

Rigid Slice-To-Volume Medical Image Registration Through Markov Random Fields

Roque Porchetto¹, Franco Stramana¹, Nikos Paragios², and Enzo Ferrante²(✉)

¹ UNICEN University, Tandil, Argentina

² CVN, CentraleSupélec-INRIA, Université Paris-Saclay,
Châtenay-Malabry, France
ferrante.enzo@gmail.com

Abstract. Rigid slice-to-volume registration is a challenging task, which finds application in medical imaging problems like image fusion for image guided surgeries and motion correction for volume reconstruction. It is usually formulated as an optimization problem and solved using standard continuous methods. In this paper, we discuss how this task be formulated as a discrete labeling problem on a graph. Inspired by previous works on discrete estimation of linear transformations using Markov Random Fields (MRFs), we model it using a pairwise MRF, where the nodes are associated to the rigid parameters, and the edges encode the relation between the variables. We compare the performance of the proposed method to a continuous formulation optimized using simplex, and we discuss how it can be used to further improve the accuracy of our approach. Promising results are obtained using a monomodal dataset composed of magnetic resonance images (MRI) of a beating heart.

1 Introduction

Slice-to-volume registration has received increasing attention during the last decades within the medical imaging community. Given a tomographic 2D slice and a 3D volume as input, this challenging problem consists in finding the slice (extracted from the input volume and specified by an arbitrary rigid transformation) that best matches the 2D input image. We stress the fact that we are working with 2D slices (e.g. ultrasound (US)) as opposed to projective 2D images (e.g. X-ray images). This is important since both problems are usually refereed as 2D/3D registration, even if they are intrinsically different. In slice-to-volume registration, every pixel from the 2D image corresponds to a single voxel in 3D space. However, in a projective 2D image every pixel represents a projection of voxels from a given viewpoint.

One can formulate different versions of slice-to-volume registration, depending on several aspects of the problem such as the *matching criterion* used to determine the similarity between the images, the *transformation model* we aim at estimating, the *optimization strategy* used to infer the optimal transformation model (continuous or discrete) and the *number of slices* given as input. In this work, we propose an iconic method (where the matching criterion is defined as

a function of the image intensity values) to infer rigid transformation models (specified using 6-DOF). The input consists of a single slice and a single volume, and we formulate it as a discrete optimization problem.

Discrete methods, where the tasks are usually formulated as a discrete labeling problem on a graph, have become a popular strategy to model vision problems [24] (and particularly, biomedical vision problems [19]) thanks to their modularity, efficiency, robustness and theoretical simplicity. This paper presents a graph-based formulation (inspired by the work of [26,27]) to solve rigid (only) slice-to-volume registration using discrete methods. As we will discuss in Sect. 1.2, other works have tackled similar problems. However, to date, no work has shown the potential of discrete methods to deal with rigid slice-to-volume registration. Our main contribution is to put a new spin on graph-based registration theory, by demonstrating that discrete methods and graphical models are suitable to estimate rigid transformations mapping slice-to-volume. We validate our approach using a dataset of magnetic resonance images (MRI) of the heart, and we compare its performance with a state-of-the-art approach based on continuous optimization using simplex method. Moreover, in the spirit of encouraging reproducible research, we make the source code of the application publicly available at the following website: <https://gitlab.com/franco.stramana1/slice-to-volume>.

1.1 Motivation

In the extensive literature of medical image registration, it is possible to identify two main problems which motivated the development of slice-to-volume registration methods during the last decades. The first one is the fusion of pre-operative high-definition volumetric images with intra-operative tomographic slices to perform diagnostic and minimally invasive surgeries. In this case, slice-to-volume registration is one of the enabling technologies for computer-aided image guidance, bringing high-resolution pre-operative data into the operating room to provide more realistic information about the patient's anatomy [17]. This technique has been used when dealing with several scenarios such as liver surgery [1], radio-frequency thermal ablation of prostate cancer [5], minimally invasive cardiac procedures [12], among many others.

