Concurrent Tumor Segmentation and Registration with Uncertainty-based Sparse non-Uniform Graphs

Sarah Parisot[d,*], William Wells III[e,f], Stéphane Chemouny[c], Hugues Duffau[g], Nikos Paragios[a,b]

[a] *Center for Visual Computing, Ecole Centrale Paris, Châtenay Malabry, France*
[b] *Equipe GALEN, INRIA Saclay - Ile de France, Orsay, France*
[c] *Intrasense SAS, Montpellier, France*
[d] *Biomedical Image Analysis Group, Department of Computing, Imperial College London, UK*
[e] *Surgical Planning Laboratory, Brigham and Women's hospital, Harvard Medical School, Boston MA, USA*
[f] *Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Boston MA, USA*
[g] *Department of Neurosurgery, Hopital Gui de Chauliac, Montpellier, France*

Abstract

In this paper, we present a graph-based concurrent brain tumor segmentation and atlas to diseased patient registration framework. Both segmentation and registration problems are modeled using a unified pairwise discrete Markov Random Field model on a sparse grid superimposed to the image domain. Segmentation is addressed based on pattern classification techniques, while registration is performed by maximizing the similarity between volumes and is modular with respect to the matching criterion. The two problems are coupled by relaxing the registration term in the tumor area, corresponding to areas of high classification score and high dissimilarity between volumes. In order to overcome the main shortcomings of discrete approaches regarding appropriate sampling of the solution space as well as important memory requirements, content driven samplings of the discrete displacement set and the sparse grid are considered, based on the local segmentation and registration uncertainties recovered by the min marginal energies. State of the art results on a substantial low-grade glioma database demonstrate the potential of our method, while our proposed approach shows maintained performance and strongly reduced complexity of the model.

Keywords: Concurrent Segmentation/Registration, Markov Random Fields, Min-Marginals, Brain Tumors

1. Introduction

Gliomas are the most common type of primary brain tumors and arise from glial cells. They are classified in 4 grades by the World Health Organization (WHO), grade I corresponding to benign tumors with excellent prognosis and Grade IV gliomas (Glioblastoma Multiforme) being the most common and lethal. WHO grade II Low Grade Gliomas (LGG) are a specific kind of glioma that represent about 30% of the brain tumors and can affect younger patients (Soffietti et al., 2010). They are characterized by a continuous slow growth and yield mild symptoms. They generally undergo anaplastic transformation into a fast growing malignant tumors and therefore have to be monitored closely via frequent MRIs. Knowing the size and extent of a brain tumor is of extreme importance in order to evaluate its growth, its reaction to therapy and for surgery planning. Currently, the physicists compute the main tumor diameters and approximate it as an ellipsoid, a highly imprecise measure that tends to overestimate the volume of the tumor (Pallud et al., 2012). The current gold standard is manual segmentation, which on top of being a tedious and time consuming task, is also subject to a high inter and intra operator variability. Automatic tumor segmentation is thus an active research field that aims at obtaining fast and robust segmentations. It is a particularly difficult subject due to the extreme heterogeneity between the tumors in appearance, shapes and size and their overlapping intensities with the healthy tissue. LGG are diffusively infiltrative tumors with extremely irregular and fuzzy boundaries, rendering the segmentation task even more difficult.

Fuzzy clustering and knowledge based methods were amongst the first considered for tumor segmentation with limited success (Clark et al., 1998; Fletcher-Heath et al., 2001). Level sets and Active Contours have been a popular approach (Ho et al., 2002; Cobzas et al., 2007; Taheri et al., 2010), but suffer from their strong sensitivity to initialization. The idea is to model the tumor boundary as a parametric curve that evolves depending on the image properties and curvature constraints. Statistical classification methods offer an efficient way of detecting tumor voxels. The voxels are treated independently and separated by a classifier that is learned from a set of training samples. Examples refer to the Support Vector Machines (SVM) (Verma et al., 2008; Zhang et al., 2004; García and Moreno, 2004), Boosting (Xuan and Liao, 2007) or the Decision Forests (Zikic et al., 2012). Despite promising performance, those methods are plagued by the i.i.d assumption that treats each voxel independently, leading to irregular segmentations. Morphological filtering (Zhang et al., 2004) or neighborhood dependent

*Corresponding author: Biomedical Image Analysis Group, Department of Computing, Imperial College London, UK. Tel: +44 (0)20 7594 8374. E-mail address: s.parisot@imperial.ac.uk. This work was carried out during the author's affiliation with Ecole Centrale Paris, INRIA Saclay - Ile de France and Intrasense from 2010 to 2013.

features (Zikic et al., 2012) offer limited improvement on the local consistency of the segmentation. Notable improvement is observed when coupling the statistical classification with local neighborhood dependencies (Lee et al., 2008; Görlitz et al., 2007; Wels et al., 2008; Bauer et al., 2011), modeled by a random field (Markov Random Field (MRF), Conditional Random Field (CRF)) (Wang et al., 2013) based spatial prior. In this context, the segmentation is locally smoothed by penalizing neighbors that are assigned different segmentation labels, but still lacks global information regarding the tumor's position and the brain boundaries. Stronger dependencies can be modeled via a hierarchical approach. Gering et al. (2002) proposed a multi layer MRF approach where the tumor is detected as an outlier from manually selected training voxels. At each layer, the segmentation is refined based on higher level information and the previous layer's segmentation. Corso et al. (2008) combine Bayesian classification using Gaussian Mixture Models with a hierarchical graph affinity model, where the spatial dependencies are modeled by assigning an affinity to each graph edge.

Atlas-based segmentation methods rely on the registration of an annotated volume to the subject in order to segment the structures of interest. The use of a brain atlas allows for structural spatial prior information, but the task is more difficult when the structure to segment is a tumor since it cannot be matched in the atlas. That is often addressed through a model for tumor detection. Kaus et al. (2001) alternate kNN classification based on intensity and anatomical location with a registration step based on the structures' segmentation, the tumor being labeled as white matter in the registration process. In (Prastawa et al., 2003) a probabilistic atlas is affinely registered to the patient, enabling to define prior probabilities on the expected intensities of the structures. The atlas is modified to account for tumor presence (detected by contrast enhancement) and edema. Similarly to Gering et al. (2002), the tumor voxels can be detected as outliers from the healthy voxels (Menze et al., 2010; Prastawa et al., 2004). The healthy structures' features are estimated from a registered healthy atlas. Additional local spatial constraints are modeled via Markov Random Fields (Menze et al., 2010) or level sets (Prastawa et al., 2004).

Atlas based methods depend on the quality of the registration. Rigid or affine registration methods are not sufficient to recover the inter patient anatomical differences, while traditional non-rigid registration methods fail in this context by attempting to find correspondences between the tumor and the healthy voxels. An efficient atlas based segmentation thus requires a registration scheme that accommodates for the presence of the tumor.

Despite extensive work in deformable image registration (Zikic et al., 2010; Ou et al., 2011; Berendsen et al., 2013; Sotiras et al., 2013), there has been limited work dedicated to registration with missing correspondences. Such a registration task is of high interest for the study of brain tumors through statistical atlases and longitudinal studies. A tumor specific probabilistic atlas, constructed through affine registration of a large database to the same reference coordinates, was notably proposed in (Parisot et al., 2011). It enabled the identification of preferential locations for the tumors and could lead to unraveling position dependent behaviors and origins. Deformable registration would enable to go further and study the interactions between the tumors and the brain structures and functional areas. Understanding the tumors growth patterns and their impact on the brain's functional organization is of key importance for therapy and surgery planning.

