the search space is continuous. Most of the published methods about slice-to-volume registration adopt a continuous formulation. Nevertheless, recent works on image registration have focused on discrete formulations (Glocker and Sotiras, 2011; Glocker et al., 2008), both for projective and sliced 2D/3D registration (Zikic et al., 2010b,a; Ferrante and Paragios, 2013; Ferrante et al., 2015a,b). The next subsections present existing work and discuss the limitations of continuous and discrete inference methods in the context of slice-to-volume registration. We also consider a third category of heuristic methods, which are independent of the nature of the variables and can be applied to a wider range of problems at the expense of not providing optimality warranty. These methods are usually applied when finding an optimal solution is impossible or impractical given the nature of the problem objective function.

4.1. Continuous

Continuous optimization algorithms are usually iterative methods. They infer the best value for a set of parameters by iteratively updating them. A common mathematical formulation for this strategy is given by:

$$\Theta_{t+1} = \Theta_t + \omega \mathbf{d}_t \quad t = 0, 1, 2, 3, \dots \quad (2)$$

where Θ is the vector of parameters, \mathbf{d}_t is the search direction at iteration t and ω is the step size or gain factor. The search direction can be calculated using different strategies.

Continuous optimization methods can be classified according to different criteria (e.g. convex or concave minimization, solvers for linear or non-linear functions, smooth or non smooth problems, etc.). Here we simply classify them depending on whether they perform gradient- or non-gradient-based optimization, i.e., whether they exploit first (or higher order) derivatives of the energy function to compute the search direction, or they rely on other

strategies. An interesting comparative analysis where methods coming from both classes are used to solve slice-to-volume registration of histological slices with MR images is presented by Osechinskiy and Kruggel (2011b).

4.1.1. Gradient based Methods

Gradient-driven methods use the derivative of the objective function to guide the optimization process. In case of convex and differentiable functions, these methods are endowed with optimality guarantees. Otherwise, convergence to local minima is possible, and they could be sensible to initialization. Their main drawback is the requirement of analytical derivation or numerical estimation of the energy function derivatives, reducing their applicability since it is usually complicated.

Gradient descent is the simplest strategy in this category, where the search direction \mathbf{d}_t is given by the negative gradient of the energy function. It refers to the standard continuous optimization method, and it has been widely applied to the problem of slice-to-volume registration (Rousseau and Glenn, 2005; Rousseau et al., 2006; Bhagalia et al., 2009; Huang et al., 2009; Tadayyon et al., 2010, 2011; Kim et al., 2010b; Marami et al., 2011; Yu et al., 2011; Zakkaroff et al., 2012; Miao et al., 2014; Xu et al., 2014a).

Conjugate gradient methods use conjugate directions instead of the local gradient to estimate \mathbf{d}_t . Energy function with the shape of a long and narrow valley, can be optimized using fewer steps than standard gradient descent approach, resulting in faster convergence. (Elen et al., 2010; Osechinskiy and Kruggel, 2011b) have applied this strategy to estimate rigid and non-rigid slice-to-volume mapping functions, respectively.

Quasi-Newton optimization strategies are used when the Jacobian or Hessian of the energy function cannot be calculated or it is too expensive to be computed. In this case, the search direction at time t is computed based on an estimation of the Hessian, calculated using information provided by the previous iteration $t - 1$. Quasi-Newton optimization was applied to perform deformable registration of brain slices to MR images in (Kim, 2005), using several similarity measures. It was also used to register 2D intra-operative US images with pre-operative volumes in (Brooks et al., 2008). One of the most popular quasi-Newton methods is the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm used in (Honal et al., 2010) to correct for breathing motion artifacts during MRI acquisitions. A limited-memory version of BFGS (L-BFGS) is particularly suited for problems involving large numbers of variables and was applied to solve multi-slice to volume registration in (Fogtmann et al., 2014; Xu et al., 2014a).

Gauss-Newton methods can deal with non-linear least squares functions. An approximation of the Hessian matrix is used which, once combined with the gradient, provides a good estimate of the search direction for such functions. Thus, it only requires computing first order derivatives (in contrast to the standard Newton method where the actual Hessian matrix must be computed). In (Heldmann et al., 2010; Olesch et al., 2011; Olesch and Fischer, 2011), Gauss-Newton method is used for multimodal registration, where 2D US slices of the liver were registered to a pre-operative CT volume. Their matching criterion computes the Sum of Squared Differences (SSD) on the segmentation of the liver vessel structure in both, 3D and 2D slices. SSD is a least square function and therefore, Gauss-Newton

method can be efficiently applied. This method was also used in (Lasowski et al., 2008), solving an Iteratively Reweighted Least Squares (IRLS) process in order to estimate a rigid transformation between fluoroCT and CT images. The Levenberg-Marquardt algorithm is an alternative approach to solve non-linear least squares problems, that requires also only first order derivatives. For well-behaved smooth functions, this algorithm can take more time than standard Gauss-Newton method. However, the method is more robust than the standard Gauss-Newton, meaning that it could find a solution even with bad initializations. To the best of our knowledge, this optimization algorithm was only applied to estimate rigid body transformations; however, it has been used to optimize several iconic matching criteria in a variety of domains such as reconstruction of 3D cell images (Yu et al., 2011), endoscopic interventions (Hummel et al., 2008) and fetal brain imaging (Kim et al., 2008b, 2010a).

4.1.2. Non-Gradient based Methods

Derivative free methods eliminate the differentiability condition of the objective function. They are applicable to a wider range of functions, including noisy, non-differentiable or even unknown functions, where we only have a black-box that returns an output value given a set of input parameters. These are cases that frequently arise when dealing with medical image registration.

The simplest non-gradient based method that has been used to solve slice-to-volume registration is the local search or best neighbor method. In this case, each parameter to be estimated is perturbed in turn using the stepsize ω , and the value of the objective function is calculated. The search direction d_t is then estimated as the one that produced the greatest improvement in the objective function. The main drawback of this approach is its dependency from the choice of the initial stepsize ω . A big stepsize can result in an algorithm moving outside its capture range. On the other hand, for stepsizes that are too small, the optimization may become trapped in a local optimum (Penney et al., 2004a). Such a greedy algorithm has been applied to estimate rigid body and affine transformations mapping slice-to-volume, which have less DOF than deformable models. Different image modalities like US to MRI (Penney et al., 2004a) or CT (Penney et al., 2006; Sun et al., 2007) images, as well as fluoroCT images (Micu et al., 2006) were considered in the clinical setting.

Nelder-Mead (Nelder and Mead, 1965) (also known as downhill simplex method or amoeba method) is the most popular derivative-free method. It relies on the notion of simplex (a $n + 1$ polytope living in a n -dimensional space) to explore the space of solutions in a systematic way. At every iteration, the method constructs a simplex over the search surface, and the objective function is evaluated on its vertices. The algorithm moves across the surface by replacing, at every iteration, the worst vertex of the current set by a point reflected through the centroid of the remaining n points. The method can converge to a local optimum for objective functions that are smooth and unimodal. However, it exhibits more robust behavior for complex parameter space compared to standard gradient-based methods, providing a good compromise between robustness and convergence time (Leung et al., 2008). It has been widely used for slice-to-volume registration applications (Chen et al., 2016; Fei et al., 2002, 2003a,b, 2004; Frühwald et al., 2009; Gill et al., 2008; Hummel et al., 2008; Kim et al., 1999, 2000; Leung et al., 2008; Noble et al., 2005; Osechinskiy and Kruggel, 2011b;

Park et al., 2004; Wein et al., 2008a; Xiao et al., 2011; Xu et al., 2014b; Yeo et al., 2006, 2004; Xiao et al., 2016).

Another popular non-gradient based method is the well-known Powell's method. It explores the search space by performing bi-directional searches along N different vectors. Usually, these vectors initially refer to the canonical directions. Then, the search directions are updated using linear combinations of the earlier ones. The algorithm iterates until no significant improvement is made. Smolíková Wachowiak et al. (2005) applied Powell's method to register two dimensional cardiac images to preoperative 3D images. Authors claim that similarly to all local techniques, Powell's method converge to local minima, but it is generally robust and accurate, and exhibits fast convergence. (Wein et al., 2008b) takes advantage of the fact that Powell's method performs line search in specific directions. They apply Principal Component Analysis (PCA) on the 12 parameters of an affine transformation, reducing the search space to the 3 most significant PCA modes. Therefore, Powell direction search is initialized with the most significant PCA modes, which assures good performance and robustness. Other papers applying Powell's method to our problem can be found here (Fei et al., 2003a; Gholipour et al., 2010; Gholipour and Warfield, 2009; Jiang et al., 2007a,b, 2009; Micu et al., 2006; Osechinskiy and Kruggel, 2011b).

4.2. Discrete

In a discrete scenario, slice-to-volume registration can be expressed as a discrete labeling problem over a Markov Random Field (MRF) (Wang et al., 2013). Discrete optimization of MRFs is, in general, an NP-hard problem (Shimony, 1994). However, in special cases, it can benefit from very efficient solutions. The trivial brute force algorithm (i.e. trying all possible combination of labels for each and every variable) has an exponential complexity that makes such an approach unsuitable. More efficient algorithms have been developed during the last two decades which boosted the use of graphical models in the field of computer vision. They can be classified in three main categories according to Kappes et al. (2013):

1. Polyhedral and combinatorial methods, solving a continuous linear programming (LP) relaxation of the discrete energy minimization problem. The central idea is to relax the integrality condition of the variables in order to simplify the problem. Once the integrality constraint is relaxed, standard linear programming methods can be applied to solve the optimization problem, and rounding strategies are used to recover the integral solution. Examples of such approaches are Linear Programming Relaxations over the Local Polytope, Quadratic Pseudo Boolean Optimization (QPBO) (Rother et al., 2007) and Dual Decomposition (Komodakis et al., 2011).
2. Message passing methods, in which messages are calculated and propagated between nodes in a graph. This propagation can be seen as a re-parametrization of the original problem aiming to establish special properties in the re-weighted function that makes inference easier. Examples are the standard Loopy Belief Propagation (LBP) (Murphy et al., 1999) and Three Re-weighted Belief Propagation (TRBP) (Wainwright et al., 2005).

3. Max-flow and move-making algorithms make use of the well known max-flow min-cut (Boykov and Kolmogorov, 2004) algorithm from graph theory, which can optimally solve some instances of discrete energies. These methods are usually combined with greedy strategies that iteratively minimize over the label space by solving a sequence of max-flow min-cut sub problems. Examples are α -expansion, $\alpha\beta$ swap (Boykov et al., 2001) and FastPD (Komodakis and Tziritas, 2007) algorithms. Simpler move-making algorithms not using max-flow, but still applying the strategy of starting with an initial labeling and iteratively moving to a better one until a convergence criterion is met, are the classical Iterated conditional modes (ICM) (Besag, 1986) and its generalization Lazy Flipper (Andres et al., 2012).

Graphical models and discrete optimization are powerful formalisms that have been successfully used during the past years in the field of computer vision (Wang et al., 2013). In general, a graphical model is represented as a graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ where vertices in \mathcal{V} correspond to the variables while \mathcal{E} is a neighborhood system (pair-wise & higher order cliques) that encodes the relationships among these variables. In a discrete optimization problem, the aim is to assign a discrete label $l_v \in L$ to every variable $v \in \mathcal{V}$, by minimizing the following energy function:

$$\operatorname{argmin}_{l_p} \mathcal{P}(g, f) = \sum_{v \in \mathcal{V}} g_p(l_v) + \sum_{(v_1, v_2) \in \mathcal{E}} f_{v_1 v_2}(l_{v_1}, l_{v_2}) + \sum_{C_i \in \mathcal{E}} f_{v_1 \dots v_n}(l_{v_{i1}}, \dots, l_{v_{i|C_i|}}), \quad (3)$$

where $g_v(l_v)$ are the unary potentials, $f_{v_1 v_2}(l_{v_1}, l_{v_2})$ are the pairwise terms and $f_{v_1 \dots v_n}(l_{v_{i1}}, \dots, l_{v_{i|C_i|}})$ are high-order terms associated to high-order cliques C_i depending on more than two variables. Note that this equation can be seen as a discrete version of Equation 1, where the transformation mapping function Θ is parameterized through the labeling $\{l_p\}$ and both similarity and regularization terms are encoded in the energy terms.

Several models have been proposed based on the general formulation from Equation 3. A discrete approach to slice-to-volume rigid registration was recently proposed by Porchetto et al. (2016). Inspired by previous works on discrete estimation of linear transformations using graphical models (Zikic et al., 2010a), they formulate it through a fully-connected pairwise MRF, where the nodes are associated to the rigid parameters, and the edges encode the relation between the variables. Deformable image registration can also be formulated as a minimal cost graph problem where the nodes of the graph correspond to the control points of a deformation grid and the graph connectivity imposes regularization constraints. Even if this technique has been applied mainly to mono-dimensional cases (2D-2D or 3D-3D) (Glocker and Sotiras, 2011), several works were published during the last years extending this theory to the case of slice-to-volume registration (Ferrante and Paragios, 2013; Ferrante et al., 2015a,b; Ferrante and Paragios, 2015b).

The first solution to the problem of deformable slice-to-volume registration using graphical models and discrete optimization was presented by Ferrante and Paragios (2013). A regular grid, which models at the same time the plane location and the in-plane deformations, is interpreted as a graphical model where every grid point corresponds to a discrete

variable. A pairwise model which combines linear and deformable parameters within a coupled formulation on a 5-dimensional label space is used. The main advantage of such a model is the simplicity provided by its pairwise structure, while the main disadvantage is the dimensionality of the label space which makes inference computationally inefficient and approximate (limited sampling of search space). FastPD was adopted as the optimization algorithm given the properties of the energy function (pairwise terms and non-submodularity). Motivated by the work of Shekhovtsov et al. (2008), in an effort to reduce the dimensionality of the label space, the work by Ferrante et al. (2015b) presents a different model (the so-called decoupled model) where linear and deformable parameters are now separated into two interconnected subgraphs which refer to lower dimensional label spaces. It reduces the dimensionality of the label space by increasing the number of edges and vertices, while keeping a pairwise graph. In this case, loopy belief propagation is used as optimization algorithm, since FastPD requires in general an equal number of labels for all nodes, which is an issue in their setting given the different dimensionality of the label spaces. Finally, in (Ferrante et al., 2015a), a high-order approach is presented, where the label space dimensionality reduction is achieved by augmenting the order of the graphical model, using third-order cliques which exploits the expression power of this type of variable interactions. Such a model provides better satisfaction of the global deformation constraints at the expense of quite challenging inference. Loopy belief propagation is also used as the optimization algorithm in this case.