The second problem is the correction of slicewise motion in volumetric acquisitions. In a variety of situations, inter slice motion may appear when capturing a volumetric image. For example, in case of fetal brain imaging (essential to understand neurodevelopmental disabilities in childhood and infancy) [21], fetus motion generates inconsistencies due to the slice acquisition time. Another case is related to functional magnetic resonance images (fMRI), usually acquired as time series of multislice single-shot echoplanar images (EPI). Patient head motion during the experiments may introduce artifacts on activation signal analysis. Slice-to-volume registration can be used to alleviate this problem by registering every slice to an anatomical volumetric reference following the well-know map-slice-to-volume (MSV) method [13].

1.2 Previous Work

Graph-based image registration, where the task is conceived as a discrete labeling problem on a graph, constitutes one of the most efficient and accurate state-of-the-art methods for image registration [22]. Even if they have shown to be particularly suitable to estimate deformable non-linear transformations [10, 11], they were also adapted to the linear case [27]. Most of the publications on the field focus on registering images which are in dimensional correspondence (2D/2D or 3D/3D). In case of projective 2D/3D image registration, only linear transformations were estimated using discrete methods by [26, 27]. More recently, several graph-based approaches to perform deformable slice-to-volume registration were introduced in [6–8]. In these works, rigid transformations were computed as a by-product of the deformable parameters, leading to unnecessary computational burden (since rigid transformations are by far lower dimensional than deformable ones). To the best of our knowledge, rigid (only) slice-to-volume registration has not been formulated within this powerful framework. To date, all the methods focusing on this challenging problem are based on continuous (e.g. simplex [5], gradient descent [21], Powell’s method [9], etc.) or heuristic approaches (evolutionary algorithms [23], simulated annealing [2]), missing the aforementioned advantages offered by discrete optimization. Based on the work of [26], we propose a discrete Markov Random Field (MRF) formulation of this problem, delivering more precise results than the state-of-the-art continuous approaches. Moreover, inspired by the work of [16] in the context of vector flow estimation, we discuss how continuous state-of-the-art approaches can be used to further refine the rigid transformations obtained through discrete optimization, resulting in more accurate solutions.

2 Rigid Slice-to-Volume Registration Through Markov Random Fields

We formulate rigid slice-to-volume registration as an optimization problem. Given a 2D image $I : \Omega_I \in \mathbb{R}^2 \rightarrow \mathbb{R}$, and a 3D image $J : \Omega_J \in \mathbb{R}^3 \rightarrow \mathbb{R}$, we aim at recovering the rigid transformation specified by $\pi = (r_x, r_y, r_z, t_x, t_y, t_z)$ that better aligns both images, by solving:

$$\hat{\pi} = \underset{\pi}{\operatorname{argmin}} \mathcal{M}(I, \pi[J]), \quad (1)$$

where $\pi[J]$ corresponds to the slice extracted from image J (using trilinear interpolation) and specified by the rigid transformation π (as explained in Fig. 1). \mathcal{M} is the so-called matching criteria, that quantifies the dissimilarity between the 2D image I and the slice $\pi[J]$. Alternative matching criteria can be adopted depending on the type of images we are registering. For example, in monomodal cases where intensities tend to be linearly correlated in both images, simple functions such as sum of absolute differences (SAD) or sum of squared differences (SSD) may make the job. However, for more complicated cases like multimodal




Fig. 1. Interpretation of the components of Eq. 1. (a) Image I corresponds to the input 2D image, which is moved by different rigid transformations π . (b) Image J corresponds to the 3D image. Given a rigid transformation π , a slice $\pi[J]$ is extracted (using trilinear interpolation). In that way, it is possible to explore the space of solutions by sampling several rigid transformations π . (c) The matching criterion \mathcal{M} quantifies the dissimilarity of both 2D images, I and $\pi[J]$. Higher values indicate dissimilar images while lower values indicate better alignment.

registration (where the relation between intensity values in both images is usually non-linear), more elaborated functions like mutual information (MI) are applied.

