A MULTI-PATH APPROACH TO HISTOLOGY VOLUME RECONSTRUCTION

Jonas Pichat[1]   *Marc Modat*[1,2]   *Tarek Yousry*[3]   *Sebastien Ourselin*[1,2]

[1] Translational Imaging Group, Centre for Medical Image Computing, University College London, UK
[2] Dementia Research Center, UCL Institute of Neurology, London, UK
[3] Department of Brain Repair and Rehabilitation, UCL Institute of Neurology, London, UK

ABSTRACT

This paper presents a method for correcting erratic pairwise registrations when reconstructing a volume from 2D histology slices. Due to complex and unpredictable alterations of the content of histology images, a pairwise rigid registration between two adjacent slices may fail systematically. Conversely, a neighbouring registration, which potentially involves one of these two slices, will work. This grounds our approach: using correct spatial correspondences established through neighbouring registrations to account for direct failures. We propose to search the best alignment of every couple of adjacent slices from a finite set of transformations that involve neighbouring slices in a transitive fashion. Using the proposed method, we obtained reconstructed volumes with increased coherence compared to the classical pairwise approach, both in synthetic and real data.

Index Terms— Rigid registration, histology, volume reconstruction

1. INTRODUCTION

The potential of histopathology image analysis is well established and embodies a two-way benefit: not only is it a gold standard in disease diagnosis but it also helps understanding the inherent reasons for making a diagnosis [1]. One goal logically derived from this reality is thus to be able to explore how the findings about a disease in histology slices correlate to the information provided by standard clinical imaging routines. This would improve the disease understanding by combining high resolution anatomical information with 3D, non-invasive though lower resolution information (e.g. using magnetic resonance (MR) imaging). This is known as data fusion, where the information gained from two modalities is usually of complementary nature. This motivates the building of 3D histological atlases as they provide *a priori* knowledge when
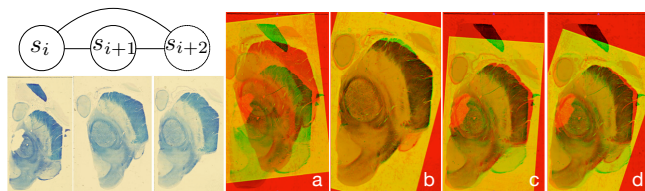
Fig. 1. Illustration of the pairwise rigid registration specificity among three consecutive histology slices s_i, s_{i+1} and s_{i+2}. The registration of slices s_{i+1} and s_i fails (a), whereas the two neighbouring registrations of slices s_{i+1} and s_{i+2}, and of slices s_{i+2} and s_i work (b)-(c). The right-most panel shows that composing two intermediate good transformations $T_{s_{i+2} \leftarrow s_{i+1}}$ and $T_{s_i \leftarrow s_{i+2}}$ results in a better alignment of slices s_{i+1} and s_i (d).

conformed to living patients MR. Atlases permit the recognition of complex anatomical structures and the identification of their corresponding signals in other modalities when the image content is not described explicitly enough. Consequently, they become very helpful in surgical planning [2] where some interventions impose high demands on intraoperative accuracy. This integration process requires bringing the modalities involved into spatial alignment. The initial concern is to globally reconstruct the histology volume as good as possible so that subsequent local adjustments are effective. Several issues arise when dealing with histology, the most obvious being that the volume information is lost as soon as the histological process is engaged. At a finer scale, tissues are independently altered due to preparation: the cutting process induces geometrical distortions, tearing and missing parts, various shapes and locations on the slide, while the staining causes artefacts and intensity inhomogeneities. In addition, the goal being to align a sequence of images of similar rather than of same objects, errors propagate and accumulate along the stack [3]. Early works tackled the tissue degradation issue by searching for the rigid transformation that locally maps the largest number of similar regions [4]. In order to improve the coherence of the reconstructed volume, attention has been directed toward neighbouring slices [5], and toward defining boundary smoothness as a gauge of quality [6]. More recent works

showed the importance of the reference slice selection [7], which as a matter of fact also showed that the generated volume depends directly on the choice of that reference slice.