We can distinguish two groups of methods for registration in the presence of a tumor. The first relies on modeling the tumor growth to evaluate the tumor induced deformation (Kyriacou et al., 1999; Mohamed et al., 2006; Zacharaki et al., 2008; Cuadra et al., 2004). Kyriacou et al. (1999) proposed a biomechanical finite element model to simulate the tumor induced deformation while assuming a radial uniform growth of the tumor. Using the tumor growth model, a healthy brain was simulated by contracting the tumor, allowing for a normal registration process. Cuadra et al. (2004) also assumed radial growth of the tumor. The registration is performed using the demons algorithm (Thirion, 1998) between healthy voxels and is based on the distance from a manually selected seed in the tumor area (that has been segmented prior to the registration process). Mohamed et al. (2006) decomposed the deformation as inter subject and tumor induced deformations. The latter was modeled via a biomechanical finite element model whose parameters are learned by statistical learning. The tumor growth is then simulated in the healthy atlas, enabling normal registration. This method was extended in (Zacharaki et al., 2008) towards a computationally efficient biomechanical model taking into account the potential infiltrative parts of the tumor by limiting the tumor growth. Growth models require either user interaction or extensive computations to evaluate the model parameters and are mostly adapted to space occupying lesions. Low grade gliomas are infiltrative tumors with little to no mass effect and edemas. The limited amount of deformation caused by the tumors renders the use of growth model not adapted and possible prone to errors assuming the tumor pushes tissue instead of infiltrating it. The second group of methods (Brett et al., 2001; Stefanescu et al., 2004) adopts a simpler approach and masks the pathology towards excluding it during registration. The tumor area is discarded during the computation of the similarity criterion and deformed by interpolation. This kind of approach offers a better modularity with respect to the pathology since no assumption is made about the pathological area nor the progression of the tumor. Both approaches require a reliable segmentation of the tumor, making the registration dependent on the quality of the segmentation of the tumor.

Registration and tumor segmentation appear as two fundamentally correlated problems, where one could benefit from the other if performed simultaneously. The idea of coupling segmentation and registration is not a new concept. Yezzi et al. (2003) used an active contour framework, estimating the registration parameters and reference volume's segmentation curve by minimizing a joint energy depending on both images. The floating image is segmented by registering the reference's segmentation. A maximum a posteriori framework was presented in (Wyatt and Noble, 2003) where the segmentation and rigid registration parameters are determined alternatively. The segmentation relies on Gaussian Mixture Models coupled with an

MRF prior, while the registration relies on the segmentation by minimizing the joint class histogram between both images. Mahapatra and Sun (2012) proposed an MRF based framework where each voxel of the image has to be assigned a displacement and segmentation label. The different classes are separated based on the intensities in both images while the registration relies on minimizing conventional similarity metrics. The registration and segmentation fields are smoothed by enforcing similar displacement among voxels of the same class. Ashburner and Friston (2005) proposed a statistical model, where a probabilistic atlas plays the part of a spatial prior for segmentation and bias field correction. The different classes are separated via a mixture of Gaussians, allowing for several modes per class. The atlas is globally registered by affine registration then locally deformed. Last but not least, Pohl et al. (2006) developed an Expectation Maximization (EM) Bayesian framework, alternatively estimating the segmentation probabilities and the rigid registration and bias field parameters.

All those methods rely on the concept that the structures to be segmented appear in both images. The joint segmentation and registration problem becomes far more challenging in the presence of a pathology due to the absence of a match in the second image. Most methods alternatively estimate the registration and segmentation maps. Chitphakdithai and Duncan (2010) proposed an EM Bayesian framework in the context of a surgical tumor resection. The resection area was detected by statistical learning on a training set based on the intensity values and deformed by interpolation (constant registration cost in the resection area). In the same clinical context, Risholm et al. (2009) coupled the demons algorithm with level sets. They alternate segmentation of the resected area by evolving a level set based on the image gradient and intensities disagreements, with a demons based registration that accommodates the resection by only allowing displacement towards the area. The problem is more challenging in the context of tumors that have complex intensity profiles. Gooya et al. (2011), inspired from the work of Zacharaki et al. (2008) and Pohl et al. (2006) introduced a method to to deal with the presence of a tumor. The tumor is simulated in a probabilistic atlas via a biomechanical model of tumor growth. The EM algorithm is used to iteratively estimate segmentation posterior probabilities and the tumor growth and registration parameters. While growth models are able to simulate the mass effect, they suffer from the computational burden of estimating the model parameters and are hardly generalizable to other pathologies. Furthermore, the quality of the registration directly depends on the quality of the model which implies extended knowledge on the pathology.

In (Parisot et al., 2012), we introduced a concurrent segmentation and registration framework that exploits the dependencies between the two problems in order to adapt the registration task to the presence of the tumor as well as increase the segmentation quality. The concurrent registration and segmentation framework is embedded in a discrete graphical model, where a sparse grid is superimposed to the volume domain and each node will be simultaneously displaced and classified. The registration term is relaxed in the tumor area that is detected by statistical classification. Pairwise constraints ensure the smoothness of the segmentation and deformation fields. This discrete approach raises the problem of defining the discrete displacement set and resolution of the sparse grid that have to be high enough to capture small details and remain computationally efficient. In this paper, we extend the proposed method through a novel content-driven hierarchical coarse to fine approach exploring segmentation and registration uncertainties as determined by the min-marginal energies. The displacement set sampling relies on the local structures anisotropy while the grid refinement is controlled by the local homogeneity of the region and the segmentation uncertainties. This yields non uniform high resolution grids with a much lower complexity. The proposed MRF based individual tumor detection and registration framework and their coupling is described in Section 2 while the uncertainty driven adaptive sampling method is introduced in Section 3. The experimental validation is part of Section 4 and is carried out on a large low-grade glioma database as well as the publicly available BRATS dataset. Discussion and future directions conclude the paper.

2. Concurrent Tumor Segmentation and Registration

2.1. Statistical Classification based Tumor Segmentation

Let us consider a volume V featuring a tumor that we seek to segment. The tumor can be efficiently detected via the construction of a classifier separating tumor voxels from healthy voxels. We adopt the Gentle Adaboost algorithm (Friedman et al., 2000) that builds a strong classifier as a linear combination of weak classifiers. Let us consider a set of N training samples $\{\mathbf{x_i}, y_i\}, i \in \{1, N\}$, where $\mathbf{x_i}$ is a voxel extracted from a tumor bearing volume, and y_i is its corresponding label (tumor or background). To each pair is associated a feature vector $\Pi(\mathbf{x_i})$ and a weight $W_i = \frac{1}{N}$.

At each iteration t, the algorithm selects a feature and a threshold, in order to build a weak classifier $h_t(\mathbf{x_i})$ as a decision stump that minimizes the classification error:

$$err = \sum_{i=1}^{N} W_i |h_t(\mathbf{x_i}) - y_i| \tag{1}$$

The weights W_i are then updated as $W_i = W_i \exp(-y_i h_t(\mathbf{x_i}))$ in order to give more importance to misclassified voxels at the next iteration. The strong classifier $H(\mathbf{x})$ is obtained by summing the weak classifiers, and yields a classification score that can be converted to probabilities as:

$$\begin{aligned} p_{tm}(\mathbf{x}) &= \frac{1}{1 + \exp(-2H(\mathbf{x}))} \\ p_{bg}(\mathbf{x}) &= 1 - p_{tm}(\mathbf{x}) \end{aligned} \tag{2}$$

The key element of the boosting algorithm is the selection of the feature vector. We adopt a high dimensional space exploring visual, phase and geometric properties. First, we rely on the intensity values using patches ($9 \times 9 \times 5$) centered on the sample voxel $\mathbf{x_i}$. Median, entropy and standard deviation values are extracted from another set of patches of sizes $k \times k \times 3$, where $k = \{7, 9, 11\}$. Second, we compute Gabor features

(Manjunath and Ma, 1996) on 2 scales and 10 orientations. We adopt the method of Zhan and Shen (2003) that approximates the 3D Gabor filters by computing two orthogonal 2D filter banks. Eventually, we compute a symmetry based feature, since the presence of the tumor will introduce an asymmetry between the hemispheres of the brain. Assuming a symmetry plane is known, the symmetry feature is computed as $Sym(\mathbf{x_i}) = \frac{1}{n}\sum_{\mathbf{x_j}\in\mathcal{N}(\mathbf{x_i})} I(\mathbf{x_j}) - \frac{1}{n}\sum_{\mathbf{x_{j,s}}\in\mathcal{N}(\mathbf{x_{i,s}})} I(\mathbf{x_{j,s}})$. $\mathcal{N}(.)$ is a neighborhood introduced to compensate the approximate symmetry plane, $\mathbf{x_{i,s}}$ is the symmetric of voxel $\mathbf{x_i}$ and $I(.)$ is the intensity value.