This optimization problem is commonly solved through continuous (gradient or non-gradient based) methods, which are considerably sensible to initialization and may be stuck in local minima. As discussed in Sect. 1.2, in this work we model rigid slice-to-volume registration as a discrete labeling problem following the discretization strategy proposed by [27].

2.1 Rigid Slice-to-Volume Registration as a Discrete Labeling Problem

Rigid slice-to-volume registration, as well as many other problems in computer vision, can be formulated as a discrete labeling problem on a pairwise Markov Random Field (MRF) [24]. Formally, a discrete pairwise MRF is an undirected graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$, where each node $v_i \in \mathcal{V}, i = 1 \dots |\mathcal{V}|$ represents a discrete variable. Any two variables v_i, v_j depend on each other if there is an edge $(v_i, v_j) \in \mathcal{E}$ linking the corresponding nodes. The range of values that can be assigned to a discrete variable is determined by the label space L . A discrete labeling problem on a pairwise MRF consists on assigning a label $l_i \in L$ to every $v_i \in \mathcal{V}$, such that the following energy is minimized:

$$\mathcal{P}(\mathbf{x}; G, F) = \sum_{v_i \in \mathcal{V}} g_i(l_i) + \sum_{(v_i, v_j) \in \mathcal{E}} f_{ij}(l_i, l_j), \quad (2)$$

where $\mathbf{x} = \{l_1, \dots, l_n\}$ is a labeling assigning a label l_i to every $v_i \in \mathcal{V}$, $G = \{g_i(\cdot)\}$ are the unary potentials associated to $v_i \in \mathcal{V}$ and $F = \{f_{ij}(\cdot, \cdot)\}$ are the pairwise potentials associated to edges $(v_i, v_j) \in \mathcal{E}$. These functions return scalar values when labels l_i are assigned to variables v_i . Since we pose the optimization as a minimization problem, potentials must associate lower values to labelings representing good solutions, and higher values otherwise.

In the formulation presented in Eq. 1, one would like to explore the space of parameters π and chose the values giving the best matching. Since we are modeling it as a discrete problem, we must adopt a discretization strategy for the (naturally) continuous space of rigid transformations π . In [27], authors proved that it is possible to estimate linear (an particularly, rigid body) transformations by solving a discrete and approximated version of this formulation. Following their proposal, we model rigid slice-to-volume registration through a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, associating every parameter of the rigid transformation $\pi = (r_x, r_y, r_z, t_x, t_y, t_z)$ to a variable $v_i \in \mathcal{V}$, giving a total of 6 variables (nodes in the graph). \mathcal{G} is a fully connected pairwise graph where $\mathcal{E} = \{(v_i, v_j)\}, \forall i \neq j$, meaning that all variables (parameters) depend on each other. Note that this pairwise model is clearly an approximation, since the real dependency between the parameters is not pairwise but high-order. However, as stated in [27], similar approximations have shown to be good enough to estimate linear transformations, while making the problem tractable.

In our discrete strategy, every parameter v_i is updated through a discrete variation d_{l_i} associated to the label l_i . Given an initial transformation $\pi_0 = (r_x^0, r_y^0, r_z^0, t_x^0, t_y^0, t_z^0)$, we explore the space of solutions by sampling discrete variations of π_0 , and choosing the one that generates the slice $\pi[J]$ best matching image I . Therefore, for a maximum size ω_i and a quantization factor κ_i , we consider the following variations to the initial estimate of v_i : $\{0, \pm \frac{\omega_i}{\kappa_i}, \pm \frac{2\omega_i}{\kappa_i}, \pm \frac{3\omega_i}{\kappa_i}, \dots, \pm \frac{\kappa_i \omega_i}{\kappa_i}\}$. The total number of labels results $|L| = 2\kappa + 1$. Note that 0 is always included since we give the possibility of keeping the current parameter estimate. For example, in case that v_0 corresponds to r_x , $\omega_0 = 0.2$ and $\kappa_0 = 2$, the label space of v_0 will correspond to $\{r_x^0, r_x^0 \pm 0.1, r_x^0 \pm 0.2\}$.