Discrete methods have several advantages when compared with continuous approaches for image registration. First, they are inherently gradient-free, while most part of continuous methods require the objective function to be differentiable. Second, continuous methods are quite often prone to be stuck in local minima when the functions are not convex. However, in case of discrete methods, even complicated functions could potentially be optimized using large neighbor search methods. Third, parallel architectures can be used to perform non-sequential tasks required by several discrete algorithms (such as message calculation in LBP) leading to more efficient implementations. Fourth, by using a discrete label space, one can explicitly control its range and resolution, what can be useful to introduce prior information. The main limitation of discrete methods when compared to continuous approaches is the accuracy, which is bounded by the discretization of the continuous space. However, as suggested by Glocker (2010), even if the optimality is bounded by the discretization, with intelligent refinement strategy the accuracy of continuous methods can be achieved.

4.3. Miscellaneous

Independently of whether the variables are continuous or discrete, different strategies can be used to explore the space of solutions. Heuristics or metaheuristics can obtain acceptable solutions especially in cases where we deal with non-linear, non-convex or black-box optimization problems, even if we know that no optimality guarantee is provided. In some cases (e.g. initializing a more complex registration procedure) having an approximately good solution is enough; in these scenarios, different strategies can be envisioned.

Evolutionary algorithms are among the most popular bio-inspired metaheuristics. In evolutionary computation, the concept of biological evolution is used to explore the search space. The individuals of a given population are the candidate solutions to our problem

and they evolve according to different laws (such as mutation, recombination and selection). Evaluation of the quality for a given solution is performed through a fitness function, which is in fact our objective function. The idea is that, over the generation sequence, individuals with better and better fitness are generated, leading to a good solution. These algorithms are usually stochastic. If this is the case, new candidate solutions might be sampled according to a multivariate normal distribution in the space of parameters. Genetic algorithms were used by Lin et al. (2013) to estimate a 6-DOF transformation between 2D MRI real-time slices of the knee and 3D MRI volumes. One evolutionary algorithm applied in several slice-to-volume registration problems is the covariance matrix adaptation evolutionary strategy (CMA-ES) (Ghanavati et al., 2010; Tadayyon et al., 2010, 2011). This method updates the covariance matrix of the aforementioned distribution, by estimating a second order model of the underlying objective function similar to the approximation of the inverse Hessian matrix. Another case based on genetic algorithms to solve our problem of interest was proposed by (Gefen et al., 2005, 2008). Different optimization strategies may be combined as well depending on the parameter space. In (Tadayyon et al., 2010, 2011), for example, the CMA-ES was applied to optimize the translation parameters of the rigid transformation model, while rotation and deformable ones were estimated using a standard gradient descent approach. In their application to prostate MR images, the CMA-ES was not able to optimize a 6-DOF search space as it diverged on rotations regardless of scaling.

Another popular metaheuristic is the well known simulated annealing (SA) method. It mimics the physical process that metal atoms suffer when the material is heated and then slowly cooled. In order to avoid local minima, the algorithm explores new directions which lower the objective function, but also, with a certain probability, those that raise the objective. This probability decreases with the number of iterations, so that it stays away from local minima in early iterations and is able to explore globally for better solutions. Birkfellner et al. (2007) used this method to register fluroCT slice to CT volumes optimizing several standard iconic matching criteria, while Cifor et al. (2013) applied it to a multi-slice to volume registration case where a robust modality-independent similarity measure was optimized.

5. Number of Slices

In this review we include single and multi slice-to-volume registration methods. While in the first case the estimated slice-to-volume mapping function considers just a single 2D slice, in the second case several slices are mapped to the 3D volume.

A trivial extension of any single slice-to-volume registration method to the multi-slice case could be achieved by simply applying the algorithm to every input slice independently. It would allow parallelizing the process -since every registration could be performed in parallel- but, at the same time, contextual information would be lost since no relation among the input slices would be considered. Several methods of this type have been proposed. We can cite (Fei et al., 2003a; Ferrante et al., 2015a; Osechinskiy and Kruggel, 2011b; Yavariabdi et al., 2015) among others.

Different strategies can be adopted to register multiple slices to a given volume. Here we only review methods that perform a consistent transformation to the set of input slices, but still conceive the process as slice-based without seeing them as a complete and unique volume. Methods that reconstruct a new volume using the input slices, and then perform a standard 3D-3D registration process, are not taken into account. Note that, in some works like (Brooks et al., 2008; Yan et al., 2012; Xu et al., 2014a), experiments have been performed to compare the performance of multi slice-to-volume versus volume-to-volume registration. In (Xu et al., 2014a), the comparison was done in the context of MR prostate images, in order to show that multi slice-to-volume registration is sufficient in capturing prostate motion intra-operatively. The conclusion was that multi slice-to-volume registration was able to produce results that were close enough to volume-to-volume registration. This is an encouraging result, considering the data computation reduction achieved when using only a few slices instead of a complete volume. Brooks et al. (2008) and Yan et al. (2012) arrived to similar conclusions in different medical scenarios. This consensus indicates the reliability of using sparse slices instead of the full volume when possible. Several advantages are associated to this strategy: (i) the computational requirements are lower thanks to the reduction in the amount of data to be processed (only a few slices vs a complete volume), (ii) the omission of the potentially complex reconstruction step, (iii) greater adaptability of the similarity measures and (iv) easier parallelization (Brooks et al., 2008).

In the presence of multiple slices, multi slice-to-volume procedure -as opposed to single slice procedure- greatly increases the quality of the matching leading to more robust registration methods (Yu et al., 2011). This is even more interesting when the relative position between the slices is known. In (Chandler et al., 2006, 2008), for example, multi slice-to-volume registration was used to correct for misaligned cardiac anatomy in Short Axis (SA) images by registering stacks of two parallel slices (which are supposed to be aligned between them) to a high-resolution 3D MR axial cardiac volume. In other cases, different configurations like orthogonal slices (Gill et al., 2008; Leung et al., 2008; Tadayyon et al., 2011; Miao et al., 2014; Xu et al., 2014a), slices aligned in arbitrary positions (obtained using a tracking system) (Wein et al., 2008a; Heldmann and Papenberg, 2009; Heldmann et al., 2010; Olesch et al., 2011; Olesch and Fischer, 2011; Cifor et al., 2013) or even temporal sequences of slices (Miao et al., 2014), are registered to a volume. Another case is Yu et al. (2008, 2011), where a bi-protocol was proposed to reconstruct a microscopic volume of a cell. The bi-protocol is composed by two sets of multiple slices imaging the same cell, which are captured with two different geometries. Both sets of slices are then registered to each other and a final volume is reconstructed using a simple interpolation strategy.

Let us recall that multi slice-to-volume registration involves lower computational complexity compared to 3D-3D registration methods. Furthermore, when compared to single slice methods, multi slice-to-volume registration certainly improves the robustness of the registration process by augmenting the image support. Dealing with several slices requires more computing power than the single slice case; however, on occasions like freehand US sweeps, slices contain redundant information that could be avoided. In that sense, as the complexity of the scheme is proportional to the number of input slices, Wein et al. (2008a)

proposed a strategy which starts out by selecting only a few key-slices. These are used to estimate a rigid or affine transformation model, mapping the US slices to a CT volume. Since neighboring frames of the freehand US sweep contain overlapping information, only a few key slices are selected. They chose the most informative slices by selecting those with the highest image entropy. Olesch et al. (2011) applied the same key-slice selection technique to the variational deformable registration framework proposed by Heldmann and Papenberg (2009). To this end, well distributed slices throughout the volume were considered, containing meaningful information in terms of entropy. In another work, Olesch and Fischer (2011) extended this key-slice selection technique by introducing a focused registration strategy that only considers slices which are in a given region of interest (ROI).

In (Jiang et al., 2007b), a multi slice-to-volume registration method -entitled Snapshot magnetic resonance imaging with Volume Reconstruction (SVR)- is proposed to deal with reconstruction of MRI of moving subjects. After imaging the target object of interest repeatedly (producing multiple overlapped stacks of slices of the same moving object) to guarantee sufficient sampling, one of the stacks is chosen. Then, an iterative multi slice-to-volume registration strategy is applied where the stacks are subsequently registered (in a rigid way) to the reference. Once all the stacks are registered, they are averaged to produce a new reference. Then, the number of temporally contiguous slices in every stack is reduced so that the data is divided into subpackages that are temporally contiguous although the slices in each subpackage may not be spatially adjacent. The process is then repeated: the subpackages are registered to the average image, which is again updated, and the number of slices per stack is once more decreased; this process is repeated until each slice is treated in isolation (reducing to the single slice-to-volume case) and the final average volume is reconstructed. In (Jiang et al., 2007a, 2009) the idea was extended to allow the reconstruction of diffusion tensor images in the same scenario. In the same work, Jiang et al. (2007b) tested the benefit of using both, parallel and orthogonally acquired slices with prospectively acquired data and simulated cases as well. More recently, a fast multi-GPU accelerated framework to perform SVR was presented (Kainz et al., 2015). They proposed a fully parallel SVR approach for the reconstruction of high-resolution volumetric data from motion corrupted stacks of images. Parallelization is performed at two levels: (i) at the slice level, multiple slices are treated separately for large parts of the reconstruction process and (ii) since pixel/voxel based operations are independent of each other, they can also be executed in parallel. Authors claimed that their approach is five to ten times faster than the fastest currently available multi-CPU frameworks.

Rousseau and Glenn (2005); Rousseau et al. (2006) proposed the registration of multiple sets of orthogonal 2D MRI slices into a high resolution MRI volume. The work of Jiang et al. (2007b) shares some similarities with them. However, instead of reducing the number of slices per subpackage until each slice is treated in isolation, a two-step approach is performed. In this setting, a global alignment of the low resolution images is followed by the registration of each slice of the low resolution images to the reconstructed high resolution volume. Only orthogonal sets of slices were considered for reconstruction.

The work by Kim et al. (2008b, 2010a) (built on top of (Rousseau and Glenn, 2005;

Rousseau et al., 2006)), proposed a new approach entitled Slice Intersection Motion Correction (SIMC). It considers the registration process directly in terms of the intersections of each pair of slices in the stacks, avoiding the intermediate volume estimation process. Independent per-slice rigid transformations are estimated by minimizing a global energy function defined by the sum of dissimilarity measures of all intersection profiles between any two orthogonal slices. SIMC can be used in this scenario, considering intersecting lines of voxels instead of standard patch or global based calculation. Kim et al. (2008b, 2010a) applied SIMC to the problem of reconstructing MR images of a moving human fetal brain. Since then, several extensions to this work have been published, like for example the work presented in (Kim et al., 2010b) expanding the method to the case of Diffusion Tensor Images and (Kio Kim et al., 2010; Kim et al., 2011) which modified it in order to account for bias field inconsistency correction on fetal brain MR images.

An approach based on particle filtering (PF) was considered in (Nir and Tannenbaum, 2011; Nir et al., 2014) to deal with multi slice-to-volume registration of histological images to MR or US volumes. PF framework derives an optimal estimation of the parameters in a Bayesian fashion, tackling two of the main issues that arise from multimodal registration, such as susceptibility to initialization and optimal solutions. PF framework was applied to the problem of head motion tracking in the context of EPI based fMRI motion correction by Chen et al. (2016). Chen and coworkers used this strategy to couple successive EPI slices during the registration process, resulting in improved performance when compared with standard volume-to-volume and single slice-to-volume registration approaches.

6. Applications

A broad number of clinical settings and applications can benefit from slice-to-volume registration. In this review we classify them in two main categories.