Spatial dependencies are introduced through an MRF model on a graph where each voxel of the image is a node and the edges connect the node to its 6 immediate neighbors. In this model, we define a binary label set $\mathcal{L}_s = \{0, 1\}$. Each node $\mathbf{x}$ (i.e image voxel) is to be assigned one label, tumor ($l_x = 1$) or background ($l_x = 0$). The optimal labeling is recovered by minimizing the MRF energy (Boykov and Funka-Lea, 2006):

$$E_{seg}(\mathbf{l}) = \sum_x V_x(l_x) + \sum_x \sum_y V_{x,y}(l_x, l_y) \tag{3}$$

The unary potentials $V_x(.)$ correspond to the classification likelihoods, seeking the most probable label according to the boosting classification decisions:

$$V_x(l) = -l_x \log(p_{tm}(\mathbf{x})) - (1 - l_x)\log(p_{bg}(\mathbf{x})) \tag{4}$$

And the pairwise term plays the part of a smoothing prior on the segmentation field, and is defined as a Potts model that penalizes neighboring nodes labeled differently:

$$V_{xy}(l) = \beta(1 - \delta(l_x, l_y)) \tag{5}$$

where β is a constant parameter describing the amount of smoothing.

The main drawback of this approach is the lack of global information on the brain structure, the spatial dependencies being encoded in a strictly local manner. Coupling segmentation with registration adds global information, but requires an efficient registration scheme.

2.2. Graph based Registration

Let us consider a source image A and a target image V defined on a domain Ω. In our case, the source image is a healthy brain and the target image is a diseased brain featuring a tumor. In the task of image registration, we want to find the geometric transformation $\mathcal{T}$ that will map the source image to the target image:

$$V(\mathbf{x}) = A \circ \mathcal{T}(\mathbf{x}) \tag{6}$$

We adopt the Free Form Deformation (FFD) approach (Rueckert et al., 1999), where a sparse grid $\mathcal{G} \subset \Omega$ is superimposed to the volume. The transformation will be evaluated on the grid's control points, and then on the whole volume by interpolation.

$$\mathcal{T}(\mathbf{x}) = \mathbf{x} + \sum_{p\in\mathcal{G}} \eta(\|\mathbf{x} - \mathbf{p}\|)\mathbf{d_p} \tag{7}$$

where $\mathbf{d_p}$ is the displacement of control point $\mathbf{p}$ and $\eta(.)$ is the projection function that describes the influence of each control point on voxel $\mathbf{x}$.

The most likely displacement should minimize the differences between the deformed image $A(\mathcal{T}(\mathbf{x}))$ and target image $V(\mathbf{x})$, evaluated by a similarity measure $\rho(.)$:

$$\mathcal{T}^{opt}(\mathbf{x}) = \arg\min_{\mathcal{T}} \sum_{p\in\mathcal{G}} \int_\Omega \bar{\eta}(\|\mathbf{x} - \mathbf{p}\|)\rho(V(\mathbf{x}), A(\mathcal{T}(\mathbf{x})))\, dx \tag{8}$$

The similarity measure is evaluated on the whole domain Ω. This information is back projected on the control points via the function $\bar{\eta}(.)$.

In order to recover the optimal control points' displacements, we adopt a discrete MRF model (Glocker et al., 2008a, 2011). Let us consider a discrete set of labels $\mathcal{L} = \{1, ..., n\}$, and a set of discrete displacements $\Delta = \{\mathbf{d^1}, ..., \mathbf{d^n}\}$. We seek to assign a label l_p to each grid node $\mathbf{p}$, where each label corresponds to a discrete displacement $\mathbf{d^{l_p}} \in \Delta$. In this setting, the deformation field is rewritten as:

$$\mathcal{T}(\mathbf{x}) = \mathbf{x} + \sum_{p\in\mathcal{G}} \eta(|\mathbf{x} - \mathbf{p}|)\mathbf{d^{l_p}} \tag{9}$$

In order to recover the optimal labeling, we need to minimize the MRF energy:

$$\mathbf{l^{opt}} = \arg\min_l E_{reg}(\mathbf{l}) \tag{10}$$

where

$$E_{reg}(\mathbf{l}) = \sum_{p\in\mathcal{G}} V_p(l_p) + \lambda \sum_{\mathbf{p}\in\mathcal{G}} \sum_{\mathbf{q}\in\mathcal{N}(\mathbf{p})} V_{p,q}(l_p, l_q) \tag{11}$$

where $\mathcal{N}(.)$ represents the neighborhood system, defined here as a 6-neighbors configuration. $V_{p,q}(.)$ is a pairwise potential, that imposes certain smoothness on the deformation.

The unary potential $V_p(.)$ is only dependent on node $\mathbf{p}$'s configuration and represent the likelihood of the node being assigned a label. To preserve the independence assumption, we can approximate the unary potentials as:

$$V_p(l_p) \approx \int_\Omega \bar{\eta}(\|\mathbf{x} - \mathbf{p}\|)\rho(V(\mathbf{x}), A(\mathbf{x} + d^{l_p}))\, dx \tag{12}$$

This approach shows great performance for the registration of healthy brains (Glocker et al., 2008a), but performs poorly in the tumor area where the similarity metric is not reliable. The most straightforward solution is to mask the pathology and not take the tumor voxels into account during the evaluation of the similarity criterion $\rho(.)$. This requires a very reliable segmentation map and would introduce a bias for the registration.

2.3. Concurrent Tumor Segmentation and Registration

Our approach aims at simultaneously performing tumor segmentation and atlas to diseased subject registration. The coupling of the segmentation with the registration of an atlas introduces global information on the brain structure, while the registration quality is improved by acknowledging the presence

of the tumor and treating it differently than healthy tissue during registration. The registration and segmentation energies are coupled in a single MRF framework, where the tumor is detected concurrently to the registration. In this combined framework, we seek to recover the optimal transformation $\mathcal{T}(\mathbf{x})$ and the segmentation map $\mathcal{S}(\mathbf{x})$.

Let us consider a sparse grid $\mathcal{G}$ superimposed to the volume, a discrete set of labels $\mathcal{L}_c = \{1, ..., 2n\}$, a predefined discrete set of displacements Δ, and the tumor $p_{tm}(\mathbf{x})$ and background $p_{bg}(\mathbf{x})$ prior probabilities learned via boosting. Each label $l \in \mathcal{L}_c$ is associated to a pair segmentation/displacement $\{s^l, \mathbf{d^l}\} \in \{0, 1\} \times \Delta$, where we define:

$$\begin{aligned} \mathbf{d^{l_p}} &= \mathbf{d^k} \text{ where } k = l_p \pmod n \\ s^{l_p} &= \left\lceil \frac{l_p}{n} \right\rceil \end{aligned} \tag{13}$$

We seek to assign a label l_p to each control point $\mathbf{p}$ of $\mathcal{G}$, simultaneously displacing the grid node and characterizing it as tumor or background. The segmentation and deformation fields are then evaluated on the whole volume by interpolation:

$$\begin{aligned} \mathcal{T}(\mathbf{x}) &= \mathbf{x} + \sum_{p \in \mathcal{G}} \eta(|\mathbf{x} - \mathbf{p}|)\mathbf{d^{l_p}} \\ \mathcal{S}(\mathbf{x}) &= \sum_{p \in \mathcal{G}} \eta(|\mathbf{x} - \mathbf{p}|)\mathbf{s^{l_p}} \end{aligned} \tag{14}$$

The MRF energy consist of segmentation and registration terms that are interdependent:

$$\begin{aligned} E_{reg,seg}(\mathbf{l}) = &\sum_{p \in \mathcal{G}} \underbrace{\alpha V_{reg}(l_p) + (1-\alpha) V_{seg}(l_p)}_{\text{Unary terms}} \\ + &\sum_{\mathbf{p} \in \mathcal{G}} \sum_{\mathbf{q} \in \mathcal{N}(\mathbf{p})} \underbrace{\alpha V_{pq,reg}(l_p, l_q) + (1-\alpha) V_{pq,seg}(l_p, l_q)}_{\text{Pairwise terms}} \end{aligned} \tag{15}$$

where α is a parameter balancing the importance of the segmentation and registration terms.