Ideally, we would like to explore the complete search space around π_0 given by all possible combinations of labels. Since we have an exponential number of potential solutions, we adopt a pairwise approximation where only variations for pairs of variables are considered. This variations are encoded in the pairwise terms of the energy defined in Eq. 2 as $f_{ij}(l_i, l_j) = \mathcal{M}(I, \pi_{l_i, l_j}[J])$. Here π_{l_i, l_j} denotes the updated version of π_0 , where only v_i and v_j were modified according to the labels l_i and l_j , while the rest of the parameters remained fixed. Unary potentials g_i are not considered since we are only interested in the interaction between variables. Therefore, the discrete version of the optimization problem introduced in Eq. 1 becomes:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \mathcal{P}(\mathbf{x}; F) = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{(v_i, v_j) \in \mathcal{E}} \mathcal{M}(I, \pi_{l_i, l_j}[J]), \quad (3)$$

where the optimal labeling $\hat{\mathbf{x}}$ represents the final rigid transformation $\hat{\pi}$ used to extract the solution slice $\hat{\pi}[J]$.

2.2 Discrete Optimization

We solve the discrete multi-labeling problem from Eq. 3 using FastPD. FastPD is a discrete optimization algorithm based on principles from linear programming and primal dual strategies, which at the same time generalizes the well known α -expansion [15]. One of the main advantages of FastPD is its modularity/scalability, since it deals with a much wider class of problems than α -expansion, being an order of magnitude faster while providing the same optimality guarantees when performing metric labeling [14]. Our problem does not fulfill the conditions to be considered a metric labeling problem (we refer the reader to [4] for a complete discussion about metric labeling); however, FastPD has shown promising results for similar formulations [27].

FastPD solves a series of max-flow min-cut [3] problems on a graph. In that sense, it is similar to α -expansion which also performs discrete inference on multi-label problems by solving successive binary max-flow min-cut problems. The main difference between these approaches is the construction of the graph where max-flow min-cut algorithm is applied. α -expansion constructs the binary problem by restricting the label space, so that the only options for a given variable are to remain in its current assignment, or to take a label α (which varies in every iteration). Instead, FastPD constructs these binary problems by performing a Linear Programming Relaxation (LPR) of the integer program that represents the discrete MRF formulation.

2.3 Incremental Approach

Discrete approximations of continuous spaces usually suffer from low accuracy (since it is bounded by the quality of the discretization). Thus, we adopt an incremental approach to explore the space of solutions in a finer way. The idea is to successively solve the problem from Eq. 3, using the solution from time t as initialization for time $t + 1$, keeping a fixed number of labels but decreasing the maximum sizes ω_i in a factor α_i . Moreover, we also adopt a pyramidal approach, where we generate a Gaussian pyramid for both input images I and J , and we run the complete incremental approach on every level of the pyramid. In that way, we increase the capture range of our method.

2.4 Simplex Refinement Step

Let us advance one of the conclusions of this work, so that we can motivate the last step of our approach. In Sect. 3.1, we compare the performance of our discrete approach with a continuous method based on simplex [18] algorithm. As we will see, when the initialization π_0 is good enough, both continuous and discrete approaches perform well. In fact, in some cases, simplex is delivering more accurate solutions than discrete. However, as we move away from the initialization, discrete optimization gives more and more significant improvements, thanks to its wider capture range. In order to improve the accuracy of our proposal, and inspired by similar conclusions discussed by [16] in the context of vector flow




Fig. 2. Visual results for two slices of the individual tests. The first column corresponds to the input 2D slice. The second column shows the difference between the input 2D slice and the initial slice. The other columns show the difference between the input and the one resulting slices applying simplex, discrete and refined approaches. Grey values indicate no difference, while white and black value indicate inconsistencies. As it can be observed, the solution given by the refined approach is outperforming the others.

estimation, we refine the results obtained with our approach by optimizing Eq. 1 through simplex, using the discrete solution as initialization.