Unlike previous approaches, the proposed method relaxes the constraint of choosing a reference slice. As in [5], we formulate the reconstruction problem using graphs, though we make sure that the set of weights—the similarity measures—applied to the graph edges is totally ordered: it is a partially ordered set that has the property of comparability. The key idea is based upon the simple observation described in Fig.1, where the transitive relation on the set N of n slices (or nodes) s is defined as: $\forall\ s_i, s_{i+1}, s_l \in N$ with $l \neq \{i, i+1\}: (s_{i+1}Rs_l \wedge s_lRs_i) \Rightarrow s_{i+1}Rs_i$, where s_l is an intermediate neighbouring slice, the binary relation R is equivalent to a rigid registration in our case and $\wedge$ is the logical conjunction. This relation can be extended to any number of intermediate slices. In practice, this means that the transformation between two adjacent slices can be expressed 1) directly, and 2) indirectly as the composition of two (or more) transformations using neighbouring slices. For every pair of adjacent slices (s_i, s_{i+1}), we thus aim at finding the transformation that minimises the dissimilarity, i.e. the cost of the link. The process is iterated until convergence to guarantee the recovery of a reconstructed histological volume with continuous structures.

2. METHODS

2.1. Preprocessing

Intensity standardisation. Inconsistencies in the staining process may exhibit large differences in the appearance of slices. In order to standardise their texture, we use cumulative distribution function (CDF) matching, for its high computational speed and the advantageous property that a CDF is a monotonically non-decreasing function. Every value from the target distribution is mapped to its nearest value in the source distribution. As for the registration problem, a reference histogram is required, to which every other histogram is matched, or as in [8] the normalisation between each pair of adjacent slices is propagated along the stack, which goes back to implicitly taking the first histogram as reference. This means the intensity-compensated stack changes with the choice of the reference slice. To bypass the non-uniqueness of the solution, we chose a groupwise approach. The CDFs c of the images of slices are formed from the normalised histograms, and the mean CDF $\widetilde{c}$ is computed using least trimmed squares (LTS), which consists of minimising: $\min_{\widetilde{c}} \sum_{m=1}^{h} (r^2)_{m:n}$ where $(r^2)_{1:n} \leq \ldots \leq (r^2)_{n:n}$ are the ordered squared residuals, $r_m = \sum(c_m - \widetilde{c})$. LTS regression consists of finding the subset of $h < n$ CDFs whose least squares produces the smallest sum of squared residuals. We set the number of inliers to $h = \lfloor n/2 \rfloor$, assuming that the outliers constitute less than 50% of the CDFs. This way, we
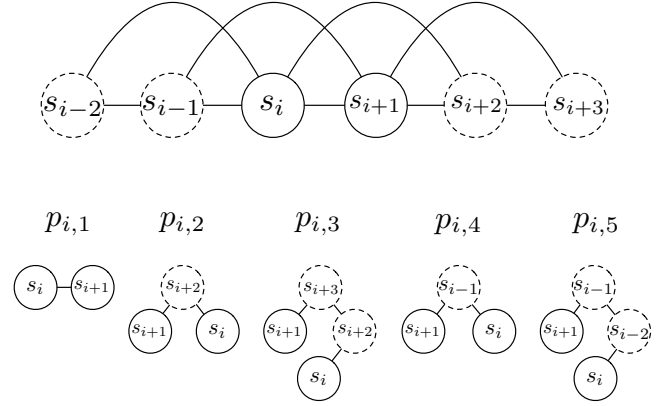

Fig. 2. Subgraph (top) and paths (bottom) associated with a pair of adjacent slices (s_i, s_{i+1}) ($\eta = 2$ and $\epsilon = 1$). Once the path $p_{i,k}$ with the minimum cost has been identified, the arc (s_i, s_{i+1}) is updated.

minimise the bias from these outliers. Every CDF is then matched to the estimated mean CDF, $\widetilde{c}$.