The pairwise costs ensure that the segmentation and registration are locally smooth. They are set as:

$$V_{pq,seg}(l_p, l_q) = \frac{|s^{l_p} - s^{l_q}|}{||\mathbf{p} - \mathbf{q}||} \tag{16}$$

$$V_{pq,reg}(l_p, l_q) = \lambda \frac{\left((\mathbf{d^{l_p}} - \mathbf{d^{l_q}}\right)^2}{||\mathbf{p} - \mathbf{q})||} \tag{17}$$

The strength of the pairwise cost depends on the distance between the connected nodes, taking into account a possible anisotropy as the distance between nodes would then differ. The closer the nodes are, the stronger the penalty imposing similar labels. The registration regularization's role is to preserve the anatomical structure of the brain. Important deformations can occur in and around the tumor area, requiring a relaxation of the pairwise cost to allow for those strong deformations.

Let us now proceed with the definitions of the unary potentials. Outside the tumor area ($s^{l_p} = 0$), the registration term seeks correspondences between the atlas and the target's healthy tissues via the similarity metric $\rho(.)$. However, this metric is not reliable in the tumor area ($s^{l_p} = 1$) since there are no existing correspondences. We use instead a constant cost C_{tm} that is independent of the chosen displacement:

$$\begin{aligned} V_{reg}(l_p) = &\int_\Omega \bar{\eta}(||\mathbf{x} - \mathbf{p}||) \\ &\left(s^{l_p} C_{tm}(\mathbf{x}) + (1 - s^{l_p})\rho(V(\mathbf{x}), A(\mathbf{x} + d^{l_p}))\right) dx \end{aligned} \tag{18}$$

The tumor probabilities and the similarity metric are evaluated on the whole volume and back projected on the control points. The use of a constant cost causes the displacement within the tumor area to be determined by interpolation with the neighboring nodes at the tumor boundary, through the pairwise regularization term. While the main role of this potential is registration of the two volumes, it allows detection of part of the tumor through the similarity measure. Indeed, if a strong dissimilarity between voxels is observed, it is likely that the area belongs to the tumor.

This potential alone is however not sufficient for a precise segmentation of the tumor, due to the fact that the tumor's local appearance can be similar to healthy tissue and that dissimilar voxels do not necessarily correspond to tumors. Additional information on the position of the tumor is introduced by coupling this registration term with a segmentation unary term. This term relies on the prior probabilities introduced in section 2.1, imposing the label that has the maximum likelihood probability:

$$\begin{aligned} V_{seg}(l_p) = &\int_\Omega \bar{\eta}(||\mathbf{x} - \mathbf{p}||) \\ &\left(-s^{l_p} \log(p_{tm}(\mathbf{x} + \mathbf{d^{l_p}})) - (1 - s^{l_p}) \log(p_{bg}(\mathbf{x} + \mathbf{d^{l_p}}))\right) dx \end{aligned} \tag{19}$$

The tumor segmentation will therefore be determined taking into account anatomical prior knowledge based on the healthy reference (introduced through the registration term) and the classification decisions. The segmentation is determined on the sparse grid associated to the moving target image, it is therefore dependent on the registration as the probability maps are not aligned with the target image at the start of the process. The position of the node after its displacement corresponds to the area that is segmented. As the registration improves, the segmentation quality does as well.

The optimal labeling is recovered using a linear programming based optimization method (Komodakis et al., 2008) that offers a great compromise between speed and accuracy.

3. Uncertainty-driven Adaptive Resampling

The main drawback of discrete approaches is the trade-off between precision and computational complexity. The search space (displacement label set) for registration would ideally cover the entire area, while a high grid resolution is required to register fine details and most importantly, to detect the tumor's irregular boundary. However, both are limited in order to maintain the computational burden manageable.

These drawbacks are usually dealt with using a hierarchical approach through the use of coarse to fine grid resolutions. This

enables to cover large and precise deformations with a limited search space, and makes the segmentation more robust by propagating segmentation decisions that are less sensitive to small variations in the image. In this context, several grid resolutions $\mathcal{G}_j$ are considered, with a series of iteration t at each resolution. At level $\{\mathcal{G}_j, t\}$, the new MRF energy is computed, based on the deformed source image, evaluated at iteration t-1:

$$l_t^{opt} = \arg\min_l E_{reg,seg}(l|V, A \circ T^{t-1}) \tag{20}$$

The displacement information is propagated from one level to the next by composing the new displacement field with the one obtained at iteration $t-1$. There are two main challenges in this approach: (i) the sampling of the discrete deformation space at each iteration and (ii) the grid resolutions. The most straightforward approach is a uniform refinement of the label set and grid resolution. This allows for precise results but ignores the local anisotropy of the structures. Furthermore, quasi voxel-level resolutions are necessary for segmentation, which cannot be considered in the context of uniform grids. Shi et al. (2012) proposed the Sparse Free Form Deformations: the multi-level grids are optimized simultaneously with a sparsity constraint across levels, ensuring that nodes in high resolution levels are given more importance in areas with discontinuities, and low importance otherwise.

Relying on local segmentation and registration uncertainties offers an alternative and enables to define an adaptive content-driven grid refinement. Such measurements can lead to computationally efficient voxel level resolutions while capturing the local anisotropy of the structure for a more efficient registration. The min-marginals measure the variations of the energy under different constraints (Kohli and Torr, 2008) and have been considered in the context of a discrete registration framework (Glocker et al., 2008b) to evaluate the local uncertainty and adapt the displacement sampling accordingly. We are inspired by this approach that we combine with segmentation uncertainty in order to define an adaptive displacement and node sampling.

3.1. Min-marginals and Displacement Sampling

Let us consider a control point $\mathbf{c_j} \in \mathcal{G}_j$ at iteration t, and its corresponding optimal labeling $l_{c_j}^{opt}$. We aim at defining the displacement sampling at the next iteration as well as the resolution of the next grid level $\mathcal{G}_{j+1}$ based on the volumes' local properties.

Our approach exploits the min-marginal energies (Kohli and Torr, 2008) that evaluate the minimum value of the MRF energy under different constraints. By imposing a label k, different from the optimal label $l_{c_j}^{opt}$, to control point $\mathbf{c_j}$, the min-marginals indicate how much a label swap costs.

$$\Psi_{\mathbf{c_j},k,t} = \min_{l, l_{c_j}=k} E_{reg,seg}(\mathbf{l}|V, A \circ \mathcal{T}^{t-1}) \tag{21}$$

Let us recall that the label l_{c_j} corresponds to a pair $\{\mathbf{d}^{\mathbf{l_{c_j}}}, s^{l_{c_j}}\}$, therefore, both segmentation and registration uncertainties can be extracted from the min-marginals. If the segmentation label is constant ($s^k = s^{l_{c_j}}$), a label swap represent a local perturbation from the optimal displacement. A small energy variation means that displacement $\mathbf{d^k}$ is almost as likely as the optimal displacement $\mathbf{d}^{\mathbf{l_{c_j}^{opt}}}$, highlighting the uncertain labeling with respect to that direction. Inversely, the labeling is quite certain in a direction where a perturbation yields a high increase of energy. By normalizing the min-marginals over all the possible displacements associated to the same segmentation label, we can compute the registration uncertainty:

$$U_{reg}(\mathbf{c_j}, k) = \frac{\exp(-\Psi_{\mathbf{c_j},k,t})}{\sum_{l \in \mathcal{L}_c, s^l = s^{l_{c_j}^{opt}}} \exp(-\Psi_{\mathbf{c_j},l,t})} \tag{22}$$

The highest $U_{reg}(.)$ correspond to the most likely labels. The registration uncertainty computed over all the possible displacements can be approximated to a Gaussian distribution (see Fig.[1]), whose covariance evaluates the local anisotropy. The search space is resampled following the covariance matrix main axes and scales, allowing for a more thorough evaluation of the deformation space in the uncertain areas. The registration uncertainty is not taken into account for the tumor label and when the parameter α is low where the deformation is mostly driven by the pairwise cost, yielding a spherical covariance matrix.