3 Experiments and Results

In this section, we present the results obtained using the proposed method (considering also the refined version), and we compare them with a state-of-the-art approach based on continuous optimization trough downhill simplex [18] (also known as Nelder-Mead or amoeba method). Simplex is one of the most popular optimization algorithms used to deal with rigid slice-to-volume registration (some examples are [2, 13, 20, 25]). It is a continuous and derivative-free method, which relies on the notion of simplex (which is a special polytope of $n + 1$ vertices living in a n -dimensional space) to explore the space of solutions in a systematic way. We used a dataset composed of MRI images of a beating heart. Given an initial sequence of 3D images $M_i, i = 0..19$ of a beating heart (with a resolution of $192 \times 192 \times 11$ voxels and a voxel size of $1.25 \times 1.25 \times 8$ mm), we generated slices which were used for two different experiments.

3.1 Implementation Details and Parameters Setting

We implemented the three versions of the algorithms discussed in this paper (simplex, discrete and refined) mainly using Python and ITK for image manipulation¹. For simplex optimization we used the method implemented

¹ The source code can be downloaded from <https://gitlab.com/franco.stramana1/slice-to-volume>.

in *scipy.optimize* package, while discrete optimization was performed using a Python wrapped version of the standard C++ implementation of FastPD. In all the experiments, we used a pyramidal approach with 4 Gaussian levels (3D images were not downsampled in z axis because of the low resolution of the images in this direction). The matching criterion adopted in all the experiments was the sum of squared differences, since pixel intensities are equivalent in both 2D and 3D images. The matching criterion \mathcal{M} based on SSD is simply defined as:

$$\mathcal{M}(I_1, I_2) = \sum_{i \in \Omega} (I_1(i) - I_2(i))^2, \quad (4)$$

For the discrete case, at every pyramid level we decreased the maximum label size for both, rotation ($\omega_{rot} = [0.02, 0.015, 0.0125, 0.01]$ rad) and translation ($\omega_{trans} = [7, 6.5, 6, 5]$ mm) parameters. Starting from these maximum sizes, we solved Eq. 3 running FastPD several times per level until no improvement is produced or a maximum number of iterations is achieved ([200, 100, 150, 600]), using different label space decreasing factors at every pyramid level ($\alpha = [0.08, 0.07, 0.05, 0.03]$). The total number of labels was fixed to 5 ($\kappa = 2$) for all the experiments. For the continuous case (where Eq. 1 was optimized using simplex), we used again a 4-levels pyramidal approach, with simplex running until convergence in every level. Finally, for the refined experiment, we just ran the simplex experiment initialized with the solution estimated with the discrete method. For every registration case, continuous approach took around 30secs while the discrete version took 9mins, running on a laptop with an Intel i7-4720HQ and 16 GB of RAM.

3.2 Experiments

We performed two different type of experiments, considering individual registration cases as well as image series. For validation, we measured three different indicators: the distance in terms of translation and rotation parameters between the estimated and ground truth transformations, together with the mean of absolute differences (MAD) between the input 2D image and the slice specified by the estimated rigid transformation.

Individual tests. The first set of experiments measures the accuracy of the three approaches using individual tests, where 100 random slices extracted from the 20 volumes are considered as single images (independently of the series), and registered to the first volume M_0 . We run the same experiment for every slice using three different initializations (resulting in 300 registration cases), where ground truth parameters were randomly perturbed in three different ranges ([5, 12), [12, 18), [18, 25) millimeters for translation and [0.1, 0.2), [0.2, 0.3), [0.3, 0.4) radians for rotation parameters) to guarantee that both, good and bad initializations, are considered for every slice. Quantitative results are reported in




Fig. 3. Individual tests where 100 2D slices (extracted at locations specified using random rigid transformations) are considered as independent registration cases. Every point from the scatter plot represents the mean of absolute differences (MAD) between the input 2D image and the slice extracted at the initial position (X axis) vs the estimated position (Y axis). We also include a linear trend estimation (fitted using least squares method) to compare the robustness of the method to bad initializations.

