Simultaneous Linear and Deformable Registration

Vivien Fécamp, Aristeidis Sotiras, Nikos Paragios

HAL Id: hal-01223991

https://inria.hal.science/hal-01223991

Submitted on 6 Nov 2015

Simultaneous Linear and Deformable Registration

Vivien Fécamp[1,2], Aristeidis Sotiras[3], and Nikos Paragios[1,2]

[1] CVC, Ecole CentraleSupélec, France
[2] GALEN Team, INRIA Saclay `vivien.fecamp,nikos.paragios@inria.fr`
[3] Section of Biomedical Image Analysis University of Pennsylvania
`aristeidis.sotiras@uphs.upenn.edu`

Abstract. In this paper, we present a new approach to tackle simultaneously linear and deformable registration between two images. Our combined formulation avoids the bias created when linear registration is performed independently before a deformable registration. Our registration problem is formulated as a discrete Markov Random Field and a higher order objective function. Usually, a grid is superimposed on the image domain where the latent variables correspond to the local image displacement vectors. Here, we decouple the linear part and the deformable part of the displacement vectors into two conjugate nodes of the grid. We enforce the smoothness of the deformable displacements vectors with binary potentials while the linearity is imposed through third and fourth order potentials. The resulting formulation is modular with respect to the image metric used to evaluate the correctness of mapping as well as with respect to the nature of the linear transformation (rigid, similarity, affine). Inference on this graph is performed efficiently through Alternating Direction Method of Multipliers. Promising results on medical 3D images demonstrate the potentials of our approach.

Keywords: Registration, Markov Random Fields, Higher Order Potentials, Alternating Direction Method of Multipliers

1 Introduction

Linear [5] and deformable registration [4,10] are among the computational pillars of medical image analysis. Linear methods aim to establish approximate correspondences using low rank models (rigid, similarity, affine, etc.). Deformable methods, often employed once global differences have been accounted for, seek an one-to-one mapping between the source and the target image.

Once linear mapping has been addressed, deformable registration is utilized to provide dense correspondences. These methods are in most of the cases image-based and one can refer to an important number of successful developments in the recent years [10]. The Demons algorithm [11] is a computationally efficient solution that iterates between correspondences and smoothing. The ELASTIX method [6] is based on a B-spline deformation model and is a good compromise

between complexity and performance. The DROP algorithm [3] is a modular, metric free, computationally efficient approach to deformable registration.

Linear and deformable registration have been considered separately up to now. Such an approach encompass various limitations with respect to both tasks. Linear transformations estimation can be addressed efficiently if and only if - especially for iconic methods - the largest portion of the image content satisfies the transformation linearity mapping. This is often the case for bones, but never or rarely the case for other organs. The use of robust estimation methods could, up to certain extent, overcome this limitation. In the context of medical imaging, this is not that trivial due to the lack of discriminative visual information. Deformable registration is built upon the result of the linear one. Therefore, one can expect that a strong bias is introduced in the deformable with respect to the choice of the target image. That is particularly disturbing when considered population studies where mapping is done to a common reference space. Furthermore, the performance of deformable registration depends heavily on the input of the linear one and might fail to provide appropriate correspondences. On top of the above concerns, one has to also consider practical issues, like for example the choice of metric spaces, the choice of the optimization process with respect to the two different parts of the deformation process, or the choice of the regularization constraint.

In this paper, our aim is to cope simultaneously with linear and deformable registration. We introduce a novel graphical model that consists of two interconnected components and estimates simultaneously the two registration components by considering a local deformation approach. The first graph consists of a regular grid where higher order constraints between control points impose the expected nature of linearity of the transformation as defined in [2]. The second graph - inspired by the one proposed in [3] - adopts an identical grid endowed with singleton terms that penalize the magnitude of displacements and pairwise constraints imposing deformation smoothness. The two graphs are interconnected and in these links the exact data term is estimated through the composition of the two transformations. The resulting formulation can deal with arbitrary types of linear mapping, arbitrary similarity criteria and various regularization terms. The optimization of such a graphical model is achieved through a master-slave framework that exploits the dual decomposition with the Alternating Direction Method of Multipliers (DD-ADMM) approach [1].

The reminder of the paper is organized as follows: Sect. 2 presents the graphical model formulation for the linear mapping, that is endowed with the deformable component. The optimization of the complete framework and the associated implementation details are presented in Sect. 3 along with experimental validation. The last section concludes the paper and provides future directions.