Masking. Every slice has its tissue masked in order to drive the registration process using relevant information only. This is solved as a two-class problem, where the large contribution of the background intensities can be identified in the matched CDFs, c_i' (i.e. largest slope). The threshold is therefore set as the bin where the derivative of c_i' is maximum. Erosion is then applied to the binary image so that the edges of the tissue block are not taken into account. This, in addition to the fact that largest connected component is kept as the final mask, is to prevent *artificial* edges (e.g. due to tearing) from impacting the registration.

2.2. Multi-path optimisation

The method finds the shortest paths in a dynamic graph. By dynamic graph, we denote the graph $G = (N, A)$ subject to a sequence of updates. A is the set of arcs (i.e. a rigid registration between two connected nodes is performed) and it depends on two parameters, η and ϵ, being respectively: the number of neighbouring slices taken into account in the shortest path search, and the maximum gap tolerated between two slices to be registered. The cost function to optimise is expressed as: $\sum_{1 \leq i \leq n-1} \min_{T_{s_i \leftarrow s_{i+1}}} d(s_i, s_{i+1} \circ T_{s_i \leftarrow s_{i+1}})$, where d is the dissimilarity and $T_{s_i \leftarrow s_{i+1}}$, an element of the special Euclidean group SE(2), is the rigid transformation that minimises d. G is static between two updates and we aim at shortening the arcs (s_i, s_{i+1}).

2.2.1. Generating the graph G

For a given pair (s_i, s_{i+1}), let K be the total number of paths that connect, directly and indirectly, s_{i+1} to s_i. A path $k \leq$

K, $p_{i,k}$, exists if it meets three conditions: 1) $p_{i,k}$ is a tree, 2) every arc of $p_{i,k}$ involves only nodes which indexes are within the range $[i-\eta; i+1+\eta]$, 3) if j is the index of the current node, a successor is within the range $j \pm (\epsilon+1)$. This inspection is done using a breadth-first search. Thus, the arc set of G, $A(G)$ is the union of the arc sets of all the subgraphs associated with pairs (s_i, s_{i+1}) and paths $p_{i,k}$ (Fig.2). The transformation $T_{s_i \leftarrow s_{i+1}}^{(k)}$ associated with the pair (s_i, s_{i+1}) and path $p_{i,k}$ is defined as the composition of the transformations of every arc in the arc set of $p_{i,k}$.

2.2.2. Registrations

The rigid registrations[1] are done using the block-matching strategy described in [9]. Prior to any registration, we use the masks obtained in section 2.1. and align the centres of mass of the two binary images. We thus ensure a better search neighbourhood. The transformations are symmetric [10], making the graph G undirected. The computation of the transformations is done once for all before the shortest path search, and the list is given by the arc set $A(G)$.

2.2.3. Updates and convergence

G is broken up into smaller overlapping subgraphs. One iteration is complete when all the arcs (s_i, s_{i+1}) of G have been shortened. For every subgraph associated with a pair (s_i, s_{i+1}), the set of dissimilarities $D_i = \{d(s_i, s_{i+1} \circ T_{s_i \leftarrow s_{i+1}}^{(k)} | 1 \leq k \leq K\}$, where $k = 1$ refers to the path connecting directly s_{i+1} to s_i, is computed and totally ordered. The arc (s_i, s_{i+1}) at the next iteration is associated with the new transformation: $T_{s_i \leftarrow s_{i+1}}^{(1)} \leftarrow \arg\min_{T_{s_i \leftarrow s_{i+1}}} D_i$. The convergence is reached when all the $p_{i,1}$ (direct links between every s_{i+1} and s_i) are the shortest links. We use the normalised mutual information as a similarity measure. An extra step consists of doing another pass through the method described, using the sequence of optimised pairwise transformations as an initialisation of the registrations. The multi-scale scheme in the block-matching strategy allows to get out of a local maxima (if still stuck despite the first pass) or to increase the smoothness of the reconstructed volume. The volume is finally reconstructed by applying the resulting sequence of transformations to the corresponding slices. The choice of the reference slice only impacts the resolution of the space in which every image is resampled.