Table 1. Average error estimated in terms of rotation (expressed in radians), translation (expressed in millimeters) and MAD for the three alternative approaches discussed in this paper. As it can be observed, the discrete and refined methods outperform the standard continuous approach optimized through simplex.

Method	R_x	R_y	R_z	T_x	T_y	T_z	MAD
Simplex	0,14	0,13	0,12	8,46	9,62	10,69	51,88
Discrete	0,12	0,08	0,09	5,87	6,72	6,18	42,36
Refined	0,10	0,07	0,08	5,09	5,96	4,92	36,45

Figs. 3 and 4 and summarized in Table 1. Visual results for qualitative evaluation are reported in Fig. 2.

Results in the scatter plot from Fig. 3 indicate that, as we go farther away from the initialization (in this case, it is quantified by the MAD between the input 2D image and the slice corresponding to the initialization), discrete and refined methods tend to be more robust. This robustness is clearly reflected by the slope of the trend lines: the refined method presents the trend line with the lower slope, meaning that even for bad initializations it converges to better solutions. The boxplot from Fig. 4 and the numerical results from Table 1 confirm that discrete and refined methods perform better not only in terms of MAD, but also with respect to the distance between the rotation/translation estimated and ground truth parameters.




Fig. 4. Boxplot corresponding to the estimated error (in terms of rotation and translation parameters) for the 300 individual tests. As it can be observed, discrete and refined approaches are reducing both the mean error (shown as a dotted line in every box) and the dispersion.

Temporal series test. The idea behind the second experiment is to simulate an image guided surgery (IGS) scenario, where a fixed pre-operative volume must be fused with consecutive intra-operative 2D images suffering deformations (in this case, due to heart beating). Given the temporal sequence of 20 volumetric MRI images M_i , $i = 0..19$, we generated a sequence of 20 2D slices to validate our method. It was extracted as in [8]: starting from a random initial translation $T_0 = (T_{x_0}, T_{y_0}, T_{z_0})$ and rotation $R_0 = (R_{x_0}, R_{y_0}, R_{z_0})$, we extracted a 2D slice I_0 from the initial volume M_0 . Gaussian noise was added to every parameter in order to generate the position used to extract the next slice from the next volume. We used $\sigma_r = 3^\circ$ for the rotation and $\sigma_t = 5\text{mm}$ for the translation parameters. All the slices were registered to the first volume M_0 . The solution of the registration problem for slice I_i was used as initialization for the slice I_{i+1} . The first experiment was initialized randomly perturbing its ground truth transformation with the same noise parameters. As it can be observed in Fig. 5, discrete and refined approaches manage to keep a good estimation error while simplex can not. Note that different strategies could be used in real scenarios to obtain good initializations for the first slice. For example, in IGS, physicians could start from a plane showing always the same anatomical structure, or identify landmark correspondences in the first slice and the 3D image useful to estimate an initial transformation.




Fig. 5. Results for the temporal series experiment. In this case, the transformation estimated for the slice i was used as initialization for the next slice of the series. We reported results in terms of MAD and rotation/translation error for the estimated transformations using the three approaches.

4 Discussion, Conclusions and Future Works

In this paper we presented a strategy to solve rigid slice-to-volume registration as a discrete graph labeling problem, following the discretization strategy introduced by [27]. We validated our proposal using a MRI dataset of a beating heart, where arbitrary 2D slices are fused with a 3D volume. The experimental results showed that our discrete approach produces more accurate and robust estimates for the rigid transformations than a continuous method based on simplex. Moreover, they also reflected that results obtained using such a method can be further refined using a continuous approach like simplex, leading to even more accurate estimations. This is coherent with the conclusions presented by [16] for the case of optical flow estimation.