6.1. Image Fusion and Image Guided Interventions (IGI)

Several medical procedures such as image guided surgeries and therapies (Fei et al., 2002), biopsies (Xu et al., 2014a), radio frequency ablation (Xu et al., 2013), tracking of particular organs (Gill et al., 2008) and minimally-invasive procedures (Liao et al., 2013; Huang et al., 2009) belong to this category. In this context, slice-to-volume registration brings high resolution annotated data into the operating room. Generally, pre-operative 3D images such as computed tomography (CT) or magnetic resonance images (MRI) are acquired for diagnosis and manually annotated by expert physicians. During the surgical procedure, 2D real time images are generated using different technologies (e.g. fluoroCT, US or interventional MRI slices). The alignment of intra-operative images with pre-operative volumes augments the information that physicians have access to, and allows them to navigate the volumetric annotation while performing the operation (see Figure 4). These intra-operative images inherit lower resolution and quality than the pre-operative ones. Moreover, tissue shift collapse as well as breathing and heart motion during the procedure, cause elastic deformation in the images, what makes slice-to-volume registration an extremely challenging task. A statistically significant improvement in alignment has been demonstrated when comparing automatic

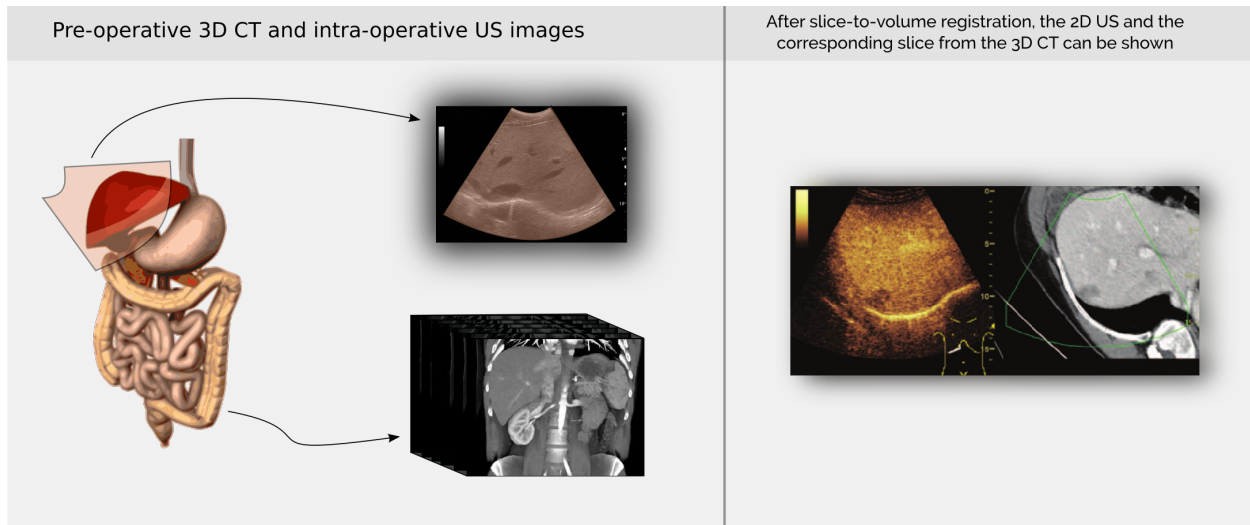


Figure 4: One of the main applications requiring slice-to-volume registration is image fusion for image guided procedures. In the figure we can observe an example where an intra-operative US slice is registered to a pre-operative CT image. Images are then shown side-by-side, providing complementary information (The side-by-side US/CT image was extracted from Ewertzen et al. (2013)).

methods to manual (human) results, showing the importance of automatic slice-to-volume registration algorithms in the context of image fusion and IGI (Frühwald et al., 2009).

The pioneering work of Fei and coworkers (Fei et al., 2001, 2002, 2003a,c, 2004) introduced iconic slice-to-volume registration to the problem of image fusion in the context of image guided surgeries. The motivation was that low-resolution Single Photon Emission Computed Tomography (SPECT) can be brought to operating by pre-registering it with a high-resolution MRI volume, which could be subsequently fused with live-time iMRI. That is how, by registering the high-resolution MR image with live-time iMRI acquisitions, Fei and coworkers could map the functional data and high-resolution anatomic information to live-time iMRI images for improved tumor targeting during thermal ablation.

Birkfellner et al. (2007) used slice-to-volume registration to fuse 2D fluoroCT with volumetric CT, which is a well know tool for image-guided biopsies in interventional radiology. In this case, the pre-interventional diagnostic high resolution CT with contrast agent is used to localize a lesion in the liver. However, during the intervention, the lesion is no longer visible. Thus, localizing the slice of the CT that corresponds to the intra-operative fluoroCT allows doctors to find the lesions during the biopsy. This approach only considers rigid transformations. However, interventional procedures like radio frequency ablation (RFA) or image-guided biopsies, which use fluoroCT as image guiding technology, are performed while the patient is breathing continuously. Therefore, deformations should also be taken into account when registering with the pre-operative static CT image. The influence of such deformations and the reliability of performing non-rigid registration in such scenario was discussed in (Micu et al., 2006; Lasowski et al., 2008). It was claimed that a 2D-3D non-rigid registration solution -based on the single low quality fluoroCT- cannot be precise as required to performed medical procedures. This is mainly due to the poor support in terms

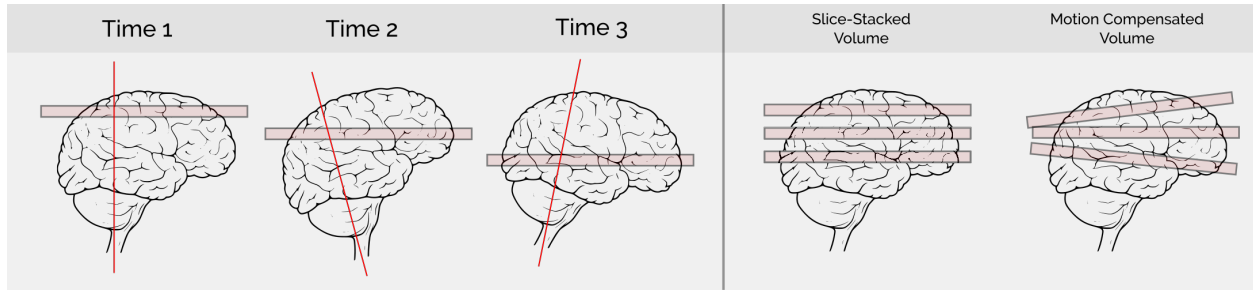


Figure 5: Slice-to-volume registration is a key task when performing motion correction for volume reconstruction. In the figure we can observe a typical case, where the patient moves while the 3D image is being acquired slice-by-slice, resulting in a corrupted slice-stacked volume. Motion correction through slice-to-volume registration is a common technique to improve the quality of the reconstructed volume (this figure is based on Figure 1 from Chen et al. (2016)).

of liver anatomical features (mainly vessels) provided by the fluoroCT slices. They proposed to overcome this limitation by providing an adaptive visualization (Lasowski et al., 2008) of the volume area surrounding the minimum estimated pose. This approach copes with the uncertainty in estimating the deformation and brings more information than a single registered slice. Their method performs rigid slice-to-volume registration, and includes views of the CT-Volume determined along flat directions of the out-of-plane motion parameters next to the minimum pose.

Laparoscopic and endoscopic interventional procedures also exploit slice-to-volume registration. San José Estépar et al. (2009), proposed a method to register endoscopic and laparoscopic US images with pre-operative computed tomography volumes in real time. It is based on a new phase correlation technique called LEPART accounting for rigid registration. Other methods applying slice-to-volume registration on minimally invasive procedures can be found here (Heldmann et al., 2010; Bao et al., 2005).

6.2. Motion Correction and Volume Reconstruction

The second category is motion correction and volume reconstruction. Here, the goal is to correct misaligned slices when reconstructing a volume of a certain modality (see Figure 5). A typical approach to solve this task consists in mapping individual slices within a volume onto another reference volume in order to correct the inter-slice misalignment. The popular map-slice-to-volume (MSV) method introduced this idea (Kim et al., 1999). More recently, applications of slice-to-volume registration to the same problem in different contexts like cardiac magnetic resonance (CMR) (Chandler et al., 2008; Elen et al., 2010), fetal images (Seshamani et al., 2013) and diffusion tensor imaging (DTI) (Jiang et al., 2009) have shown promising results. In these problems, it is usually assumed that a single slice is coherent, in the sense that spatial inconsistency only happens at the inter-slice level.

Slice motion correction in the context of volume reconstruction, typically involves iterative registration of slices to a target volume. The target volume may be an anatomical reference or it could be reconstructed at each iteration using current estimates of slice motion, considering all possible views of the subject.

Kim and coworkers (Kim et al., 1999) introduced the map-slice-to-volume (MSV) approach for inter-slice motion correction in time series of fMRI image. In such cases, head motion represents the major source of error in measuring intensity changes related to given stimuli in fMRI time series (Yeo et al., 2004). The aim of MSV is to retrospectively remap slices that are shifted by head motion to their spatially correct location using an anatomical MRI volume as reference. The MSV approach estimates 6-DOF rigid transformations independently for every fMRI slice, by minimizing the energy based on mutual information, using a Nelder-Mead downhill simplex method for optimization. This method was presented as an alternative to the slice-stack approach. Instead of considering slice-wise registration, it assumes stacks of slices being already registered among them, and ignores the inter-slice motion inherent to multi slice echo planar imaging (EPI) acquisition sequence (since each slice is excited at a sequential time interval). The MSV showed better performance than previous volume-to-volume registration methods, but the reliability of the estimated position parameters for the end cap slices was low due to the limited information support of smaller regions (less textured area and more background). Park et al. (2004) proposed to overcome this limitation using Joint Mapping of Slices into Volume (JMSV). JMSV is a multi-slice registration method that jointly estimates a rigid body transformations per slice, while penalizing the implied acceleration in the motion trajectory of the subject -i.e. abrupt changes in the motion parameters of sequentially acquired slices. Other extensions to the standard MSV include (i) accounting for deformable registration (Kim et al., 2000) through TPS transformation models; (ii) improving motion correction capability of MSV with concurrent iterative field-corrected reconstruction (Yeo et al., 2004, 2008); (iii) extensive evaluation of the activation detection performance of MSV and effects of temporal filtering of motion parameter estimates (Yeo et al., 2006) and spin saturation effect (Kim et al., 2008a); (iv) improving the MI matching criterion estimation in the low-complexity end-slices (slices near top or bottom of the head scans, presenting poorer information) by incorporating joint probability density functions of image intensities estimated from successfully registered center-slices in the same time-series.

After the fundamental work of Kim et al. (1999) and its extensions (most of them related to fMRI image reconstruction), another problem requiring to correct for inter-slice motion started to attract attention from the community: fetal brain MR imaging. In this case, there is no anatomical reference available to be used as a target volume; therefore, a reference volume is calculated during the registration process using the current estimate of the slices. Jiang et al. (2007a) and Rousseau et al. (2006) established the basis of a new family of methods which rely on iterations of slice-to-volume registration and scattered data interpolation (SDI) to perform super-resolution reconstruction of moving objects -in particular, fetal brain MRI-. Both approaches share a similar and iterative slice-to-volume registration scheme, but differ on the SDI method: while Rousseau et al. (2006) use Gaussian kernel-based SDI, Jiang et al. (2007a) relied on a regular grid of control-points that acts as a cubic B-spline to perform SDI. In a more recent work, Gholipour and Warfield (2009); Gholipour et al. (2010) criticized SDI-based approaches stating that none of these techniques can guarantee the convergence of the reconstruction to, at least, a local optimal solution. Thus,

they proposed the use of a maximum likelihood (ML) error minimization approach, which guarantees the convergence of the reconstructed image to match the motion-corrected slices through steepest descent error minimization. An extension to this method was presented in (Gholipour et al., 2011), using a real-time sensor based tracking (i.e. a non-image based approach) to estimate an initial position of every slice, which was then refined through a retrospective slice-to-volume registration approach based on ML error minimization.

7. Validation of Slice-to-Volume Registration Methods

A main issue identified during this study is the lack of open/public benchmarks with gold standard annotations specially conceived to validate slice-to-volume registration methods. This type of datasets, generally available for most medical image analysis problems¹, enable fair comparison of new approaches with the state-of-the-art, leading to more constructive contributions. One of the few public datasets used in several slice-to-volume registration studies (Rivaz et al., 2014c,b,a; Ferrante and Paragios, 2013; Pardasani et al., 2016) is the MNI BITE (Montreal Neurological Institute’s Brain Images of Tumors for Evaluation database) introduced by Mercier et al. (2012). It contains 2D US, 3D US and MRI brain images of several patients, with ground truth provided in form of anatomical landmarks. Even if this dataset was used in several studies, it was not specially designed to validate slice-to-volume registration methods, particularly not in case of single slice methods. As stated by Pardasani et al. (2016), one caveat is that these landmarks were identified in the reconstructed US volumes (not in the independent slices), and therefore different workarounds (see Section 7.2.2) have to be used to adapt it to single slice-to-volume registration validation (where we need homologous annotations in both, 2D and 3D images).

New validation datasets specifically designed to evaluate slice-to-volume registration methods should include annotations in form of anatomical landmarks, manual segmentation of structures of interest and gold standard transformations mapping slice-to-volume in both, 2D and 3D images. Public datasets for benchmarking proved to be essential to build communities around vision problems (like Pascal VOC dataset (Everingham et al., 2010) for object class recognition or IBSR dataset (IBSR) for brain image segmentation) and they would certainly contribute to make slice-to-volume registration a well established problem.

In what follows, we analyze the standard approaches that have been used to validate the methods considered in this review, discussing advantages and disadvantages in the context of slice-to-volume registration.

7.1. Type of images

There are two main type of images used for validation: phantoms and clinical images. Phantoms are artificial objects which can be scanned under controlled conditions to simulate the acquisition of real images of a patient. They are usually designed to mimic certain

¹An extensive list of datasets for benchmarking several biomedical image analysis tasks can be found in the Grand Challenge website (<https://grand-challenge.org/>). Note that none of these datasets are specially conceived for slice-to-volume registration.

properties of human or animal tissue, and used to calibrate the parameters of medical imaging devices or for benchmarking purposes. Not only artificial anatomical structures (like breast (Marami et al., 2011), heart (Huang et al., 2009), abdominal structures (Yang et al., 2015; Birkfellner et al., 2007), lumbar spine (Yan et al., 2012), pelvic organs (Fei et al., 2004), etc) but also other type of non conventional objects (like Lego bricks (Xiao et al., 2016) or even pineapples (Gholipour et al., 2011)) have been considered as phantoms to be imaged in the context of slice-to-volume registration. Moreover, digitally simulated phantoms (like the BrainWeb (Cocosco et al., 1997) considered in Gholipour et al. (2010) or the well-known analytical Shepp-Logan phantom (Shepp and Logan, 1974) adopted by Kim et al. (2011)) have also been used. One of the main advantages of digitally simulated phantoms is that, since we known in detail how the image was generated, we can provide extra information like segmentation masks or probability maps for certain structures of interest, which can then be used for validation as we will see in Section 7.2. The other type of images that can be considered are clinical studies. This is the most common and valuable case, where images of real patients are captured and used to validate a registration method.