3. VALIDATION

In order to assess the efficiency of the proposed method, two MR volumes were sampled, thereby creating a sequence of 2D slices perfectly aligned: MR1, with 26 slices and MR2 with 29 slices. This volume was used as our ground truth.
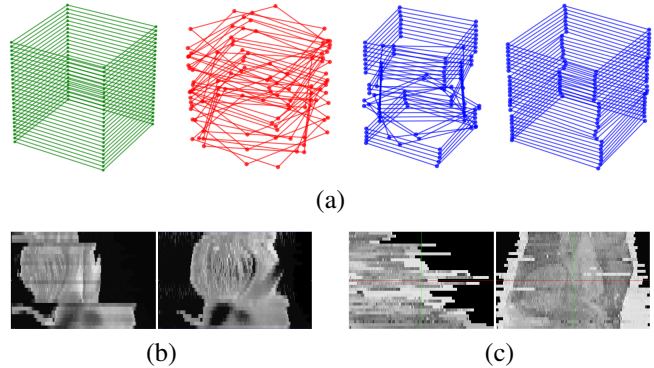
(a)

(b)   (c)

Fig. 3. Validation and results. (a) Using four landmarks (the four corners of a 2D MR image), the original MR volume is shown in green, the rigidly corrupted volume, in red, and the reconstructed volumes using pairwise transformations and our multi-path approach, in blue (respectively on the left and on the right). (b) shows the reconstruction of the MR volumes using pairwise transformations (left) and our multi-path approach (right). (c) shows the histology volume reconstructions in the same cases as in (b).

The resolution of 2D slices was 512×512. Every 2D slice was corrupted with a random rigid transformation and adding Gaussian noise ($\sigma_{noise} = 0.05$), although most details were already lost by nature of the image itself. The rotation (around

Table 1. Target registration error (TRE) with respect to four landmarks (image corners) between original and reconstructed MR volumes after pairwise (pw) and multi-path (mp) approach for 2 datasets (in voxels).

	mean TRE $\pm \sigma$		max TRE	
	pw	mp	pw	mp
MR1	81.92 ± 37.86	20.48 ± 5.47	152.07	56.32
MR2	84.48 ± 42.94	53.76 ± 28.74	148.69	76.8

the stacking dimension) and translation (in-plane) parameters were respectively limited to the ranges $[-25, 25]$ (in degrees) and $[-50; 50]$ (in voxels). However, the translation parameter is secondary here, as the centres of mass of tissues to be registered are aligned by default. The resulting sequence of slices was used as an input in the described method and the target registration error (TRE) was measured with respect to the four corners of the images, in order to check the accuracy of the reconstruction. As expected, the method allowed to reduce the mean TRE as well as correcting erratic pairwise registrations (Table 1). TREs need to be considered as relative rather than absolute measures because they depend on where the landmarks have been taken. MR1 exhibited several major failures (see in Fig.3 the blue volume on the left (a) and the left panel in (b)) after pairwise registrations, while MR2

was unique in that the pairwise reconstruction was inducing a spiral movement of the stack. By taking into account registrations involving pairs of non-adjacent slices, the latter effect was constrained and reduced. The small number of slices made us choose $\eta = 2$ and $\epsilon = 1$.

The method was also applied to three histology datasets of around 30 slices each. Similar conclusions were drawn (Fig.3c). We chose $\eta = 2$ and $\epsilon = 1$, and in the registration step: percentage of blocks used in the optimisation scheme: 70%; number of levels used to generate the pyramids in the coarse-to-fine approach: 5; maximum number of iterations in the LTS to perform per level: 50; standard deviation (in voxels) of the Gaussian kernel used to smooth the pair of images to be registered: 5 (3 in the second pass). The multi-path approach took between 4 and 8 iterations to converge (including the second pass).

4. DISCUSSION

Intensity standardisation. The choice made here for intensity standardisation was only intended to ease the registration process. It has no biological ground: a quantitative analysis of the reconstructed histology volume may require more advanced normalisations that take into account the variations in absorbance of staining compounds, and address the problem of retrieving the biological component of a pixel value [11]. Another interesting track would be to use color decomposition as pre-processing for separating stained structures and drive alignments using specific tissue types.