2 Method

The registration problem consists in finding a transformation T that aligns an image J (typically referred to as source) to a reference image I (typically re-

ferred to as target). A common approach for modeling this problem is by energy minimization:

$$\hat{T} = \arg\min_{T} \xi\left(I, J \circ T\right), \tag{1}$$

where $\hat{T}$ is the optimal transformation and ξ is a similarity measure. Our purpose is to model the image transformation through the displacements of a small set of control points. Therefore, we will formulate (1) as a discrete label assignment problem with the use of Markov Random Fields theory.

$G = (V, E, C)$ a hypergraph, where V denotes the set of nodes, E the set of edges, and C the set of higher order cliques. Let $L = \{l_1, \ldots l_n\}$ be the set of labels that corresponds to a quantized version of the solution space, and l_p denotes the label assigned to node p. The algorithm associates a label l_p to each control point p, in order to align the two images. The energy of the MRF can be written as:

$$E_{MRF} = \sum_{p \in V} U_p(l_p) + \sum_{(p,q) \in E} V_{p,q}(l_p, l_q) + \sum_{c \in C} H_c(l_c) , \tag{2}$$

where $U_p(l_p)$ denotes the *unary potentials*, $V_{p,q}(l_p, l_q)$ denotes the *binary potentials*, $l_c = \{l_p, p \in c\}$, is the set of labels assigned to the nodes in the clique c and $H_c(l_c)$ denotes the *higher order potentials*.

The alignment is measured by a cost function. The choice of the cost function depends on the images; this choice is not imposed in our framework under the mild assumption that it can be evaluated over a patch. The control points correspond to the center of the patches into which the image has been parcellated and are placed according to a grid on the image.

We want to separate and determine simultaneously the linear and deformable transformations. To do that, we decompose the displacement of each control point into a linear and a deformable part. In order to infuse that knowledge into our MRF framework, we will duplicate (see Fig. 1a) the original grid of control points: the first part will encode the linear displacements, the second part will encode the deformable displacements. Therefore, each control point is represented by two corresponding nodes of the graph, one representing its linear displacement, the other one representing its deformable displacement. That duplication allows to keep a reasonable number of nodes and labels, and therefore to greatly decrease the computational cost of performing inference in the graph. Let us note $V1$ the nodes in the first part of the duplicated graph (linear part), and $V2$ those in the second part (deformable part):

$$V = V1 \cup V2. \tag{3}$$

2.1 Graph Construction

Let us now define C, the set of cliques. In our framework, the cliques have very different goals. Those in $V2$ ensure the smoothness of the deformable displacements. So there is an edge between between each pair of neighbour nodes, which

form a grid as used for computing deformable displacement alone as in [9]. Moreover, each node in $V2$ have an unary potential so deforming linearly the image in this part of the graph is penalized. The cliques in $V1$ ensures that the linear displacements of all the points form a coherent linear transformation of the image. What is left to be detailed is the data term. The data term should capture the interactions between pairs of linear and deformable displacements so each pair of duplicated nodes (one in $V1$, one in $V2$) will be linked by an edge.

Unary Potentials. To ensure the algorithm prefers large linear displacements instead of large deformable ones, we employ a unary potential penalizing the norm of the vector of the displacement vector.

$$U_p(l_p) = \|l_p\|. \tag{4}$$

This potential is defined for every node in $V2$, where l_p corresponds to a deformable displacement.

Binary Potentials.

Smoothing Term. A regularization term operating between nodes in $V2$ is necessary in order to ensure the deformable registration is smooth. This can be achieved by penalizing the vector differences between neighboring nodes:

$$V_{p,q}(l_p, l_q) = \frac{\|q - p - l_p + l_q\|}{\|q - p\|} \quad , \tag{5}$$

where p and q represent two neighbour control points, both in $V2$.

Data Term. In order to quantify the alignment of the two images, we employ a patch-based similarity criterion, or we compare a patch from the source image $B_{p,q}$ with a patch in the target domain $B_{l_p+l_q}$ that is chosen based on the composition of the affine and deformable part of the deformation. In order to model the composition of the two parts, it is necessary to use a binary term involving the corresponding nodes p and q in the two parts of the graph. The data term is defined as:

$$V_{p,q} = \rho(B_{p,q} B_{l_p+l_q}) \quad . \tag{6}$$

There is no constraint imposed on the choice of the matching criterion ρ. The proposed model can encompass a wide choice of intensity-based similarity measures, from the sum of absolute difference (SAD) to statistical measures for multimodal registration like mutual information [12].

Higher Order Terms. The higher order potentials are defined as in [2]. Triplets and a special clique called λ-clique ensure the linearity of the transformation. An example of those cliques for a 2D grid for a 2D registration is shown in Fig. 1a.