Most of the methods considered in this review use clinical images for validation, since it proves the potential of the method to be applied in real clinical context. However, it must be noted that one of the main drawbacks of clinical datasets is the lack of annotations. Finding annotations for clinical images is harder than in case of phantoms, since we usually require specialists who can manually annotate them. Therefore, a common strategy followed by several authors like Xu et al. (2008); Huang et al. (2009); San José Estépar et al. (2009); Kim et al. (2011); Gholipour et al. (2011); Schulz et al. (2014); Kainz et al. (2015) is to combine phantom and clinical datasets, in a effort to get the best of both worlds: accurate measurements based on phantom data, and realistic validation using clinical data.

7.2. Measuring the Performance of Slice-to-Volume Registration Methods

Alternative indicators have been considered to quantify the performance of slice-to-volume registration. While some of them measure global properties (like the estimated parameters in rigid registration) other methods focus on local aspects (for example, the error in specific anatomical landmarks or local differences in terms of image intensity before and after registration). Table 3 includes a summary of the most relevant metrics used to validate the methods included in this survey. Here we discussed them in detail.

7.2.1. Distance Between Transformation Parameters

Given a single (or multiple) slices and a volume, if the transformation Θ_{GT} that maps both images is known, we can estimate the distance between Θ_{GT} and the estimated transformation $\hat{\Theta}$. This approach is mostly used to validate global linear transformations (see for example Kim et al. (1999); Park et al. (2004); Dalvi and Abugharbieh (2008); Ferrante and Paragios (2013); Lin et al. (2013)) where the number of parameters to estimate is low and distances per parameter $|\Theta_{GT}^i - \hat{\Theta}^i|$ can be reported. However, in Po Su et al. (2013), the same strategy is applied to measure the accuracy of the estimated deformation field. In that context, they simply consider the average of the euclidean distances over the components of the deformation field in every pixel, given by $\frac{1}{N} \sum_{x \in \Omega} \|d_{GT}(x) - \hat{d}(x)\|$ where $d_{GT}(x), \hat{d}(x)$

are the corresponding displacement vectors of the ground-truth and estimated deformation fields in the position x respectively, N is the total number of pixels and Ω is the image domain.

Note that this strategy requires to know a priori the exact transformation (rigid or non-rigid) Θ_{GT} mapping slice-to-volume. In both phantom and clinical scenarios, this is rarely the case. To overcome this limitation, a common strategy applied in several works (e.g. Kim et al. (1999); Park et al. (2004); Yeo et al. (2004); Smolíková Wachowiak et al. (2005); Yeo et al. (2006); Birkfellner et al. (2007); Dalvi and Abugharbieh (2008); Gefen et al. (2008); Kim et al. (2008a); Yeo et al. (2008); Elen et al. (2010); Gholipour et al. (2010); Zakkaroff et al. (2012); Ferrante and Paragios (2013); Ferrante et al. (2015b,a); Porchetto et al. (2016)) consists in generating synthetic ground truth by extracting arbitrary slices from known transformations given a 3D image. These parameters are then perturbed to initialize the registration process, and used as ground-truth to calculate the error with respect to the estimated transformation. Another alternative is to generate bronze-standard annotations. In this case, the idea is to use more reliable registration methods when possible (e.g. volume-to-volume registration when multiple slices are available, or landmark based registration when accurate anatomical landmarks are provided) to obtain the transformations, which are then considered as ground-truth. Alternative methods to produce bronze standard annotations for slice-to-volume registration can be found in Fei et al. (2002, 2003a, 2004, 2003c); Penney et al. (2006); Jiang et al. (2009).

7.2.2. Point-based Registration Error

Another common strategy frequently used in the literature is based on landmarks. The idea is to annotate points of interest which are visible in both, the slice and the volume images, so that we can measure the distance between the corresponding points before and after registration. The mean distance between the ground truth and registered anatomical landmarks is commonly referred as target registration error (TRE), and calculated as:

$$\sqrt{\frac{1}{M} \sum_{k=1}^M \| p_{GT}^k - T_{\hat{\Theta}} \circ p^k \|^2}, \quad (4)$$

where p^k, p_{GT}^k are the corresponding landmarks in the source and target domains, M is the number of landmarks and $T_{\hat{\Theta}}$ is the estimated transformation after registration parameterized by $\hat{\Theta}$. This coefficient is sometimes referred as root-mean-squared error (RMS). If fiducial markers are used instead of anatomical points, then the metric is referred as fiducial registration error (FRE). Both concepts were first introduced in the seminal work of Fitzpatrick et al. (1998) to measure the performance of a geometric registration method, but have been adopted by the registration community as a general validation metric. Different variations of TRE have been used to validate slice-to-volume registration methods in Bao et al. (2005); Rousseau and Glenn (2005); Penney et al. (2006); Rousseau et al. (2006); Hummel et al. (2008); Leung et al. (2008); Wein et al. (2008a); Huang et al. (2009); Tadayyon et al. (2011); Yan et al. (2012); Cifor et al. (2013); Fuerst et al. (2014); Nir et al. (2014);

Rivaz et al. (2014c,b,a); Xu et al. (2014a,b); Yavariabdi et al. (2015); Xiao et al. (2016), including fiducial and anatomical landmarks.

Landmark-based error is a reliable indicator of the registration quality. However, it must be noted that in slice-to-volume registration cases where slices are too sparse, annotating such points of interest can be extremely difficult and sometimes impossible. A simple workaround was adopted by Pardasani et al. (2016) to adapt the MNI BITE dataset (where landmarks are identified in the reconstructed US volumes and the 3D MRI images) to the slice-to-volume registration case: they only used the US slices that were within 0.3mm of an expert-identified landmark. This reduced the number of available slices, but allowed them to compute the TRE indicator.

An alternative measurement based on point distances can be adopted even if landmarks are not available. In this case, the position of the voxels within a region of interest is compared with the corresponding voxels in the target before and after registration. Approaches considering distances between voxels have been used in several slice-to-volume registration methods like Fei et al. (2002, 2004, 2003c,a); Chandler et al. (2006); Kim et al. (2008b); Tadayyon et al. (2010); Kuklisova-Murgasova et al. (2012); Fogtman et al. (2014). However, we still require a way to establish correspondences between the voxels in both images, which is not always available.

7.2.3. Segmentation-based Metrics

When segmentation masks are available for certain structures of interest in both slice and volume images, segmentation-based statistics computed before and after registration can be used for validation. The rationale is that structures that were not aligned before registration, should be afterwards. Thus, the usual strategy consists in extracting a slice from the volume at the initial position, and doing the same after registration. At this point, the problem reduces to computing segmentation-based statistics between two 2D images. Alternative segmentation-based coefficients have been considered to deal with slice-to-volume registration validation.

Dice similarity coefficient (Dice, 1945) (DSC, also known as Sørensen-Dice coefficient) quantifies the amount of overlapping between two given segmentations masks A and B following $DSC(A, B) = \frac{2|A \cap B|}{|A| + |B|}$. Its value ranges from 0, meaning no spatial overlap, to 1, indicating complete overlap. This coefficient has been widely used in several slice-to-volume registration methods, where segmentation masks of liver tumor (Cifor et al., 2013), brain structures (Ferrante and Paragios, 2013; Ferrante et al., 2015a,b; Kainz et al., 2015), prostate (Nir et al., 2014), cardiac left ventricle (Xu et al., 2014b) and pelvic organs (Yavariabdi et al., 2015) were considered. Alternative overlapping measures were used in slice-to-volume studies by (Brooks et al., 2008; Chandler et al., 2008; Osechinskiy and Kruggel, 2011a; Museyko et al., 2014). As stated by (Yavariabdi et al., 2015), note that a high Dice does not mean a good contour overlap, which is a desired property after registration of structures of interest. In this case, indicators like Hausdorff distance (maximum distance between contours) and contour mean distance (CMD) can be reported (examples in the case of slice-to-volume registration can be found in Sun et al. (2007); Brooks et al. (2008); Elen et al. (2010); Tadayyon et al. (2010, 2011); Zakkaroff et al. (2012); Nir et al. (2014); Yavariabdi et al. (2015)).

The main advantage of using segmentation masks for slice-to-volume registration validation is their availability: annotations in the form of segmentation masks are more frequent than landmarks or ground-truth transformations. Moreover, if we aim to evaluate non-rigid registration, the estimated deformation fields can be used to warp the segmentation masks, enabling quantification for deformable registration error. However, as discussed in Rohlfing (2012), attention must be given to the structures of interest used for validation. Statistics computed using wide and non well-localized structures segmentation (like brain tissue for example), may not provide sufficient evidence to validate registration accuracy. According to the authors, only smaller, more localized regions of interest which approximate point landmarks, can provide such an evidence, since their overlap approximates point-based registration error. Combining area overlapping and contour distance based coefficients (which provide complementary information about the registration quality) has been suggested as a way to alleviate such inconsistencies Ferrante and Paragios (2015a); Yavariabdi et al. (2015).

7.2.4. Appearance-based metrics

Appearance information can also be used for validation. When the slice-to-volume registration method is an intermediate step towards performing motion correction/volume reconstruction, an indirect way to evaluate the registration quality is by quantifying the reconstruction accuracy. In this case, reconstruction error measured through intensity differences (usually the root mean square differences or RMSD) or signal-to-noise ratio (SNR) indicators between estimated and ground-truth volumes were used (Zarow et al., 2004; Noble et al., 2005; Chandler et al., 2006; Jiang et al., 2007a; Gholipour and Warfield, 2009; Gholipour et al., 2010; Honal et al., 2010; Gholipour et al., 2011; Yu et al., 2011; Kuklisova-Murgasova et al., 2012; Kainz et al., 2015). An alternative considered in several works (Gholipour and Warfield, 2009; Gholipour et al., 2010, 2011) is to measure the sharpness of the reconstruction. The intuition is that when uncorrected motion or error residuals between the input image acquisitions are present, the average image will be an out-of-focus motion blurred version of the imaged structure (Gholipour and Warfield, 2009). One of the main advantages of such measure is that it does not require ground-truth reconstructed images, since it is computed directly on the estimated volume.

Another case where simple intensity differences or correlation metrics between the estimated and target slice can be used is when performing monomodal slice-to-volume registration. In this case SAD, SSD or CC can be used to measure registration accuracy through visual error quantification since slice and volume intensities tend to be linearly correlated (see for example Marami et al. (2011); Porchetto et al. (2016)). In multimodal cases, more complex metrics (e.g. MI or NMI) or even ad-hoc similarity measures defined for specific modalities (like the LC^2 metric used by Wein et al. (2008b)) can also be adopted for validation.

The last strategy consists in quantified visual inspection, where one or multiple experts are asked to visually rate the quality of the registration assigning a score according to a given scale. This approach was taken in Birkfellner et al. (2007); Boer et al. (2007); Frühwald et al. (2009); Kim et al. (2010a); Po Su et al. (2013). The main disadvantage of such strategy is that it is highly time consuming and require human intervention. However, this approach can be

combined with any other listed in this survey, since it provides complementary information obtained from expert knowledge, which is difficult to quantify through objective/mathematical metrics.

8. Discussion and future directions

In this survey we reviewed and discussed some of the most important works in the literature related to slice-to-volume registration. These papers were classified according to different principles including matching criterion, transformation model, optimization method, number of slices, application and validation strategy.

One of the conclusions that emerges when analyzing Table 1 (which presents a compact summary of this work) is that the majority of the methods discussed in this survey focus on rigid registration, using iconic matching criteria and continuous optimization techniques. It is, therefore, worthwhile to wonder why this tendency is so clear and marked when dealing with slice-to-volume registration.

The fact that most of the works estimate rigid transformations may be related to two main reasons. The first one has to do with simple requirements from the application viewpoint. As stated in section 3.1, such a basic transformation model is expressive enough to explain simple slice-to-volume mapping, since it can deal with in-plane and out-of-plane translations and rotations. In clinical scenarios that do not inherit image distortion-like basic cases of inter-slice motion correction- this type of models may be enough. In more complex cases like image guided surgeries, it is a matter of fact that intra-operative images are deformed with respect to the pre-operative volume, due to tissue shift collapse and breathing/heart motion during the procedure. Moreover, not only in intra-operative cases but also during non-interventional imaging, many body organs motion (e.g. cardiac, lung or tongue) have natural elastic motion. Even so, rigid models still dominate in the literature. This is related to the second reason we identified: the widespread distrust of the non-rigid transformations among the physicians. From a medical perspective, it is sometimes preferable to supply reliable but approximate cues than unrealistic solutions. Non-rigid and elastic models might lead to solutions which are correct from a geometrical point of view but they are not anatomically meaningful. Further research on realistic deformation models reflecting physical tissue properties will certainly improve the accuracy of our estimations and develop trust in our methods inside the medical community.

Iconic matching criteria (based on alternative similarity measures) turned out to be the choice of preference to describe the quality of the solutions. This is somehow unexpected if we consider the lack of image support naturally associated to slice-to-volume registration, and the image noise frequently present in intra-operative, real-time and low-quality modalities, normally corresponding to the input 2D image, which makes it difficult to explain image similarities solely based on intensity information. In that sense, we identify two strategies helping to alleviate this limitation. The first one is to use (when possible) multiple slices instead of a single one. In that way, one can improve the matching quality by augmenting the image support, while keeping lower computational complexity when compared to standard 3D-3D registration methods (as explained in Section 5, this strategy proved to be particularly

useful when dealing with motion correction for image reconstruction). The second one is to combine iconic matching criteria with geometric or sensor-based strategies (as discussed in section 2.4). While the last ones are more robust to image changes and discontinuities, the former ones contribute to produce more accurate solutions. When used in tandem with sensor based methods (for example in image guided interventions where sensors can be attached to imaging devices or surgery tools), iconic similarity measures showed to be useful to refine initial rigid estimations and provide more precise results.