An interesting discussion about the limitations of our approach, emerges when we observe the results obtained in previous work by [6–8] using similar

images. In these works, both rigid and local deformable parameters are estimated in a one-shot discrete optimization procedure, delivering results which are considerably better than ours, even for the refined approach. Since we are dealing with 2D images which are deformed with respect to the static 3D volume (due to heart beating), estimating both rigid and deformable parameters at the same time seems to be the correct solution since there is a clear mutual dependence between them. However, if we look at the results corresponding to the first slices of the temporal series in Fig. 5 (where there is almost no deformation, and even null deformation for the 1st slice), we can see that the quality of the solution is significantly better than in the other cases. In fact, the error is almost 0. It suggests that when the 2D image is not deformed with respect to the input volume, our method is enough to capture slice-to-volume mapping. This limitation is somehow inherent to the model we are using: rigid transformations can not deal with local deformations. To improve the results in these cases, we plan to extend our approach to linear transformations where also anisotropic scaling and shearing can be considered. Following the strategy by [27], it will result straightforward.

Finally, a future line of research has to do with applying discrete rigid (or linear) slice-to-volume registration to other problems. As discussed in Sect. 1.1 motion correction for volume reconstruction is another problem requiring mapping slice-to-volume. It would be interesting to explore how our approaches performs in this case.

References

1. Bao, P., Warmath, J., Galloway, R., Herline, A.: Ultrasound-to-computer-tomography registration for image-guided laparoscopic liver surgery. *Surg. Endosc.* **19**, 424–429 (2005)
2. Birkfellner, W., Figl, M., Kettenbach, J., Hummel, J., Homolka, P., Schernthaner, R., Nau, T., Bergmann, H.: Rigid 2D/3D slice-to-volume registration and its application on fluoroscopic CT images. *Med. Phys.* **34**(1), 246 (2007)
3. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE TPAMI* **26**(9), 1124–1137 (2004)
4. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. *IEEE Trans. Pattern Anal. Mach. Intell.* **23**(11), 1222–1239 (2001)
5. Fei, B., Duerk, J.L., Boll, D.T., Lewin, J.S., Wilson, D.L.: Slice-to-volume registration and its potential application to interventional MRI-guided radio-frequency thermal ablation of prostate cancer. *IEEE Trans. Med. Imaging* **22**(4), 515–525 (2003)
6. Ferrante, E., Fecamp, V., Paragios, N.: Implicit planar and in-plane deformable mapping in medical images through high order graphs. In: *ISBI 2015*, pp. 721–724, April 2015. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7163974>
7. Ferrante, E., Fecamp, V., Paragios, N.: Slice-to-volume deformable registration: efficient one-shot consensus between plane selection and in-plane deformation. *IJCARS* **10**, 791–800 (2015)