Fig. 1: a) The new graph in 2D (some triplets are omitted for clarity). Binary are links, triplets are ellipses, the only quadruplet, the λ-clique, is the green shape. b) The decomposition of the original problem in slave subproblems in 2D: one for each line in $V1$, one for each λ-clique, one for each duplicated control point (vertical rectangles here), and one for $V2$.

2.2 Optimization algorithm

To solve the MRF, we use DD-ADMM [1]. Dual Decomposition [7] consists in decomposing a global difficult problem into smaller solvable subproblems (referred to as slaves) and then extracting a solution by cleverly combining the solutions from these subproblems. DD-ADMM is an improvement of Dual Decomposition which accelerates the convergence. In this case, the difficulty of the inference of the optimization displacements lies in the presence of the higher order cliques. Here, the graph is decomposed into trees that constitute the set of subproblems and into a grid. The only requirement for the choice of the subproblems is that they cover (at least once) every node and hyperedge of the hypergraph G.

In our case, in $V1$, a slave problem is defined for each line parallel to a coordinate axis in the first part of the grid, and a slave for each λ-clique. An example of the different slaves in 2D is illustrated in Fig. 1b. A big slave contains all the nodes of the graph but only the edges of $V2$ and the edges encoding the data term. In this big slave, those edges contain one node(in $V1$) link to nothing else. Thus we can send a message from it to the other end of the edge, like in Message-Passing algorithm, to put all needed information into the unary of the node in $V2$. Then we have a simple slave and we optimize it using the Fast-PD algorithm [8]. One additional advantage of this optimization method is the independance of the slave problems, which allows a parallelization of the computation.

3 Experimental Validation

3.1 Implementation Details

The algorithm uses an iterative coarse-to-fine refinement process. The quality of the image is reduced at the first steps to accelerate the computation. The label space is successively refined to explore a large number of displacements

while keeping a reasonable execution time. The label space corresponds to a discretization of potential displacement vectors, regularly distributed on a grid around the 0-displacement vector. The maximal length of the displacement vectors is 0.4 multiplied by the distance between two control points along each axis. The length is iteratively reduced along the iterations. We used up to 7 iterations in our experiments. The successive label space refinement allows to keep the number of labels quite small, 3^3 or 5^3, while reaching sub-millimeter registration accuracy. The grid contains 3^3 control points at the first iterations and is increased to 9^3.

The algorithm is implemented in C++. The tests were performed on a 64 bits machine with a Intel Xeon W3670 processor and 16 Go of RAM. The mean running time for 3D volumes was about 160 seconds when using the similarity criterion SAD.

3.2 Affine transformed images

We use a database of abdomen 3D CT images, containing 6 images of the same patient at different moments. Two organs have been manually segmented by medical doctors, the sigmoid and the bladder. The image dimension is about 512*512*121 with a physical spacing of 0.92*0.92*4 mm, with small variations on the images. We perform several affine transformations of one image. We then applied a small deformation field to the transformed image. This deformation field is small in the sense it should not contain any global linear transformation. We then try to register these deformed images to the original one. Rotations lies between 0 ° and 5 ° and translations reach 20mm. We performed 22 different transformations, with a Sum of Absolute Differences (SAD) similarity measure. We want to compare the affine transformation we find with the one we initially applied. So we fixed 6 points in the images at some extremities of the bodies, and compute the mean distance between the two transformations. Our results show a mean distance of 2.61 mm. Most of the error come from rotations which are not captured by the data term. The results could be improved by using a rotation invariant measure. One example of registration is shown in Fig. 2.

3.3 Real images

We then use intra-patient images from the same database to compare our method with a sequential linear and deformable registration. So images are initially aligned with a linear registration. Then we apply a deformable registration algorithm, DROP [3]. In parallel, we apply our algorithm. We compare the DICE we get from the two methods. Our results show a small improvement (cf. 1) of the DICE.