In terms of optimization strategy, even if the communities of computer vision and medical imaging started to massively adopt discrete approaches during the last decade, the same did not happen on the particular case of slice-to-volume registration. While most of the published methods adopt continuous optimizers (gradient and non-gradient based), only a few works recently published started to conceive this problem from a discrete perspective. Continuous methods are well established and showed to be good enough to deal with basic slice-to-volume registration problems. Nevertheless, results obtained in other biomedical image analysis problems modeled through discrete methods (Paragios et al., 2016) suggest that these strategies have a great potential which remains to be exploited for slice-to-volume registration. In cases where the number of parameters to be estimated is too high (e.g. deformable registration where dense deformation fields are adopted as transformation models), the energy functions are not convex (e.g. multimodal image registration) or the search space is remarkably wide (e.g. poor initializations), discrete methods may make a difference.

The research community of computer vision and medical imaging has made major efforts towards developing more accurate slice-to-volume registration strategies, mostly considering the geometric aspects of this problem. However, the massive production of image data that we are witnessing during the 21st century, combined with the latest advances in artificial intelligence and, in particular, deep convolutional networks (LeCun et al., 2015), open the possibility to conceive slice-to-volume image registration within an entirely different paradigm. Deep convolutional networks trained using massive amounts of data outperformed all existent state-of-the-art strategies in other fundamental vision tasks like image segmentation (Long et al., 2015) and object detection (Szegedy et al., 2013). Some recent works on image registration and vector flow estimation using CNNs (where not only the similarity measure (Simonovsky et al., 2016; Zagoruyko et al., 2015) but also the actual registration process (Fischer et al., 2015; Miao et al., 2016; Yang et al., 2016) is learned from examples) suggest that these problems are not an exception. The formulation of slice-to-volume image registration under this paradigm has just started to be explored and remains largely unexplored. In that sense, the promising results already achieved (Hou et al., 2017) suggest that shifting to a learning-based slice-to-volume registration paradigm may lead to faster and more accurate predictions, opening new doors in the field.

Acknowledgments

This research was partially supported from the European Research Council Grant 259112.

Table 3: Alternative validation metrics considered in the context of slice-to-volume registration.

Metric	Paper	
Distance between transformation parameters	Kim et al. (1999); Park et al. (2004); Yeo et al. (2004); Smolřková Wachowiak et al. (2005); Dalvi and Abugharbieh (2008); Gefen et al. (2008); Kim et al. (2008a); Yeo et al. (2008); Frřhwald et al. (2009); Jiang et al. (2009); San Jose Estepar et al. (2009); Elen et al. (2010); Gholipour et al. (2010); Zakkaroff et al. (2012); Ferrante and Paragios (2013); Lin et al. (2013); Seshamani et al. (2013); Po Su et al. (2013); Schulz et al. (2014); Ferrante et al. (2015a,b); Porchetto et al. (2016); Chen et al. (2016)	
Point-based metrics	Fei et al. (2002, 2003a, 2004, 2003c); Bao et al. (2005); Rousseau and Glenn (2005); Chandler et al. (2006); Penney et al. (2006); Rousseau et al. (2006); Kuklisova-Murgasova et al. (2012); Kim et al. (2008b); Hummel et al. (2008); Leung et al. (2008); Wein et al. (2008a); Huang et al. (2009); Tadayyon et al. (2010, 2011); Yan et al. (2012); Cifor et al. (2013); Chicherova et al. (2014); Fogtmann et al. (2014); Fuerst et al. (2014); Nir et al. (2014); Rivaz et al. (2014c,b,a); Xu et al. (2014a,b); Hallack et al. (2015); Yavariabdi et al. (2015); Xiao et al. (2016); Chen et al. (2016)	
Segmentation	Overlapping error	Brooks et al. (2008); Chandler et al. (2008); Osechinskiy and Kruggel (2011a); Cifor et al. (2013); Ferrante and Paragios (2013); Museyko et al. (2014); Nir et al. (2014); Xu et al. (2014b); Ferrante et al. (2015b); Kainz et al. (2015); Yavariabdi et al. (2015)
	Contour distances	Sun et al. (2007); Brooks et al. (2008); Elen et al. (2010); Tadayyon et al. (2010, 2011); Zakkaroff et al. (2012); Nir et al. (2014); Ferrante et al. (2015b); Yavariabdi et al. (2015)
Appearance	Intensity / reconstruction	Zarow et al. (2004); Noble et al. (2005); Chandler et al. (2006); Jiang et al. (2007a); Wein et al. (2008b); Gholipour and Warfield (2009); Osechinskiy and Kruggel (2009); Gholipour et al. (2010); Honal et al. (2010); Gholipour et al. (2011); Marami et al. (2011); Yu et al. (2011); Kuklisova-Murgasova et al. (2012); Po Su et al. (2013); Kainz et al. (2015); Porchetto et al. (2016); Marami et al. (2016b,a)
	Visual inspection	Birkfellner et al. (2007); Boer et al. (2007); Frřhwald et al. (2009); Kim et al. (2010a); Po Su et al. (2013)

References

- Andres, B., Kappes, J.H., Beier, T., Köthe, U., Hamprecht, F.A., 2012. The Lazy Flipper: Efficient Depth-Limited Exhaustive Search in Discrete Graphical Models, pp. 154–166.
- Arbel, T., Morandi, X., Comeau, R.M., Collins, D.L., 2001. Automatic Non-linear MRI-Ultrasound Registration for the Correction of Intra-operative Brain Deformations. *MIC-CAI* , 913–922.
- Ballard, D., 1981. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition* 13, 111–122. URL: <http://linkinghub.elsevier.com/retrieve/pii/0031320381900091>, doi:10.1016/0031-3203(81)90009-1.
- Bao, P., Warmath, J., Galloway, R., Herline, A., 2005. Ultrasound-to-computer-tomography registration for image-guided laparoscopic liver surgery. *Surgical Endoscopy* 19, 424–429. doi:10.1007/s00464-004-8902-1.
- Bay, H., Ess, A., Tuytelaars, T., Van Gool, L., 2008. Speeded-Up Robust Features (SURF). *Computer Vision and Image Understanding* 110, 346–359. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1077314207001555>, doi:10.1016/j.cviu.2007.09.014.
- Belongie, S., Malik, J., Puzicha, J., 2002. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 509–522. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=993558>, doi:10.1109/34.993558.
- Besag, J., 1986. Statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)* , 259—302 URL: <http://www.tandfonline.com/doi/abs/10.1080/026647693000000059>, doi:10.1080/026647693000000059.
- Besl, P., McKay, H., 1992. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 239–256. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=121791>, doi:10.1109/34.121791.
- Bhagalia, R., Fessler, J.a., Kim, B., Meyer, C.R., 2009. Improved fMRI time-series registration using joint probability density priors, in: *SPIE*, pp. 72590J–72590J–9. URL: <http://link.aip.org/link/PSISDG/v7259/i1/p72590J/s1&Agg=doi>, doi:10.1117/12.811421.
- Bhattachayya, A., 1943. On a measure of divergence between two statistical population defined by their population distributions. *Bulletin Calcutta Mathematical Society* .

- Birkfellner, W., Figl, M., Kettenbach, J., Hummel, J., Homolka, P., Schernthaner, R., Nau, T., Bergmann, H., 2007. Rigid 2D/3D slice-to-volume registration and its application on fluoroscopic CT images. *Medical Physics* 34, 246.
- Birkfellner, W., Hummel, J., Wilson, E., 2008. Image-Guided Interventions Technology and Applications (Chapter 2, Tracking Devices), in: Peters, T., Cleary, K. (Eds.), *Image-Guided Interventions Technology and Applications*. Springer. chapter Chapter 2, pp. 23–45.
- Boer, B., Ruijter, J., Voorbaak, F.P., 2007. Towards the automatic registration of histological sections into a 3D reference mode, in: *Belgian-Dutch Conference on Artificial Intelligence*. doi:10.1177/1745691612459060.
- Boykov, Y., Kolmogorov, V., 2004. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 1124–1137. doi:10.1109/TPAMI.2004.60.
- Boykov, Y., Veksler, O., Zabih, R., 2001. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 1222–1239. doi:10.1109/34.969114.
- Brooks, R., Collins, D.L., Morandi, X., Arbel, T., 2008. Deformable Ultrasound Registration without Reconstruction, in: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. URL: <http://link.springer.com/10.1007/978-3-540-85990-1>, doi:10.1007/978-3-540-85990-1.
- Chan, T., Vese, L., 2001. Active contours without edges. *IEEE Transactions on Image Processing* 10, 266–277. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=902291>, doi:10.1109/83.902291.
- Chandler, A., Pinder, R., Netsch, T., Schnabel, J., Hawkes, D., Hill, D., Razavi, R., 2006. Correction of Misaligned Slices in Multi-Slice MR Cardiac Examinations by Using Slice-To-Volume Registration, in: *IEEE International Symposium on Biomedical Imaging (ISBI)*, IEEE. pp. 474–477. URL: <http://www.ncbi.nlm.nih.gov/pubmed/18312619> <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1624956>, doi:10.1109/ISBI.2006.1624956.
- Chandler, A.G., Netsch, T., Cocosco, C.A., Schnabel, J.A., Hawkes, D.J., 2004. Slice-to-volume registration using mutual information between probabilistic image classifications, in: Fitzpatrick, J.M., Sonka, M. (Eds.), *SPIE*, pp. 1120–1129. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=841250>, doi:10.1117/12.534966.
- Chandler, A.G., Pinder, R.J., Netsch, T., Schnabel, J.A., Hawkes, D.J., Hill, D.L.G., Razavi, R., 2008. Correction of misaligned slices in multi-slice cardiovascular magnetic resonance

using slice-to-volume registration. *Journal of Cardiovascular Magnetic Resonance* 10, 13. URL: <http://www.jcmr-online.com/content/10/1/13>.

- Chen, Y.h., Mittelman, R., Kim, B., Meyer, C., Hero, A., 2016. Particle filtering for slice-to-volume motion correction in EPI based functional MRI, in: 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE. pp. 679–683. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7471761>, doi:10.1109/ICASSP.2016.7471761.
- Chicherova, N., Fundana, K., Bert, M., Cattin, P.C., 2014. Histology to CT Data Matching Using Landmarks and a Density Biased RANSAC, in: MICCAI 2014, pp. 243–250.
- Cifor, A., Risser, L., Heinrich, M.P., Chung, D., Schnabel, J.A., 2013. Rigid Registration of Untracked Freehand 2D Ultrasound Sweeps to 3D CT of Liver Tumours. *Abdominal Imaging. Computation and Clinical Applications, MICCAI*.
- Cocosco, C.A., Kollokian, V., Kwan, R.K.S., Pike, G.B., Evans, A.C., 1997. Brainweb: Online interface to a 3d mri simulated brain database. *NeuroImage* 5, 425.
- Dalal, N., Triggs, B., 2005. Histograms of Oriented Gradients for Human Detection, in: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), IEEE. pp. 886–893. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1467360>, doi:10.1109/CVPR.2005.177.
- Dalvi, R., Abugharbieh, R., 2008. Fast feature based multi slice to volume registration using phase congruency. *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS) 2008*, 5390–3. URL: <http://www.ncbi.nlm.nih.gov/pubmed/19163936>, doi:10.1109/IEMBS.2008.4650433.
- Dice, L.R., 1945. Measures of the Amount of Ecologic Association Between Species. *Ecology* 26.
- Elen, A., Hermans, J., Ganame, J., Loeckx, D., Bogaert, J., Maes, F., Suetens, P., 2010. Automatic 3-D Breath-Hold Related Motion Correction of Dynamic Multislice MRI. *IEEE Transactions on Medical Imaging* 29, 868–878. doi:10.1109/TMI.2009.2039145.
- Eresen, A., Peng Li, Jim Xiuquan Ji, 2014. Correlating 2D histological slice with 3D MRI image volume using smart phone as an interactive tool for muscle study, in: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, IEEE. pp. 6418–6421. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6945097>, doi:10.1109/EMBC.2014.6945097.
- Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A., 2010. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision* 88, 303–338. doi:10.1007/s11263-009-0275-4.

- Ewertsen, C., Saftoiu, A., Gruionu, L.G., Karstrup, S., Nielsen, M.B., 2013. Real-time image fusion involving diagnostic ultrasound. *AJR. American journal of roentgenology* 200, W249–55. URL: <http://www.ncbi.nlm.nih.gov/pubmed/23436869>, doi:10.2214/AJR.12.8904.
- Fei, B., Duerk, J.L., Boll, D.T., Lewin, J.S., Wilson, D.L., 2003a. Slice-to-volume registration and its potential application to interventional MRI-guided radio-frequency thermal ablation of prostate cancer. *IEEE Transactions on Medical Imaging* 22, 515–525. doi:10.1109/TMI.2003.809078.
- Fei, B., Duerk, J.L., Wilson, D.L., 2002. Automatic 3D registration for interventional MRI-guided treatment of prostate cancer. *Computer Aided Surgery* 7, 257–267. doi:10.1002/igs.10052.
- Fei, B., Frinkley, K., Wilson, D.L., 2003b. Registration Algorithms for Interventional MRI-guided Treatment of the Prostate, in: *SPIE*, pp. 192–201. doi:10.1117/12.479696.
- Fei, B., Lee, Z., Boll, D.T., Duerk, J.L., Sodee, D.B., Lewin, J.S., Wilson, D.L., 2003c. Image Registration and Fusion for Interventional MRI Guided Thermal Ablation of the Prostate Cancer. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* 51, 177–183. doi:10.1109/TNS.2003.823027.
- Fei, B., Lee, Z., Boll, D.T., Duerk, J.L., Sodee, D.B., Lewin, J.S., Wilson, D.L., 2004. Registration and fusion of SPECT, high-resolution MRI, and interventional MRI for thermal ablation of prostate cancer. *IEEE Transactions on Nuclear Science* 51, 177–183. doi:10.1109/TNS.2003.823027.
- Fei, B., Wheaton, A., Lee, Z., Nagano, K., Duerk, J.L., Wilson, D.L., 2001. Robust Registration Method for Interventional MRI-Guided Thermal Ablation of Prostate Cancer, in: Mun, S.K. (Ed.), *SPIE*, pp. 53–60. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=905627>, doi:10.1117/12.428092.
- Ferrante, E., Fecamp, V., Paragios, N., 2015a. Implicit planar and in-plane deformable mapping in medical images through high order graphs, in: *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, IEEE. pp. 721–724. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7163974>, doi:10.1109/ISBI.2015.7163974.
- Ferrante, E., Fecamp, V., Paragios, N., 2015b. Slice-to-volume deformable registration: efficient one-shot consensus between plane selection and in-plane deformation, in: *International Journal of Computer Assisted Radiology and Surgery*, Springer Berlin Heidelberg. pp. 791–800. doi:10.1007/s11548-015-1205-2.