3.2. Uncertainty-driven Graph Refinement

Consider a uniform grid $\mathcal{G}_{j,max}$ of resolution $M \times N \times P$, and $\mathcal{A}_j : \mathcal{G}_{j,max} \rightarrow \{0, 1\}$ an activation function describing the resolution of the current adaptively sampled grid $\mathcal{G}_j \subset \mathcal{G}_{j,max}$. At the next resolution level, $\mathcal{G}_{j,max}$ is refined as a grid $\mathcal{G}_{j+1,max}$ of resolution $2M-1 \times 2N-1 \times 2P-1$, splitting all existing edges in two. The grid $\mathcal{G}_{j+1,max}$ represents the new level's maximal resolution, corresponding to a uniform sampling. The new grid $\mathcal{G}_{j+1}$ resolution is determined by activating relevant nodes while ignoring the ones that are not necessary to increase the quality of the registration or segmentation.

A node $\mathbf{p} \in \mathcal{G}_{j+1,max}$ can be activated ($\mathcal{A}_{j+1}(\mathbf{p}) = 1$) if it satisfies at least one of those three conditions: (i) the node has a direct correspondent $\mathbf{c_j} \in \mathcal{G}_{j,max}$ (same coordinates) that is activated ($\mathcal{A}_j(\mathbf{c_j}) = 1$), (ii) it is connected to nodes in $\mathcal{G}_j$ that have a high segmentation uncertainty (segmentation activation), (iii) it is connected to nodes in $\mathcal{G}_j$ that have a high registration uncertainty (registration activation). The registration and segmentation activations are determined via the definition of two activation terms $A_r(\mathbf{p})$ and $A_s(\mathbf{p})$ respectively, both taking value in $\{0, 1\}$. To propagate the min-marginals and activation information, we define an inter level neighborhood system $\mathcal{N}_i(.)$ by connecting $\mathbf{c_j} \in \mathcal{G}_{j,max}$ to its 27 closest neighbors (based on the image's spatial coordinates) in $\mathcal{G}_{j+1,max}$. The neighborhood system is shown in Fig.[2]

The registration activation criterion relies on the idea that small and precise displacements are necessary around salient structures, while increasing the resolution on homogeneous regions is not necessary. A node should be activated if it is at the interface of adjacent structures. Considering a node $\mathbf{c_j} \in \mathcal{G}_j$, it covers an image region delimited by its maximum displacement. If the region is not homogeneous, there will be strong

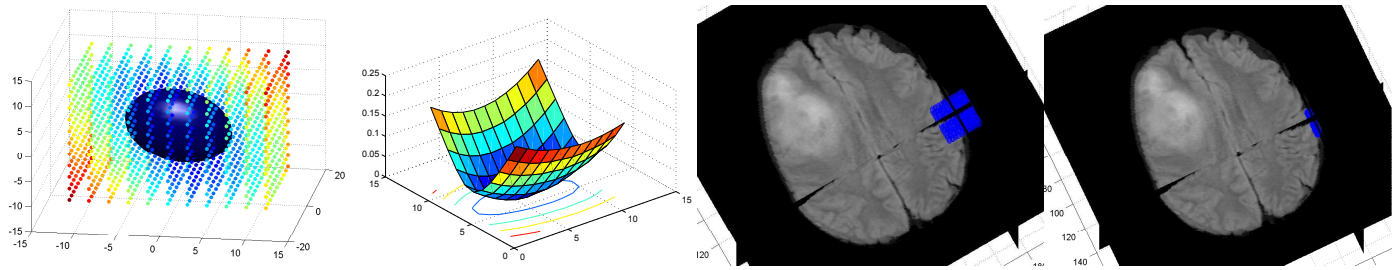

(a) (b) (c) (d)

Figure 1: Registration uncertainty and displacement set resampling for one control point: (a) Min-marginal values per displacement label (blue: low, red: high energy) associated covariance matrix centered at the optimal label, (b) Min marginals visualization on a 2D slice, (c) Original isotropic displacement set and (d) Uncertainty driven displacement set, following the brain boundaries.

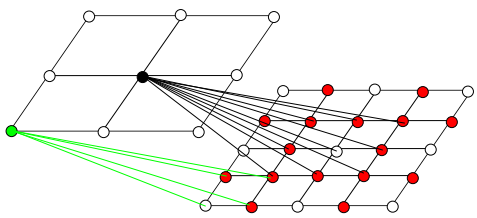

Figure 2: Visual representation of the grid refinement from level j (left) to level j+1 (right). Grid resampling: the nodes that have direct correspondences appear in white, and the new nodes are red. The edges connecting the 2 grids represent the nodes' neighborhood. The grid is shown in 2D for increased visibility.

min-marginal energies variations with respect to the displacement label. The activation criterion is based on the node's energy range and is defined as:

$$A_r(\mathbf{p}) = H\left(\sum_{\mathbf{c_j}\in\mathcal{N}_i(\mathbf{p})} \frac{1}{N}\left(\max_{\substack{k\in\mathcal{L}_c,\\ s^k=s^{l_{c_j}^{opt}}}} \Psi_{\mathbf{c_j},k,t} - \min_{\substack{k\in\mathcal{L}_c,\\ s^k=s^{l_{c_j}^{opt}}}} \Psi_{\mathbf{c_j},k,t}\right) - \mu\right) \quad (23)$$

μ is the mean value over all activated nodes in $\mathcal{G}_j$, H(.) is the heaviside step function, and N is the number of nodes in the neighborhood of $\mathbf{p}$. The node $\mathbf{p}$ will be activated if the mean energy range among its neighbors in $\mathcal{G}_j$ is higher than the mean range over all nodes.

Similarly, the segmentation node activation is based on the segmentation uncertainty that can be evaluated by measuring the energy variation when the segmentation label changes. The uncertainty with respect to one segmentation label S can be computed by normalizing over all labels:

$$U_{seg}(\mathbf{c_j}|S) = \frac{\sum_{l\in\mathcal{L}_c, s^l=S} \exp(-\Psi_{\mathbf{c_j},l,t})}{\sum_{l\in\mathcal{L}_c} \exp(-\Psi_{\mathbf{c_j},l,t})} \quad (24)$$

In the case of a binary segmentation, we can simply reformulate the uncertainty as:

$$U(\mathbf{c_j}) = 1 - |U_{seg}(\mathbf{c_j}|S) - 0.5| \quad (25)$$

This term measures how certain the chosen label is. A low value of $U(\mathbf{c_j})$ infers a highly reliable labeling. We seek to propagate the segmentation decisions to the next grid level $\mathcal{G}_{j+1}$ based on their reliability, so that the focus is on uncertain areas. This is achieved by adding an inter-level pairwise potential to the global energy :

$$V_{ucy}(l_{c_j}^{opt}, l_p) = \begin{cases} 0 & \text{if } s^{l_{c_j}^{opt}} = s^{l_p} \\ -log(U(\mathbf{c_j})) & \text{Otherwise} \end{cases} \quad (26)$$

where $\mathbf{c_j}$ is a control point in G_j in the neighborhood of $\mathbf{p}$. This potential penalizes nodes in G_{j+1} that are assigned a label different than their neighbor in $\mathcal{G}_j$. The amount of penalty depends on how certain the labeling of $\mathcal{G}_j$ is. In this neighborhood configuration, a node in $\mathcal{G}_{j+1}$ can be influenced by several nodes in $\mathcal{G}_j$, so that there is no penalty when the nodes labels are different and equally likely, the new node being situated at the tumor boundary. The segmentation activation criterion is controlled by the strength of the penalty and defined as:

$$A_s(\mathbf{p}) = H\left(\left|\sum_{\mathbf{c_j}\in\mathcal{N}_i(\mathbf{p})} (-1)^{s^{l_p}} V_{ucy}(l_{c_j}^{opt}, l_p)\right| - tsh\right) \quad (27)$$

Where N is the number of nodes in $\mathcal{N}_i(\mathbf{p})$, tsh is a threshold parameter and H(.) is the Heaviside step function. This term measures how strong the penalty is on node $\mathbf{p}$, taking into account the fact that there is no penalty if its neighbors are labeled differently with equally confident labels. Nodes with a low overall penalty will be activated.