4   Conclusion

In this paper, we have a discrete MRF formulation to solve the problems of linear and deformable registrations simultaneously, using a local higher order graphical
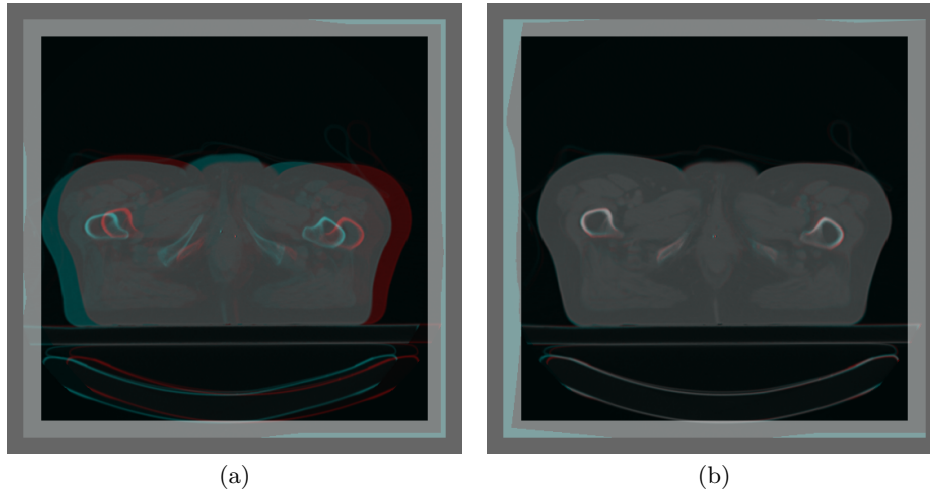
<table>
<tr><td>(a)</td><td>(b)</td></tr>
</table>

Fig. 2: A registration, the two images are superimposed in different colours: (a) Before registration. (b) After registration.

DICE	Bladder	Sigmoid
Before registration	45.61	39.383
Consecutive registration	78.15	68.55
Our registration	78.47	68.64

Table 1: Results of the DICE of two organs while comparing our simultaneous registration with a consecutive one.

model acting with hidden variables being the displacement vectors as labels. The proposed formulation is metric-free (can deal with arbitrary similarity criterion), modular with respect to the nature of the linear transformation (rigid, similarity, affine and could be extended to projective) and exhibits computational efficiency due to its relative local nature and the designed search space. We use a multi-level strategy, in a coarse-to-fine manner. The performance of the method on 3D multi-modal medical data along with comparisons with state of the art methods demonstrate its potential for applications. Opposed to the usual sequential linear/deformable registration, our scheme is based on a sound mathematical framework, so it gives a better accuracy to both the linear transformation and the deformable field. Moreover this approach is fast compared to state of the art methods.

Furthermore, mapping from 2D to 3D is a great problem of phenomenal interest either in vision or in medical imaging towards image-based navigation/guidance. The same concept as the one considered when decomposing 2D-2D or 3D-3D deformations in linear and non-linear components can be also

applied to 2D-3D. The clinical impact of such a component in computer assisted surgery is currently under investigation.

References

1. Pedro Aguiar, Eric P Xing, Mário Figueiredo, Noah A Smith, and André Martins. An augmented lagrangian approach to constrained map inference. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 169–176, 2011. 2, 5

2. B. Anonymous, Anonymous C., and Anonymous D. Modular linear iconic matching using higher order graphs. *ISBI, IEEE Transactions on*, 2015. 2, 4

3. Ben Glocker, Nikos Komodakis, Georgios Tziritas, Nassir Navab, and Nikos Paragios. Dense image registration through MRFs and efficient linear programming. *Medical image analysis*, 12(6):731–41, dec 2008. 2, 6

4. Ben Glocker, Aristeidis Sotiras, Nikos Komodakis, and Nikos Paragios. Deformable medical image registration: setting the state of the art with discrete methods. *Annual review of biomedical engineering*, 13:219–244, August 2011. 1

5. Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith. Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *NeuroImage*, 17(2):825–841, oct 2002. 1

6. S. Klein, M. Staring, K. Murphy, M.A. Viergever, and J. Pluim. elastix: A toolbox for intensity-based medical image registration. *Medical Imaging, IEEE Transactions on*, 29(1):196 –205, jan. 2010. 1

7. Nikos Komodakis, Nikos Paragios, and Georgios Tziritas. MRF energy minimization and beyond via dual decomposition. *PAMI*, 33(3):531–52, March 2011. 5

8. Nikos Komodakis, Georgios Tziritas, and Nikos Paragios. Fast, Approximately Optimal Solutions for Single and Dynamic MRFs. *IEEE Conference on Computer Vision and Pattern Recognition*, D(2):1–8, 2007. 5

9. D Rueckert, L I Sonoda, C Hayes, D L Hill, M O Leach, and D J Hawkes. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE transactions on medical imaging*, 18(8):712–21, August 1999. 4

10. Aristeidis Sotiras, Davatzikos Christos, and Nikos Paragios. Deformable Medical Image Registration: A Survey. Research Report RR-7919, INRIA, Sep 2012. 1

11. Tom Vercauteren, Xavier Pennec, Aymeric Perchant, and Nicholas Ayache. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*, 45(1, Supplement 1):S61 – S72, 2009. 1

12. Paul Viola and William M. Wells. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24(2):137–154, 1997. 4