- Ferrante, E., Paragios, N., 2013. Non-rigid 2D-3D Medical Image Registration Using Markov Random Fields, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI). volume 8151, pp. 163–170.
- Ferrante, E., Paragios, N., 2015a. Graph based slice-to-volume deformable registration. Technical report. INRIA. .
- Ferrante, E., Paragios, N., 2015b. Graph-based slice-to-volume registration. [Research Report] RR-8803, INRIA Saclay; Center for Visual Computing, CentraleSupélec, Inria, Université Paris-Saclay, pp.30 .
- Fischer, P., Dosovitskiy, A., Ilg, E., ..., Brox, T., 2015. FlowNet: Learning Optical Flow with Convolutional Networks. ICCV 2015 .
- Fitzpatrick, J.M., West, J.B., Maurer, C.R., 1998. Predicting error in rigid-body point-based registration. IEEE transactions on medical imaging 17, 694–702.
- Fogtmann, M., Seshamani, S., Kroenke, C., Cheng, X., Chapman, T., Wilm, J., Rousseau, F., Studholme, C., 2014. A unified approach to diffusion direction sensitive slice registration and 3-D DTI reconstruction from moving fetal brain anatomy. IEEE Transactions on Medical Imaging 33, 272–289. doi:10.1109/TMI.2013.2284014.
- Frühwald, L., Kettenbach, J., Figl, M., Hummel, J., Bergmann, H., Birkfellner, W., 2009. A comparative study on manual and automatic slice-to-volume registration of CT images. European radiology 19, 2647–53.
- Fuerst, B., Wein, W., Müller, M., Navab, N., 2014. Automatic ultrasound-MRI registration for neurosurgery using the 2D and 3D LC(2) Metric. Medical image analysis 18, 1312–9. URL: <http://www.ncbi.nlm.nih.gov/pubmed/24842859>, doi:10.1016/j.media.2014.04.008.
- Gefen, S., Bertrand, L., Kiryati, N., Nissanov, J., 2005. Localization of Sections Within the Brain Via 2D to 3D Image Registration, in: IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), IEEE. pp. 733–736. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1415509>, doi:10.1109/ICASSP.2005.1415509.
- Gefen, S., Kiryati, N., Nissanov, J., 2008. Atlas-based indexing of brain sections via 2-D to 3-D image registration. IEEE Transactions on Biomedical Engineering 55, 147–156. doi:10.1109/TBME.2007.899361.
- Ghanavati, S., Mousavi, P., Fichtinger, G., Abolmaesumi, P., 2011. Phantom validation for ultrasound to statistical shape model registration of human pelvis, in: Wong, K.H., Holmes III, D.R. (Eds.), SPIE, p. 79642U. doi:10.1117/12.876998.

- Ghanavati, S., Mousavi, P., Fichtinger, G., Foroughi, P., Abolmaesumi, P., 2010. Multi-Slice to Volume Registration of Ultrasound Data to a Statistical Atlas of Human Pelvis, in: SPIE, pp. 76250O–76250O–10. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=747944>, doi:10.1117/12.844080.
- Gholipour, A., Estroff, J.A., Warfield, S.K., 2010. Robust Super-Resolution Volume Reconstruction From Slice Acquisitions: Application to Fetal Brain MRI. *IEEE Transactions on Medical Imaging* 29, 1739–1758. doi:10.1109/TMI.2010.2051680.
- Gholipour, A., Polak, M., van der Kouwe, A., Nevo, E., Warfield, S.K., 2011. Motion-robust MRI through real-time motion tracking and retrospective super-resolution volume reconstruction, in: International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS), IEEE. pp. 5722–5725. doi:10.1109/IEMBS.2011.6091385.
- Gholipour, A., Warfield, S.K., 2009. Super-resolution reconstruction of fetal brain MRI. *MICCAI Workshop on Image Analysis for the Developing Brain (IADB'2009)*. , 45–52.
- Gill, S., Abolmaesumi, P., Vikal, S., Mousavi, P., Fichtinger, G., 2008. Intraoperative Prostate Tracking with Slice-to-Volume Registration in MR, in: Proceedings of the 20th International Conference of the Society for Medical Innovation and Technology.
- Glocker, B., 2010. Random Fields for Image Registration. PhD Thesis .
- Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N., 2008. Dense image registration through MRFs and efficient linear programming. *Medical image analysis* 12, 731–41.
- Glocker, B., Sotiras, A., 2011. Deformable Medical Image Registration: Setting the State of the Art with Discrete Methods. *Annual review of biomedical engineering* 13, 219–244.
- Gourdon, A., Ayache, N., 1994. Registration of a curve on a surface using differential properties, in: European Conference on Computer Vision.
- Guzman, L., Commandeur, F., Acosta, O., Simon, A., Fautrel, A., Rioux-Leclercq, N., Romero, E., Mathieu, R., de Crevoisier, R., 2016. Slice correspondence estimation using surf descriptors and context-based search for prostate whole-mount histology mri registration, in: *IEEE Engineering in Medicine and Biology*, pp. 1163–1166.
- Hallack, A., Papiez, B.W., Wilson, J., Wang, L.M., Maughan, T., Gooding, M., Schnabel, J., 2015. Correlating Tumour Histology and ex vivo MRI Using Dense Modality-Independent Patch-Based Descriptors, in: *Patch-MI 2015 MICCAI Workshop*, pp. 137–145. doi:10.1007/978-3-319-28194-0.
- Harris, C., Stephens, M., 1988. A combined corner and edge detector, in: *Proc. of Fourth Alvey Vision Conference*.

- Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, S.M., Schnabel, J.a., 2012. MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration. *Medical Image Analysis* 16, 1423–1435. URL: <http://dx.doi.org/10.1016/j.media.2012.05.008>, doi:10.1016/j.media.2012.05.008.
- Heinrich, M.P., Jenkinson, M., Papież, B.W., Brady, S.M., Schnabel, J.A., 2013. Towards Realtime Multimodal Fusion for Image-Guided Interventions Using Self-similarities, in: *MICCAI 2013*, pp. 187–194. URL: http://link.springer.com/10.1007/978-3-642-40811-3_24.
- Heldmann, S., Beuthien, B., Olesch, J., 2010. Improved Minimal-Invasive Laparoscopic Liver Surgery by Registration of 3D CT and 2D Ultrasound Slices. *Proc BMT* , 2–4doi:10.1515/BMT.2010.368.
- Heldmann, S., Papenberg, N., 2009. A Variational Approach for Volume-to-Slice Registration, in: *Scale Space and Variational Methods in Computer Vision*, pp. 624–635.
- Hoerth, R., Baum, D., Knötel, D., Prohaska, S., Willie, B., Duda, G., Hege, H.C., Fratzl, P., Wagermaier, W., 2015. Registering 2D and 3D imaging data of bone during healing. *Connective Tissue Research* 56, 133–143. doi:10.3109/03008207.2015.1005210.
- Holden, M., 2008. A Review of Geometric Transformations for Nonrigid Body Registration. *IEEE Transactions on Medical Imaging* 27, 111–128. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4359047>, doi:10.1109/TMI.2007.904691.
- Honal, M., Leupold, J., Huff, S., Baumann, T., Ludwig, U., 2010. Compensation of breathing motion artifacts for MRI with continuously moving table. *Magnetic resonance in medicine* 63, 701–12. URL: <http://www.ncbi.nlm.nih.gov/pubmed/20187180>, doi:10.1002/mrm.22162.
- Hou, B., Alansary, A., Mcdonagh, S., Davidson, A., Hajnal, J.V., Rueckert, D., Glocker, B., Kainz, B., 2017. Predicting Slice-to-Volume Transformation in Presence of Arbitrary Subject Motion. *ArXiv Preprint 1702.08891v2* .
- Huang, X., Moore, J., Guiraudon, G., Jones, D.L., Bainbridge, D., Ren, J., Peters, T.M., 2009. Dynamic 2D ultrasound and 3D CT image registration of the beating heart. *IEEE Transactions on Medical Imaging* 28, 1179–1189.
- Hummel, J., Figl, M., Bax, M., Bergmann, H., Birkfellner, W., 2008. 2D/3D registration of endoscopic ultrasound to CT volume data. *Physics in Medicine and Biology* 53, 4303–4316. doi:10.1088/0031-9155/53/16/006.
- IBSR, . IBSR. Internet Brain Segmentation Repository [Online]. Available: <http://www.cma.mgh.harvard.edu/ibsr/> URL: Available: <http://www.cma.mgh.harvard.edu/ibsr/>.

- Jiang, S., Xue, H., Counsell, S., Anjari, M., Allsop, J., Rutherford, M., Rueckert, D., Hajnal, J.V., 2009. Diffusion tensor imaging (DTI) of the brain in moving subjects: Application to in-utero fetal and ex-utero studies. *Magnetic Resonance in Medicine* 62, 645–655. URL: <http://doi.wiley.com/10.1002/mrm.22032>, doi:10.1002/mrm.22032.
- Jiang, S., Xue, H., Counsell, S.J., Anjari, M., Allsop, J., Rutherford, M.A., Rueckert, D., Hajnal, J.V., 2007a. In-utero 3D high resolution fetal brain diffusion tensor imaging, in: Ayache, N., Ourselin, S., Maeder, A. (Eds.), *Medical Image Computing and Computer-Assisted Interventions (MICCAI)*, Berlin, Heidelberg. URL: <http://link.springer.com/10.1007/978-3-540-75757-3>, doi:10.1007/978-3-540-75757-3.
- Jiang, S., Xue, H., Glover, A., Rutherford, M., Rueckert, D., Hajnal, J.V., 2007b. MRI of moving subjects using multislice Snapshot images with Volume Reconstruction (SVR): Application to fetal, neonatal, and adult brain studies. *IEEE Transactions on Medical Imaging* 26, 967–980. doi:10.1109/TMI.2007.895456.
- Joshi, S.C., Miller, M.I., 2000. Landmark matching via large deformation diffeomorphisms. *IEEE transactions on image processing* 9, 1357–70. URL: <http://www.ncbi.nlm.nih.gov/pubmed/18262973>, doi:10.1109/83.855431.
- Kainz, B., Steinberger, M., Wein, W., Kuklisova-Murgasova, M., Malamateniou, C., Keraudren, K., Torsney-Weir, T., Rutherford, M., Aljabar, P., Hajnal, J.V., Rueckert, D., 2015. Fast Volume Reconstruction From Motion Corrupted Stacks of 2D Slices. *IEEE Transactions on Medical Imaging* 34, 1901–1913. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7064742>, doi:10.1109/TMI.2015.2415453.
- Kalman, R., 1960. A new approach to linear filtering and prediction problems. *Trans ASME J Basic Eng* 82, 35–45.
- Kappes, J.H., Andres, B., Hamprecht, F.A., Schnorr, C., Nowozin, S., Batra, D., Kim, S., Kausler, B.X., Lellmann, J., Komodakis, N., Rother, C., 2013. A Comparative Study of Modern Inference Techniques for Discrete Energy Minimization Problems, in: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE. pp. 1328–1335. doi:10.1109/CVPR.2013.175.
- Kim, B., Boes, J.L., Bland, P.H., Chenevert, T.L., Meyer, C.R., 1999. Motion correction in fMRI via registration of individual slices into an anatomical volume. *Magnetic Resonance in Medicine* 41, 964–972.
- Kim, B., Meyer, C.R., Bland, P.H., 2000. Correction of local deformations in fMRI by 3D non-linear warping in map-slice-to-volume approach, in: *Proc. of Intl. Soc. Mag. Reson. Med.*, p. 59412.

- Kim, B., Yeo, D.T.B., Bhagalia, R., 2008a. Comprehensive mathematical simulation of functional magnetic resonance imaging time series including motion-related image distortion and spin saturation effect. *Magnetic resonance imaging* 26, 147–59. URL: <http://www.ncbi.nlm.nih.gov/pubmed/17662548>, doi:10.1016/j.mri.2007.05.007.
- Kim, K., Habas, P., Rousseau, F., Glenn, O., Barkovich, A., Studholme, C., 2010a. Intersection Based Motion Correction of Multislice MRI for 3D Fetal Brain Image Formation. *IEEE Transactions on Medical Imaging* 29, 146–158. doi:10.1109/TMI.2009.2030679.
- Kim, K., Habas, P.A., Rajagopalan, V., Scott, J.A., Corbett-Detig, J.M., Rousseau, F., Barkovich, A.J., Glenn, O.A., Studholme, C., 2011. Bias field inconsistency correction of motion-scattered multislice MRI for improved 3D image reconstruction. *IEEE transactions on medical imaging* 30, 1704–12. doi:10.1109/TMI.2011.2143724.
- Kim, K., Habas, P.a., Rousseau, F., Glenn, O.a., Barkovich, a.J., Koob, M., Dietemann, J.L., Robinson, A.J., Poskitt, K.J., Miller, S.P., Studholme, C., 2010b. Reconstruction of a geometrically correct diffusion tensor image of a moving human fetal brain, in: *SPIE*, pp. 76231I–76231I–9. URL: <http://link.aip.org/link/PSISDG/v7623/i1/p76231I/s1&Agg=doi>, doi:10.1117/12.844542.
- Kim, K., Hansen, M., Habas, P., Rousseau, F., Glenn, O., Barkovich, a.J., Studholme, C., 2008b. Intersection based registration of slice stacks to form 3D images of the human fetal brain. *IEEE International Symposium on Biomedical Imaging (ISBI)* , 1167–1170doi:10.1109/ISBI.2008.4541209.
- Kim, T.S., 2005. Co-registration of Multiple Postmortem Brain Slices to Corresponding MRIs Using Voxel Similarity Measures and Slice-to-Volume Transformation. *Journal of Biomedical Engineering Research* , 231–241.
- Kio Kim, Habas, P., Rajagopalan, V., Scott, J., Corbett-Detig, J., Rousseau, F., Glenn, O., Barkovich, J., Studholme, C., 2010. Non-iterative relative bias correction for 3D reconstruction of in utero fetal brain MR imaging, in: *IEEE Engineering in Medicine and Biology Society (EMBS)*, IEEE. pp. 879–882. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5627876>, doi:10.1109/IEMBS.2010.5627876.
- Komodakis, N., Paragios, N., Tziritas, G., 2011. MRF Energy Minimization and Beyond via Dual Decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 531–552. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5467090>, doi:10.1109/TPAMI.2010.108.
- Komodakis, N., Tziritas, G., 2007. Approximate Labeling via Graph Cuts Based on Linear Programming. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29, 1436–1453. URL:

<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4250468>,
doi:10.1109/TPAMI.2007.1061.

- Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M.A., Hajnal, J.V., Schnabel, J.A., 2012. Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Medical Image Analysis* 16, 1550–1564. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1361841512000965>, doi:10.1016/j.media.2012.07.004.
- Lasowski, R., Benhimane, S., Vogel, J., Jakobs, T.F., Zech, C.J., Trumm, C., Clason, C., Navab, N., Solutions, M., Tomography, C., 2008. Adaptive visualization for needle guidance in RF liver ablation: taking organ deformation into account, in: *SPIE*, pp. 69180A–69180A–12. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=829479>, doi:10.1117/12.769847.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521, 436–444. URL: <http://www.nature.com/doi/finder/10.1038/nature14539>, doi:10.1038/nature14539.
- Leung, K.Y.E., van Stralen, M., Nemes, A., Voormolen, M.M., van Burken, G., Geleijnse, M.L., Ten, F.J., Cate, J., Reiber, H.C., 2008. Sparse Registration for Three-Dimensional Stress Echocardiography. *IEEE transactions on medical imaging* .
- Liao, R., Zhang, L., Sun, Y., Miao, S., Chefd’Hotel, C., 2013. A Review of Recent Advances in Registration Techniques Applied to Minimally Invasive Therapy. *IEEE Transactions on Multimedia* 15, 983–1000. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6428702>, doi:10.1109/TMM.2013.2244869.
- Lin, C.C., Zhang, S., Frahm, J., Lu, T.W., Hsu, C.Y., Shih, T.F., 2013. A slice-to-volume registration method based on real-time magnetic resonance imaging for measuring three-dimensional kinematics of the knee. *Medical Physics* .
- Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation. *CVPR* .
- Long, J.L., Zhang, N., Darrell, T., 2014. Do Convnets Learn Correspondence?, in: *Advances in Neural Information Processing Systems*. NIPS., pp. 1601–1609.
- Lorensen, W.E., Cline, H.E., 1987. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics* 21, 163–169. doi:10.1145/37402.37422.
- Lowe, D., 1999. Object recognition from local scale-invariant features, in: *Proceedings of the Seventh IEEE International Conference on Computer Vision*, IEEE. pp. 1150–1157 vol.2. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=790410>, doi:10.1109/ICCV.1999.790410.

- Maes, F., Vandermeulen, D., Marchal, G., Suetens, P., 1997. Multimodality Image Registration by Maximization of Mutual Information 16, 187–198. doi:10.1109/42.563664.
- Marami, B., Scherrer, B., Afacan, O., Erem, B., Warfield, S.K., Gholipour, A., 2016a. Motion-Robust Diffusion-Weighted Brain MRI Reconstruction Through Slice-Level Registration-Based Motion Tracking. *IEEE Transactions in Medical Imaging* 35, 2258–2269.
- Marami, B., Scherrer, B., Afacan, O., Warfield, S.K., 2016b. Motion-Robust Reconstruction Based on Simultaneous Multi-slice Registration for Diffusion-Weighted MRI of Moving Subjects, in: *MICCAI 2016*, pp. 544–552. doi:10.1007/978-3-319-46726-9.
- Marami, B., Sirouspour, S., Capson, D.W., 2011. Model-based 3D/2D deformable registration of MR images. *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS) 2011*, 4880–3. doi:10.1109/IEMBS.2011.6091209.
- Markelj, P., Tomaževič, D., Likar, B., Pernuš, F., 2010. A review of 3D/2D registration methods for image-guided interventions. *Medical Image Analysis* 16, 642–661.
- Mercier, L., Del Maestro, R.F., Petrecca, K., Araujo, D., Haegelen, C., Collins, D.L., 2012. Online database of clinical MR and ultrasound images of brain tumors. *Medical Physics* 39, 3253. URL: <http://scitation.aip.org/content/aapm/journal/medphys/39/6/10.1118/1.4709600>, doi:10.1118/1.4709600.
- Miao, S., Liao, R., Moran, G., Butler, J., Pan, L., Wang, Z.J., 2014. Dynamic MR-based respiratory motion compensation for hybrid PET/MR system. *IEEE Conference on Industrial Electronics and Applications* , 1915–1920 URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6931481>.
- Miao, S., Wang, Z.J., Liao, R., 2016. A CNN Regression Approach for Real time 2D 3D Registration. *IEEE Transactions on Medical Imaging* 35.
- Micu, R., Jakobs, T., Urschler, M., Navab, N., 2006. A new registration/visualization paradigm for ct-fluoroscopy guided rf liver ablation. *International Conference on Medical Image Computing and Computer Aided Intervention (MICCAI)* , 882–890 URL: http://link.springer.com/chapter/10.1007/11866565_108.
- Mitra, J., Ghose, S., Sidibe, D., Oliver, A., Marti, R., Llado, X., Vilanova, J.C., Comet, J., Meriaudeau, F., 2012. Weighted likelihood function of multiple statistical parameters to retrieve 2D TRUS-MR slice correspondence for prostate biopsy, in: *IEEE International Conference on Image Processing (ICIP)*, IEEE. pp. 2949–2952. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6467518>, doi:10.1109/ICIP.2012.6467518.

- Moradi, M., Mahdavi, S., 2013. Ultrasound RF time series for tissue typing: first in vivo clinical results. SPIE 8670, 1–8. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1669561>, doi:10.1117/12.2007672.
- Murgasova, M., Estrin, G.L., Rutherford, M., Rueckert, D., Hajnal, J., 2016. Distortion correction in fetal EPI using non-rigid registration with Laplacian constraint, in: 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI), IEEE. pp. 1372–1375. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7493522>, doi:10.1109/ISBI.2016.7493522.
- Murphy, K.P., Weiss, Y., Jordan, M.I., 1999. Loopy Belief Propagation for Approximate Inference: An Empirical Study, in: Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence (UAI'99),.
- Museyko, O., Marshall, R.P., Lu, J., Hess, A., Schett, G., Amling, M., Kalender, W.A., Engelke, K., Oleg Museykoa*, Robert Percy Marshallb, Jing Lua, Andreas Hessc, Georg Schettd, Michael Amlingb, W.A.K.&K.E., 2014. Registration of 2D histological sections with 3D micro-CT datasets from small animal vertebrae and tibiae. *Computer Methods in Biomechanics and Biomedical Engineering* 5842, 1–16. doi:10.1080/10255842.2014.941824.
- Nelder, J.A., Mead, R., 1965. A Simplex Method for Function Minimization. *The Computer Journal* 7, 308–313. URL: <http://comjnl.oxfordjournals.org/cgi/doi/10.1093/comjnl/7.4.308>, doi:10.1093/comjnl/7.4.308.
- Nir, G., Sahebjavaher, R.S., Kozlowski, P., Chang, S.D., Jones, E.C., Goldenberg, S.L., Salcudean, S.E., 2014. Registration of whole-mount histology and volumetric imaging of the prostate using particle filtering. *IEEE Transactions on Medical Imaging* 33, 1601–1613.
- Nir, G., Tannenbaum, A., 2011. Temporal registration of partial data using particle filtering, in: *IEEE International Conference on Image Processing*, IEEE. pp. 2177–2180. doi:10.1109/ICIP.2011.6116064.
- Noble, N.M.I., Boubertakh, R., Razavi, R.S., Hill, D.L.G., 2005. Inter-breath-hold Registration for the Production of High Resolution Cardiac MR Volumes, in: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 894–901. doi:10.1007/11566489_110.
- Oktaç, O., Bai, W., Lee, M., Guerrero, R., Kamnitsas, K., Caballero, J., de Marvao, A., Cook, S., O'Regan, D., Rueckert, D., 2016. Multi-input cardiac image super-resolution using convolutional neural networks, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer International Publishing. pp. 246–254.

- Olesch, J., Beuthien, B., Heldmann, S., Papenberg, N., Fischer, B., 2011. Fast intra-operative non-linear registration of 3D-CT to tracked, selected 2D-ultrasound slices, in: SPIE, pp. 3–8. doi:10.1117/12.878105.
- Olesch, J., Fischer, B., 2011. Focused Registration of Tracked 2D US to 3D CT Data of the Liver, in: Bildverarbeitung für die Medizin 2011, pp. 79–83. doi:10.1007/978-3-642-19335-4_18.
- Osechinskiy, S., Kruggel, F., 2009. Quantitative comparison of high-resolution MRI and myelin-stained histology of the human cerebral cortex. International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS) , 85–89doi:10.1109/IEMBS.2009.5334695.
- Osechinskiy, S., Kruggel, F., 2011a. Deformable registration of histological sections to brain MR images using a hybrid boundary-based slice-to-volume approach. International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS) 2011, 4876–9. URL: <http://www.ncbi.nlm.nih.gov/pubmed/22255431>, doi:10.1109/IEMBS.2011.6091208.
- Osechinskiy, S., Kruggel, F., 2011b. Slice-to-Volume Nonrigid Registration of Histological Sections to MR Images of the Human Brain. Anatomy Research International 2011, 1–17. URL: <http://www.hindawi.com/journals/ari/2011/287860/>, doi:10.1155/2011/287860.
- Paragios, N., Ferrante, E., Glocker, B., Komodakis, N., Parisot, S., Zacharaki, E.I., 2016. (Hyper)-graphical models in biomedical image analysis. Medical Image Analysis URL: <http://linkinghub.elsevier.com/retrieve/pii/S1361841516301062>, doi:10.1016/j.media.2016.06.028.
- Paragios, N., Rousson, M., Ramesh, V., 2003. Non-rigid registration using distance functions. Computer Vision and Image Understanding 89, 142–165. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1077314203000109>, doi:10.1016/S1077-3142(03)00010-9.
- Pardasani, U., Baxter, J.S.H., Peters, T.M., Khan, A.R., 2016. Single slice US-MRI registration for neurosurgical MRI-guided US, p. 97862D. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2214265>, doi:10.1117/12.2214265.
- Park, H., Meyer, C.R., Kim, B., 2004. Improved Motion Correction in fMRI by Joint Mapping of Slices into an Anatomical Volume. International Conference on Medical Image Computing and Computer Aided Intervention (MICCAI) , 745–751.
- Penney, G., Blackall, J., Hamady, M., Sabharwal, T., Adam, A., Hawkes, D., 2004a. Registration of freehand 3D ultrasound and magnetic resonance liver images. Medical Image Analysis 8, 81–91. URL:

<http://linkinghub.elsevier.com/retrieve/pii/S1361841503000744>,
doi:10.1016/j.media.2003.07.003.

- Penney, G.P., Barratt, D.C., Chan, C.S.K., Slomczykowski, M., Carter, T.J., Edwards, P.J., Hawkes, D.J., 2006. Cadaver validation of intensity-based ultrasound to CT registration. *Medical image analysis* 10, 385–95. URL: <http://www.ncbi.nlm.nih.gov/pubmed/16520083>, doi:10.1016/j.media.2006.01.003.
- Penney, G.P., Schnabel, J.A., Rueckert, D., Viergever, M.A., Niessen, W.J., 2004b. Registration-based interpolation. *IEEE Transactions on Medical Imaging* 23, 922–926.
- Po Su, Jianhua Yang, Kongkuo Lu, Nam Yu, Wong, S.T., Zhong Xue, 2013. A Fast CT and CT-Fluoroscopy Registration Algorithm With Respiratory Motion Compensation for Image-Guided Lung Intervention. *IEEE Transactions on Biomedical Engineering* 60, 2034–2041. doi:10.1109/TBME.2013.2245895.
- Porchetto, R., Stramana, F., Paragios, N., Ferrante, E., 2016. Rigid Slice-To-Volume Medical Image Registration through Markov Random Fields. *BAMBI Workshop, MICCAI* .
- Rivaz, H., Chen, S., Collins, D.L., 2014a. Automatic Deformable MR-Ultrasound Registration for Image-Guided Neurosurgery. *IEEE transactions on medical imaging* 0062, 1–15. doi:10.1109/TMI.2014.2354352.
- Rivaz, H., Karimaghloo, Z., Collins, D.L., 2014b. Self-similarity weighted mutual information: A new nonrigid image registration metric. *Medical Image Analysis* 18, 343–358. doi:10.1016/j.media.2013.12.003.
- Rivaz, H., Karimaghloo, Z., Fonov, V.S., Collins, D.L., 2014c. Nonrigid registration of ultrasound and MRI using contextual conditioned mutual information. *IEEE Transactions on Medical Imaging* 33, 708–725. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6687247>, doi:10.1109/TMI.2013.2294630.
- Roche, A., Malandain, G., Pennec, X., Ayache, N., 1998. The correlation ratio as a new similarity measure for multimodal image registration, pp. 1115–1124. URL: <http://link.springer.com/10.1007/BFb0056301>, doi:10.1007/BFb0056301.
- Rohlfing, T., 2012. Image similarity and tissue overlaps as surrogates for image registration accuracy: Widely used but unreliable. *IEEE Transactions on Medical Imaging* 31, 153–163. doi:10.1109/TMI.2011.2163944.
- Rother, C., Kolmogorov, V., Lempitsky, V., Szummer, M., 2007. Optimizing Binary MRFs via Extended Roof Duality, in: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE. pp. 1–8. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4270228>, doi:10.1109/CVPR.2007.383203.