Registrations. As in [12], we believe that it is better to preserve the shape of the tissue at this step rather than arbitrarily and possibly wrongly compensating for distortions. Thus, the number of degrees of freedom was restricted to 3, as no external information about the true shape was used here. The number of rigid registrations is $\sum_{j=0}^{\epsilon} N - (j+1)$, making the method have a linear time complexity (although all the registrations can be performed in parallel). The previous number does not depend upon η because redundant registrations are discarded and the transformations are symmetric (only the inverse of a computed transformation is taken when needed).

Tissue degradations. Our framework is proposing alternative ways to reach alignment despite various degradations, rather than addressing any of them specifically (Fig.1). This means involving slices with slightly different features, positions and orientations that are potentially more conducive to reaching global optimum. Specifically, though, our method showed good results for tissues torn into several pieces when the largest connected component was used to drive the registration of the entire image. One may also want to consider each connected component as an independent tissue element, and then fuse them into a unique image after separate alignments. Future works may involve more complex image correction techniques able to differentiate between cases where tissue is missing or just torn open.

5. REFERENCES

[1] M. N. Gurcan, L. E. Boucheron, A. Can, A. Madabhushi, N. M. Rajpoot, and B. Yener, "Histopathological image analysis: A review," *Biomedical Engineering, IEEE Reviews in*, vol. 2, pp. 147–171, 2009.

[2] M. M. Chakravarty, G. Bertrand, C. P. Hodge, A. F. Sadikot, and D. L. Collins, "The creation of a brain atlas for image guided neurosurgery using serial histological data," *Neuroimage*, vol. 30, no. 2, pp. 359–376, 2006.

[3] G. Malandain, E. Bardinet, K. Nelissen, and W. Vanduffel, "Fusion of autoradiographs with an MR volume using 2-D and 3-D linear transformations," *NeuroImage*, vol. 23, no. 1, pp. 111–127, 2004.

[4] S. Ourselin, A. Roche, G. Subsol, X. Pennec, and N. Ayache, "Reconstructing a 3D structure from serial histological sections," *Image and vision computing*, vol. 19, no. 1, pp. 25–31, 2001.

[5] P. Yushkevich, B. Avants, L. Ng, M. Hawrylycz, P. Burstein, H. Zhang, and J. Gee, "3D Mouse Brain Reconstruction from Histology Using a Coarse-to-Fine Approach," in *Biomedical Image Registration*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2006, vol. 4057, pp. 230–237.

[6] A. Cifor, L. Bai, and A. Pitiot, "Smoothness-guided 3-D reconstruction of 2-D histological images," *Neuroimage*, vol. 56, no. 1, pp. 197–211, 2011.

[7] U. Bagci and L. Bai, "Automatic best reference slice selection for smooth volume reconstruction of a mouse brain from histological images," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 9, pp. 1688–1696, 2010.

[8] G. Malandain and E. Bardinet, "Intensity compensation within series of images," in *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2003*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2003, vol. 2879, pp. 41–49.

[9] S. Ourselin, A. Roche, S. Prima, and N. Ayache, "Block Matching: A General Framework to Improve Robustness of Rigid Registration of Medical Images," in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2000*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2000, vol. 1935, pp. 557–566.

[10] M. Modat, D. M. Cash, P. Daga, G. P. Winston, J. S. Duncan, and S. Ourselin, "Global image registration using a symmetric block-matching approach," *Journal of Medical Imaging*, vol. 1, no. 2, pp. 024 003–024 003, 2014.

[11] M. Macenko, M. Niethammer, J. Marron, D. Borland, J. T. Woosley, X. Guan, C. Schmitt, and N. E. Thomas, "A method for normalizing histology slides for quantitative analysis." in *ISBI*, vol. 9, 2009, pp. 1107–1110.

[12] J. S. Lee, S.-H. Ahn, D. S. Lee, S. H. Oh, C. S. Kim, J. M. Jeong, K. S. Park, J.-K. Chung, and M. C. Lee, "Voxel-based statistical analysis of cerebral glucose metabolism in the rat cortical deafness model by 3D reconstruction of brain from autoradiographic images," *European journal of nuclear medicine and molecular imaging*, vol. 32, no. 6, pp. 696–701, 2005.