Eventually, we can rewrite the MRF energy at resolution level j and iteration t:

$$\begin{aligned} E_{reg,seg}(l|V, A\circ\mathcal{T}^{t-1}, \mathcal{G}_j) &= \frac{1}{|\mathcal{G}_j|}\sum_{p\in\mathcal{G}_j} V_p(l_p|V, A\circ\mathcal{T}^{t-1}) \\ &+ \sum_{p\in\mathcal{G}_j}\sum_{q\in\mathcal{N}(\mathbf{p})} V_{pq,}(l_p, l_q) + \sum_{p\in\mathcal{G}_j}\sum_{c\in\mathcal{N}_i(\mathbf{p})} \frac{1}{N} V_{ucy}(l_c^{opt}, l_p) \end{aligned} \quad (28)$$

where N is the number of nodes in $\mathcal{G}_j$ that are connected to node $\mathbf{p} \in \mathcal{G}_{j+1}$, and $\mathcal{N}_i(\mathbf{p})$ is the corresponding neighborhood.

4. Experimental Validation

Our data set consisted of 110 3D FLAIR MRI volumes of different patients featuring a low-grade glioma prior any treatment. The complete tumors have been manually segmented in

all volumes by experts. Although additional modalities would have offered increased segmentation quality, there were not systematically available and provided by our clinical partners. We therefore focused on the FLAIR modality.

36 volumes were randomly selected for boosting learning. We tested our joint segmentation and registration framework on the 74 remaining volumes. The reference pose for registration was a 3D FLAIR MRI volume of a single healthy subject of size $256 \times 256 \times 24$ and resolution $0.9 \times 0.9 \times 5.5\ mm^3$. The absence of existing multi subject healthy atlases of FLAIR modality has motivated the use of a single subject as reference pose to evaluate the algorithm.

We ran an additional set of experiments on the 10 real low-grade gliomas cases of the BRATS training database for an easier comparison with existing methods. We segmented the complete tumors (including active tumor, necrotic core and occasional edema). The results can be compared for the low-grade glioma case (complete tumor) with the results presented on the training set in the BRATS proceedings. In order to maintain consistency with our FLAIR database, only FLAIR images were considered for boosting training which was carried out through leave one out cross validation experiments. Registration was performed using the T2-weighted images due to the insufficient quality of the FLAIR images. Furthermore, this allows the use of the T2-weighted MNI-ICBM multi subject atlas (Fonov et al., 2009) of size $193 \times 229 \times 193$ and resolution $1 \times 1 \times 1$ as reference for registration, which is more adapted to anatomical differences between subjects than a single subject reference pose. This results in exploiting two of the four available modalities for segmentation and the T2-weighted modality for registration. The reference poses for registration are shown in Fig. [3].

As preprocessing, all volumes were skullstripped and rigidly registered to their reference pose (Ourselin et al., 2000). Their intensity was regularized by simply setting all volumes to the same median and interquartile range as the reference pose. Since all volumes are rigidly registered, an approximate symmetry plane of the reference pose is used for all volumes to evaluate the boosting symmetry feature. We compared the joint registration and segmentation framework with a sequential approach, where the tumor is segmented using the single boosting based MRF method and the segmentation is used as a mask for registration (not taking into account the segmented area). To demonstrate the potential of the adaptive resampling framework, we compared the results without uncertainties where the segmentation is propagated from one resolution to the next using a manually set penalty cost:

$$V^t_{pen}(l_p) = \int_\Omega \bar{\eta}(\|\mathbf{x} - \mathbf{p}\|) s^{l_p} \mathcal{S}^{t-1}(\mathbf{x}) \frac{2}{\exp(j)}\, dx \qquad (29)$$

where j is the resolution level and t the current iteration. The uncertainty based framework was compared with this approach at the maximal and same final grid resolution.

4.1. Implementation

The same set of parameters were used for all volumes in the FLAIR database and were determined heuristically in order to obtain the best possible results over the whole database. Our coarse to fine hierarchical approach consisted of 3 image levels and 4 grid levels, where the resolution of the image increases with the grid resolution. The maximal grid resolution increased from $9 \times 9 \times 5$ to $65 \times 65 \times 37$. In accordance with the Free Form Deformation framework, the projection function used was cubic B-splines. We set the parameter α so that the presence of the tumor has an increasing impact on the registration. It is progressively diminished from 1 to 0.015, the focus being on segmentation at the finest level. This setting enables to focus on aligning the main brain structures at coarse resolutions where the tumor is only roughly detectable then progressively increase the segmentation precision. The constant cost C_{tm} for registration is progressively increased, initially set to 5 and 6 times the mean value of the similarity criterion without and with uncertainties respectively. The parameter λ describing the influence of the registration smoothing was set to 20 and relaxed in the tumor area to allow for the potentially important displacements induced by the tumor. The threshold tsh for node activation was set to 1.6.

We perform 3 iterations at each grid level. Without exploitation of the uncertainty information, the displacement sampling is sparse (31 labels, sampled along the main axes) and refined at each iteration by reducing the maximum displacement. We adopt a dense sampling (1331 labels) to compute the local uncertainties at the first iteration, and a sparse sampling at the 2 remaining iterations, the labels being sampled along the covariance matrix main axes. This enables to exploit the uncertainty information with limited impact on the run time. When α is low, the local anisotropy cannot be captured efficiently by the min marginals. A sparse sampling is adopted for all iterations at the last 2 grid levels.

The parameters were adapted to the BRATS dataset, setting the maximal grid resolution from $11 \times 12 \times 11$ to $81 \times 96 \times 81$, the constant cost to 6 and 7 times the mean value of the similarity criterion and progressively increased and the threshold for grid activation as the mean penalty value over all active nodes. The same parameters were used for the 10 volumes. Experiments were only carried out with sparse sampling.

4.2. Uncertainty based Grid Nodes Activation

The percentage of activated nodes, with respect to the maximum uniform resolution is shown in Fig.[4]. The complexity of the framework is considerably reduced, only activating less than 20 % of the nodes at the last level for both datasets. Our current implementation associates changes of labels of inactive nodes with infinite costs and runs for less than a minute (sparse sampling) or 2 to 4 minutes (sparse/dense sampling) on the FLAIR database, and 3 to 8 minutes on the BRATS database of higher resolution. A direct construction of the grid would significantly impact the run time and memory cost. Considering an MRF with L, E and N being respectively the number of labels, edges and nodes, we provide a complexity analysis (excluding the time required to build the data term potentials):

- Computational cost:
  $O\left(L \times O\left(E \times N^2\right)\right) \sim O\left(L \times E \times N^2\right)$ per iteration