- Rousseau, F., Glenn, O., 2005. A novel approach to high resolution fetal brain MR imaging. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* 3749, 548–555.
- Rousseau, F., Glenn, O.A., Iordanova, B., Rodriguez-Carranza, C., Vigneron, D.B., Barkovich, J.A., Studholme, C., 2006. Registration-Based Approach for Reconstruction of High-Resolution In Utero Fetal MR Brain Images. *Academic Radiology* 13, 1072–1081. URL: <http://www.sciencedirect.com/science/article/pii/S1076633206002753>, <http://linkinghub.elsevier.com/retrieve/pii/S1076633206002753>, doi:10.1016/j.acra.2006.05.003.
- Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J., 1999. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE transactions on medical imaging* 18, 712–21. URL: <http://www.ncbi.nlm.nih.gov/pubmed/10534053>, doi:10.1109/42.796284.
- San José Estépar, R., Westin, C.F., Vosburgh, K.G., 2009. Towards real time 2D to 3D registration for ultrasound-guided endoscopic and laparoscopic procedures. *International Journal of Computer Assisted Radiology and Surgery* 4, 549–560. URL: <http://link.springer.com/10.1007/s11548-009-0369-z>, doi:10.1007/s11548-009-0369-z.
- Schulz, J., Siegert, T., Bazin, P.L., Maclaren, J., Herbst, M., Zaitsev, M., Turner, R., 2014. Prospective slice-by-slice motion correction reduces false positive activations in fMRI with task-correlated motion. *NeuroImage* 84, 124–132. doi:10.1016/j.neuroimage.2013.08.006.
- Sederberg, T.W., Parry, S.R., 1986. Free-form deformation of solid geometric models. *ACM SIGGRAPH Computer Graphics* 20, 151–160. URL: <http://portal.acm.org/citation.cfm?doid=15886.15903>, doi:10.1145/15886.15903.
- Seshamani, S., Fogtman, M., Cheng, X., Thomason, M., Gatenby, C., Studholme, C., 2013. Cascaded slice to volume registration for moving fetal FMRI, in: *IEEE International Symposium on Biomedical Imaging (ISBI)*, IEEE. pp. 796–799. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6556595>.
- Shakeri, M., Tsogkas, S., Ferrante, E., Lippe, S., Kadoury, S., Paragios, N., Kokkinos, I., Paris-saclay, U., Montreal, P., 2016. Sub-cortical brain structure segmentation using F-CN, in: *ISBI 2016*.
- Shekhovtsov, A., Kovtun, I., Hlaváč, V., 2008. Efficient MRF deformation model for non-rigid image matching. *Computer Vision and Image Understanding* 112, 91–99. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1077314208000970>, doi:10.1016/j.cviu.2008.06.006.
- Shepp, L.A., Logan, B.F., 1974. The fourier reconstruction of a head section. *IEEE Transactions on Nuclear Science* 21, 21–43.

- Shimony, S.E., 1994. Finding MAPs for belief networks is NP-hard. *Artificial Intelligence* .
- Simonovsky, M., Gutiérrez-Becker, B., Mateus, D., Navab, N., Komodakis, N., 2016. *A Deep Metric for Multimodal Registration*. Springer International Publishing, Cham.
- Smolíková Wachowiak, R., Wachowiak, M.P., Fenster, A., Drangova, M., 2005. Registration of two-dimensional cardiac images to preprocedural three-dimensional images for interventional applications. *Journal of Magnetic Resonance Imaging* 22, 219–228.
- Sotiras, A., Davatzikos, C., Paragios, N., 2013. Deformable Medical Image Registration: A Survey. *IEEE Transactions on Medical Imaging* 32, 1153–1190.
- Sotiras, A., Ou, Y., Glocker, B., Davatzikos, C., Paragios, N., 2010. Simultaneous geometric–iconic registration. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)* 13, 676–683.
- Studholme, C., Hill, D., Hawkes, D., 1999. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition* 32, 71–86. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0031320398000910>, doi:10.1016/S0031-3203(98)00091-0.
- Sun, Y., Kadoury, S., Li, Y., John, M., Resnick, J., Plambeck, G., Liao, R., Sauer, F., Xu, C., 2007. Image guidance of intracardiac ultrasound with fusion of pre-operative images. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)* 10, 60–67. doi:10.1007/978-3-540-75757-3.
- Szegedy, C., Toshev, A., Erhan, D., 2013. Deep Neural Networks for Object Detection, in: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (Eds.), *Advances in Neural Information Processing Systems 26*. Curran Associates, Inc., pp. 2553–2561.
- Tadayyon, H., 2010. MRI-Guided Prostate Motion Tracking using Multislice-to-Volume Registration .
- Tadayyon, H., Lasso, A., Gill, S., Kaushal, A., Guion, P., Fichtinger, G., 2010. Target Motion Compensation in MRI-guided Prostate Biopsy with Static Images, in: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, pp. 1199–1216. doi:10.1016/j.micinf.2011.07.011.Innate.
- Tadayyon, H., Lasso, A., Kaushal, A., Guion, P., Fichtinger, G., 2011. Target motion tracking in MRI-guided transrectal robotic prostate biopsy. *IEEE transactions on biomedical engineering* 58, 3135–42. doi:10.1109/TBME.2011.2163633.
- Taron, M., Paragios, N., Jolly, M.P., 2009. Registration with Uncertainties and Statistical Modeling of Shapes with Variable Metric Kernels. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31, 99–113. doi:10.1109/TPAMI.2008.36.

- Wainwright, M., Jaakkola, T., Willsky, A., 2005. MAP Estimation Via Agreement on Trees: Message-Passing and Linear Programming. *IEEE Transactions on Information Theory* 51, 3697–3717. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1522634>, doi:10.1109/TIT.2005.856938.
- Wang, C., Komodakis, N., Paragios, N., 2013. Markov Random Field modeling, inference & learning in computer vision & image understanding: A survey. *Computer Vision and Image Understanding* 117, 1610–1627. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1077314213001343>, doi:10.1016/j.cviu.2013.07.004.
- Wein, W., Brunke, S., Khamene, A., Callstrom, M.R., Navab, N., 2008a. Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Medical Image Analysis* 12, 577–585. doi:10.1016/j.media.2008.06.006.
- Wein, W., Cheng, J., Khamene, A., 2008b. Ultrasound based respiratory motion compensation in the abdomen. *Workshop on Image Guidance and Computer Assistance for Softtissue Interventions. MICCAI* 32, 294. URL: http://www.igsti08.de/paper/igsti08_011.pdf.
- Wein, W., Khamene, A., Clevert, D.A., Kutter, O., Navab, N., 2007. Simulation and Fully Automatic Multimodal Registration of Medical Ultrasound, in: *Medical Image Computing and Computer-Assisted Intervention MICCAI 2007*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 136–143. URL: http://link.springer.com/10.1007/978-3-540-75757-3_17, doi:10.1007/978-3-540-75757-3_17.
- Xiao, G., Bloch, B.N., Chappelow, J., Genega, E.M., Rofsky, N.M., Lenkinski, R.E., Tomaszewski, J., Feldman, M.D., Rosen, M., Madabhushi, A., 2011. Determining histology-MRI slice correspondences for defining MRI-based disease signatures of prostate cancer. *Computerized Medical Imaging and Graphics* 35, 568–578. URL: <http://dx.doi.org/10.1016/j.compmedimag.2010.12.003>, doi:10.1016/j.compmedimag.2010.12.003.
- Xiao, Y., Yan, C.X.B., Drouin, S., de Nigris, D., Kochanowska, A., Collins, D.L., 2016. User-friendly freehand ultrasound calibration using Lego bricks and automatic registration. doi:10.1007/s11548-016-1368-5.
- Xiaolei Huang, Paragios, N., Metaxas, D., 2006. Shape registration in implicit spaces using information theory and free form deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28, 1303–1318. doi:10.1109/TPAMI.2006.171.
- Xu, H., Lasso, A., Fedorov, A., Tuncali, K., Tempany, C., Fichtinger, G., 2014a. Multi-slice-to-volume registration for MRI-guided transperineal prostate biopsy. *International Journal of Computer Assisted Radiology and Surgery* 10, 563–572.

- Xu, L., Liu, J., Zhan, W., Gu, L., 2013. A novel algorithm for CT-ultrasound registration. *IEEE Point-of-Care Healthcare Technologies (PHT)* , 101–104 URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6461294>, doi:10.1109/PHT.2013.6461294.
- Xu, R., Athavale, P., Nachman, A., Wright, G.A., 2014b. Multiscale Registration of Real-Time and Prior MRI Data for Image-Guided Cardiac Interventions. *IEEE Transactions on Biomedical Engineering* 61, 2621–2632. doi:10.1109/TBME.2014.2324998.
- Xu, S., Kruecker, J., Turkbey, B., Glossop, N., Singh, A.K., Choyke, P., Pinto, P., Wood, B.J., 2008. Real-time MRI-TRUS fusion for guidance of targeted prostate biopsies. *Computer Aided Surgery* 13, 255–264. URL: <http://www.tandfonline.com/doi/full/10.3109/10929080802364645>, doi:10.3109/10929080802364645.
- Yan, C.X.B., Goulet, B., Tampieri, D., Collins, D.L., 2012. Ultrasound-CT registration of vertebrae without reconstruction. *International Journal of Computer Assisted Radiology and Surgery* 7, 901–909. doi:10.1007/s11548-012-0771-9.
- Yang, M., Ding, H., Kang, J., Zhu, L., Wang, G., Yang, M., Ding, H., Wang, G., Kang, J., Zhu, L., Shen, B., 2015. Subject-specific real-time respiratory liver motion compensation method for ultrasound-MRI/CT fusion imaging. *Int J CARS* 10, 517–529. URL: <http://dx.doi.org/10.1007/s11548-014-1085-x>, doi:10.1007/s11548-014-1085-x.
- Yang, X., Kwitt, R., Niethammer, M., 2016. Fast Predictive Image Registration, in: *Deep Learning and Data Labeling for Medical Applications Workshop (MICCAI, 2016)*.
- Yavariabdi, A., Bartoli, A., Samir, C., Artigues, M., Canis, M., 2015. Mapping and Characterizing Endometrial Implants by Registering 2D Transvaginal Ultrasound to 3D Pelvic Magnetic Resonance Images. *Computerized Medical Imaging and Graphics* 45, 11–25. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0895611115001056>, doi:10.1016/j.compmedimag.2015.07.007.
- Yeo, D.T., Fessler, J.A., Kim, B., 2004. Motion Correction in fMRI by Mapping Slice-to-Volume with Concurrent Field-Inhomogeneity Correction, in: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 752–760.
- Yeo, D.T.B., Bhargava, R.R., Kim, B., 2006. Improved map-slice-to-volume motion correction with B0 inhomogeneity correction: validation of activation detection algorithms using ROC curve analyses. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)* 9, 276–83. URL: <http://www.ncbi.nlm.nih.gov/pubmed/17354782>.
- Yeo, D.T.B., Fessler, J.a., Kim, B., 2008. Concurrent correction of geometric distortion and motion using the map-slice-to-volume method in echo-planar imaging. *Magnetic resonance imaging* 26, 703–14. doi:10.1016/j.mri.2007.11.001.

- Yu, Y., Trouve, A., Chalmond, B., 2008. A Fast 3D Volume Reconstruction for Confocal Micro-rotation Cell Imaging. IEEE International Conference on BioMedical Engineering and Informatics , 775–778doi:10.1109/BMEI.2008.201.
- Yu, Y., Trouvé, a., Chalmond, B., Renaud, O., Shorte, S.L., 2011. Confocal bi-protocol: A new strategy for isotropic 3D live cell imaging. Journal of Microscopy 242, 70–85. doi:10.1111/j.1365-2818.2010.03440.x.
- Zagoruyko, S., Komodakis, N., Paristech, P., Paris Est, U., 2015. Learning to Compare Image Patches via Convolutional Neural Networks. CVPR 2015 .
- Zakkaroff, C., Radjenovic, A., Greenwood, J., Magee, D., 2012. Recovery of Slice Rotations with the Stack Alignment Transform in Cardiac MR Series, in: British Machine Vision Conference (BMVC), British Machine Vision Association. pp. 35.1–35.11. URL: <http://www.bmva.org/bmvc/2012/BMVC/paper035/index.html>, doi:10.5244/C.26.35.
- Zarow, C., Kim, T.S., Singh, M., Chui, H.C., 2004. A standardized method for brain-cutting suitable for both stereology and MRI-brain co-registration. Journal of Neuroscience Methods 139, 209–215. doi:10.1016/j.jneumeth.2004.04.034.
- Zikic, D., Glocker, B., Kutter, O., Groher, M., Komodakis, N., Kamen, A., Paragios, N., Navab, N., 2010a. Linear intensity-based image registration by Markov random fields and discrete optimization. Medical image analysis 14, 550–62. doi:10.1016/j.media.2010.04.003.
- Zikic, D., Glocker, B., Kutter, O., Groher, M., Komodakis, N., Khamene, A., Paragios, N., Navab, N., 2010b. Markov random field optimization for intensity-based 2D-3D registration, in: SPIE Medical Imaging. International Society for Optics and Photonics. doi:10.1117/12.837232.