8. Ferrante, Enzo, Paragios, Nikos: Non-rigid 2D-3D medical image registration using markov random fields. In: Mori, Kensaku, Sakuma, Ichiro, Sato, Yoshinobu, Barillot, Christian, Navab, Nassir (eds.) MICCAI 2013. LNCS, vol. 8151, pp. 163–170. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-40760-4_21](https://doi.org/10.1007/978-3-642-40760-4_21)
9. Gholipour, A., Estroff, J.A., Warfield, S.K.: Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE TMI* **29**(10), October 2010. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3694441&tool=pmcentrez&rendertype=abstract>, <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5482022>
10. Glocker, B., Sotiras, A.: Deformable medical image registration: setting the state of the art with discrete methods. *Annu. Rev. Biomed. Eng.* **13**, 219–244 (2011)
11. Heinrich, M.P., Jenkinson, M., Brady, M., Schnabel, J.A.: MRF-based deformable registration and ventilation estimation of lung CT. *IEEE TMI* **32**(7), 1239–1248 (2013)
12. Huang, X., Moore, J., Guiraudon, G., Jones, D.L., Bainbridge, D., Ren, J., Peters, T.M.: Dynamic 2D ultrasound and 3D CT image registration of the beating heart. *IEEE TMI* **28**(8), 1179–1189 (2009)
13. Kim, B., Boes, J.L., Bland, P.H., Chenevert, T.L., Meyer, C.R.: Motion correction in fMRI via registration of individual slices into an anatomical volume. *Magn. Reson. Med.* **41**(5), 964–972 (1999)
14. Komodakis, N., Tziritas, G., Paragios, N.: Fast, approximately optimal solutions for single and dynamic MRFs. In: CVPR (2007)
15. Komodakis, N., Tziritas, G., Paragios, N.: Performance vs computational efficiency for optimizing single and dynamic MRFs: setting the state of the art with primal-dual strategies. *Comput. Vis. Image Underst.* **112**(1), 14–29 (2008). <http://linkinghub.elsevier.com/retrieve/pii/S1077314208000982>
16. Lempitsky, V., Roth, S., Rother, C.: FusionFlow: discrete-continuous optimization for optical flow estimation. In: CVPR, pp. 1–22 (2008)
17. Liao, R., Zhang, L., Sun, Y., Miao, S., Chefed’Hotel, C.: A review of recent advances in registration techniques applied to minimally invasive therapy. *IEEE TMM* **15**(5), 983–1000 (2013)
18. Nelder, J.A., Mead, R.: A simplex method for function minimization. *Comput. J.* **7**(4), 308–313 (1965). <http://comjnl.oxfordjournals.org/cgi/doi/10.1093/comjnl/7.4.308>
19. Paragios, N., Ferrante, E., Glocker, B., Komodakis, N., Parisot, S., Zacharaki, E.I.: (Hyper)-graphical models in biomedical image analysis. *Med. Image Anal.* (2016). <http://linkinghub.elsevier.com/retrieve/pii/S1361841516301062>
20. Park, H., Meyer, C.R., Kim, B.: Improved motion correction in fMRI by Joint mapping of slices into an anatomical volume. In: MICCAI, pp. 745–751 (2004)
21. Rousseau, F., Glenn, O.: A novel approach to high resolution fetal brain MR imaging. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, vol. 3749, pp. 548–555 (2005). <http://link.springer.com/10.1007/11566465>
22. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: a survey. *IEEE Trans. Med. Imaging* **32**, 1153–1190 (2013)
23. Tadayyon, H., Lasso, A., Kaushal, A., Guion, P., Fichtinger, G.: Target motion tracking in MRI-guided transrectal robotic prostate biopsy. *IEEE TBE* **58**(11), 3135–3142 (2011)
24. Wang, C., Komodakis, N., Paragios, N.: Markov random field modeling, inference & learning in computer vision & image understanding: a survey. *CVIU* **117**(11), 1610–1627 (2013). <http://linkinghub.elsevier.com/retrieve/pii/S1077314213001343>

25. Xu, R., Athavale, P., Nachman, A., Wright, G.A.: Multiscale registration of real-time and prior MRI data for image-guided cardiac interventions. *IEEE TBE* **61**, 2621–2632 (2014)
26. Zikic, D., Glocker, B., Kutter, O., Groher, M., Komodakis, N., Khamene, A., Paragios, N., Navab, N.: Markov random field optimization for intensity-based 2D–3D registration. In: *SPIE Medical Imaging*, p. 762334. International Society for Optics and Photonics (2010)
27. Zikic, D., Glocker, B., Kutter, O., Groher, M., Komodakis, N., Kamen, A., Paragios, N., Navab, N.: Linear intensity-based image registration by Markov random fields and discrete optimization. *Med. Image Anal.* **14**(4), 550–562 (2010)