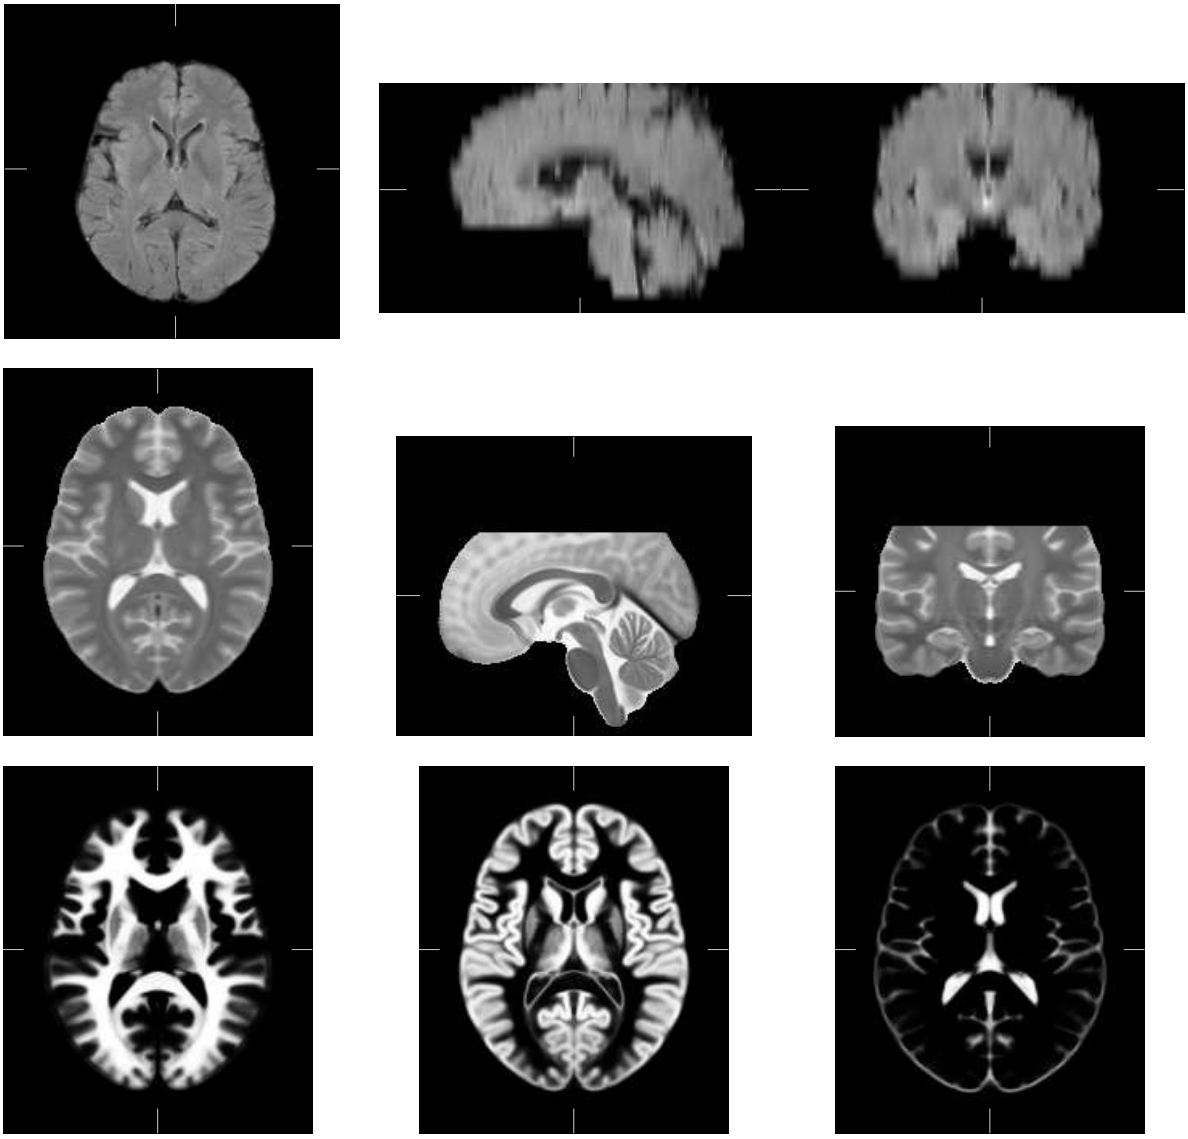

Figure 3: Healthy reference poses used for registration. FLAIR (First row) and T2 MNI-ICBM atlas (second row). Bottom row: MNI atlas Probability maps, from left to right: White Matter, Gray Matter and CSF.

- Memory cost: $\mathcal{O}(L \times \mathcal{O}(N + E))$

Reducing the number of nodes to approximately 20% at the finer resolution scale leads to:

- Computational cost:
  $\mathcal{O}\left(L \times \mathcal{O}\left(0.2E \times (0.2N)^2\right)\right) \sim \mathcal{O}\left(0.008L \times E \times N^2\right)$
  When taking into account the number of iterations, we can obtain a complexity that is approximately 3-4 orders of magnitude lower.

- Memory cost: $\mathcal{O}(L \times \mathcal{O}(0.2(N + E)))$, approximately one order of magnitude lower.

Fig. [11] shows visual examples of the last two grid levels' resolution. Nodes are activated around the brain's structures and the tumor's boundary, demonstrating the adequacy of the registration and segmentation activation terms.

4.3. Segmentation Evaluation

The segmentation results were evaluated by comparing the automatic segmentation AS to the manual segmentation M. To this end, we compute the Dice score $Dc = \frac{2\|M \cap AS\|}{\|M\| + \|AS\|}$, the rate of false positives $FPR = \frac{\|AS\| - \|M \cap AS\|}{\|AS\|}$ and the rate of true positives $TPR = \frac{\|M \cap AS\|}{\|M\|}$ and mean absolute distance between contours (MAD). Segmentations were evaluated after reverting to the patient's space (before rigid registration) where the manual segmentations were performed. Segmentations of higher quality are obtained using the joint framework on the FLAIR dataset (especially highlighted by the MAD score), and are equivalent with a high resolution uniform grid and an adaptively sampled low resolution grid, while the low resolution uniform grid yields poor tumor detection. Among the test set, 27 volumes have been manually segmented by 2 different experts. The inter expert Dice score reaches 89 % in median and gets as low as 76 %,

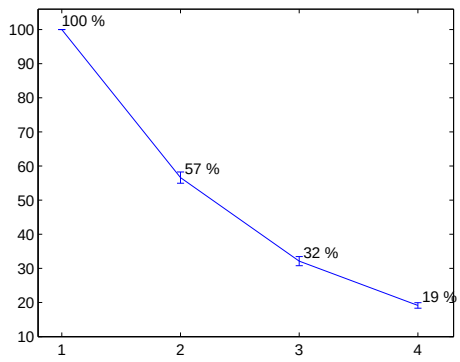

(a)

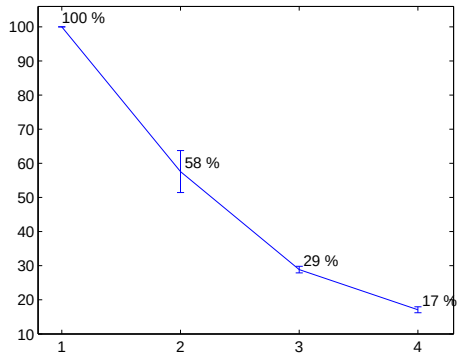

(b)

Figure 4: Mean percentage of activated nodes per level. (a) FLAIR database, (b) BRATS database.

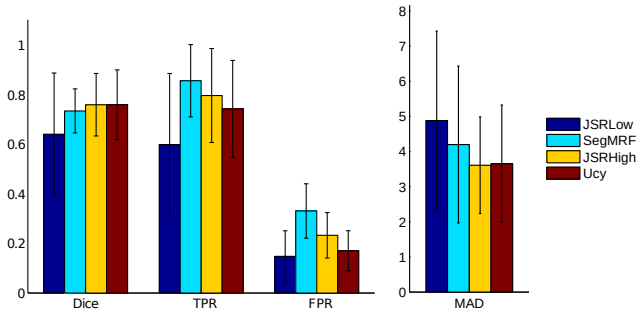

Figure 5: Quantitative Segmentation Results, FLAIR database: Error bars (mean and standard deviation) of the Dice score, False Positive (FPR) and True Positive (TPR) rates and MAD score (in millimeters) for the joint framework with low (JSRLow) and high resolution (JSRHigh), the individual segmentation framework (SegMRF) and the uncertainty based approach (Ucy).

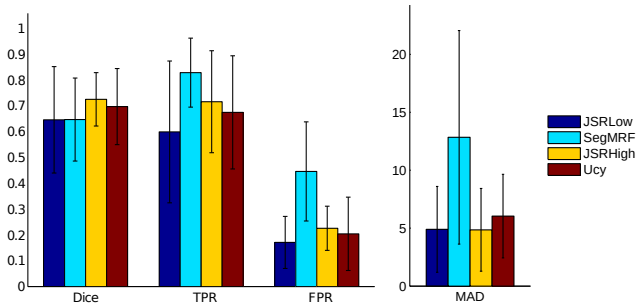

Figure 6: Quantitative Segmentation Results, BRATS database: Error bars (mean and standard deviation) of the Dice score, False Positive (FPR) and True Positive (TPR) rates and MAD score (in millimeters) for the joint framework with low (JSRLow) and high resolution (JSRHigh), the individual segmentation framework (SegMRF) and the uncertainty based approach (Ucy).

which highlights the high inter expert variability with respect to the manual segmentations of tumors and is close to the obtained automatic segmentations (81% median over the 27 volumes). Error bars of the different scores are shown in Fig.[5].

Segmentation results on the BRATS dataset are on par with results obtained by the BRATS 2012 challenge winners, (mean Dice score 70-72 %, median 72-73 % with and without uncertainties respectively), while a strong increase of quality is obtained with respect to the regularized boosting results (mean Dice score 65 %). The tumors are poorly detected using a low uniform resolution (mean Dice score 64 %, reduction of the true positive rate of 7% (with uncertainties) and 11% (without uncertainties)). Error bars of the different scores are shown in Fig.[6].

Visual segmentation results for both datasets are shown in Fig.[8].

4.4. Registration Evaluation

The registration was evaluated mostly qualitatively. For quantitative analysis, the ventricles where manually segmented for 33 volumes of the FLAIR dataset and the Dice score, False and True Positive rate, and MAD were computed between the registered source image and the target image outside the tumor area and are shown in Fig.[7]. Fig.[9] shows visual registration results comparing the joint registration and segmentation framework to the individual registration where the pathology has been masked. Quantitative results show equivalent performance outside the tumor area, while visual examples show a high increase in quality of registration in and around the tumor area, in cases where the individual framework fails. Quantitative results also highlight the maintained performance outside the tumor area using an adaptively sampled grid, and a lower quality registration considering a uniform grid of equivalent low resolution.

Visual examples of registration results on the BRATS dataset are shown in Fig. [10]. Cases where the individual framework and low resolution registration fail are illustrated, as well as the obtained complete segmentation of an image.

5. Conclusion and Discussion

In this paper we have presented a concurrent registration and tumor segmentation framework that exploits the interdependencies between the two problems. We adopt a discrete graphical model on a sparse grid superimposed to the image domain. Each grid node is simultaneously classified and displaced

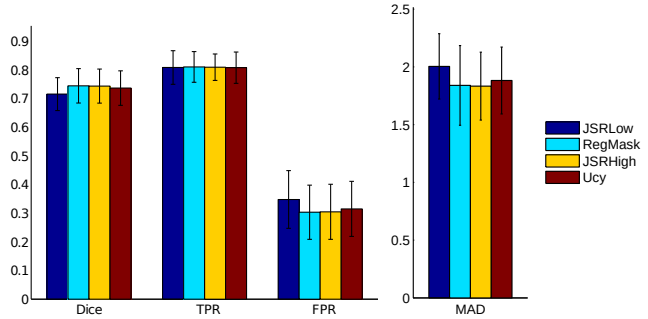

Figure 7: Quantitative Registration results, FLAIR database: Error bar graphs of the Dice, True Positives (TPR), False Positives (FPR) and MAD scores (in millimeters) obtained for the joint framework with low (JSRLow) and high resolution (JSRHigh), the individual registration framework with masked pathology (RegMask) and the uncertainty based approach (Ucy)

based on a boosting classifier and image similarity. The detected tumor area is registered by interpolation with the neighboring nodes. The progressive impact of the tumor segmentation on the registration allows to deal with the presence of the tumor without the introduction of an initial bias that can lead to registration errors, while the introduction of spatial information on the brain structures significantly reduces the false detections. The inclusion of uncertainties enables to deal with the main drawback of discrete approaches, that is the trade off between precision and computational complexity. Adaptive refinement of the sparse grid yields a much lower complexity framework. While our current implementation simply discards inactive nodes, direct construction of the non uniform grid would lead to significant diminution of the run time.

The framework offers great modularity with respect to the similarity measure for registration, the segmentation prior probabilities estimation, the image modality and the clinical context. We presented the method in the context of diffuse low-grade gliomas and registration/segmentation of a healthy subject/atlas to a subject with a tumor. Aside for enhanced segmentation quality through the healthy brain's anatomical information (obtained segmentation results are close to the inter expert variability), this offers the possibility to build statistical atlases of tumor appearances in the brain and to evaluate the impact of the tumors on the brain's functional organization. Furthermore, the method's modularity allows easy adaptation to different problems where correspondences are missing, such as registration between pre operative and intra/post operative images with tumor resection for surgical guidance.

The choice of pathology masking (instead of growth models) is justified by the infiltrative nature of the low-grade gliomas, and coupled with the discrete formulation and adaptive sampling, results in a fast algorithm that shows great performance. It is however not adapted to fast growing and space occupying tumors that yield strong deformation that would have to be modeled accordingly.

One limitation of the method is its dependency on the boosting classification output. It is progressively refined through increasing resolution levels and segmentation propagation/penalty

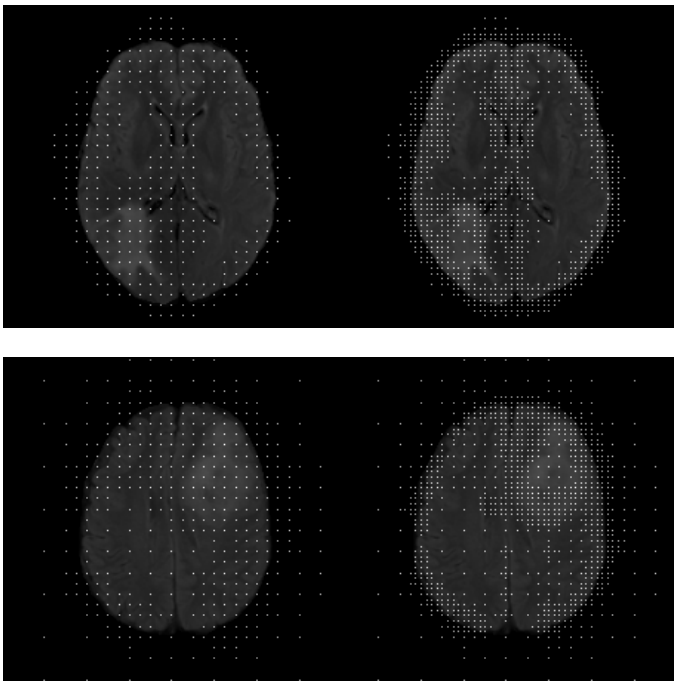

Figure 11: Visual examples of the activated nodes for the last 2 levels of the incremental approach. The nodes are superimposed to the target image.

across levels but still constitutes the baseline of the obtained segmentation. Two natural extensions of the algorithm are the inclusion of multimodal information and multiclass segmentation (to separate tumor core, necrosis and edema). Both can easily be introduced in the model through the boosting feature vector and by increasing the number of labels (both during boosting classification and construction of the MRF model). Such extensions are likely to increase the quality of the detection and segmentation.

Last but not least, a drawback of the method is its requirement for manual setting of important model parameters that can have a strong impact on the results. Optimal learning of the weights ((Komodakis, 2011) of the graphical model being considered in this paper from training data will allow to eliminate the need of manual parameter setting.

Acknowledgments

This work was supported by ANRT (grant 147/2010), Intrasense, the European Research Council Starting Grant Diocles (ERC-STG- 259112) and by NIH grant P41EB015898.

References

Ashburner, J., Friston, K.J., 2005. Unified segmentation. Neuroimage 26, 839–851.

Bauer, S., Nolte, L.P., Reyes, M., 2011. Fully automatic segmentation of brain tumor images using support vector machine classification in combination with hierarchical conditional random field regularization, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI. Springer, pp. 354–361.

Berendsen, F.F., Van Der Heide, U.A., Langerak, T.R., Kotte, A.N., Pluim, J.P., 2013. Free-form image registration regularized by a statistical shape model: application to organ segmentation in cervical mr. Computer Vision and Image Understanding 117, 1119–1127.

Boykov, Y., Funka-Lea, G., 2006. Graph cuts and efficient n-d image segmentation. International Journal of Computer Vision 70, 109–131.

Brett, M., Leff, A.P., Rorden, C., Ashburner, J., 2001. Spatial normalization of brain images with focal lesions using cost function masking. Neuroimage 14, 486–500.

Chitphakdithai, N., Duncan, J.S., 2010. Non-rigid registration with missing correspondences in preoperative and postresection brain images, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI. Springer, pp. 367–374.

Clark, M.C., Hall, L.O., Goldgof, D.B., Velthuizen, R., Murtagh, F.R., Silbiger, M.S., 1998. Automatic tumor segmentation using knowledge-based techniques. IEEE Transactions on Medical Imaging 17, 187–201.

Cobzas, D., Birkbeck, N., Schmidt, M., Jagersand, M., Murtha, A., 2007. 3d variational brain tumor segmentation using a high dimensional feature set, in: IEEE 11th International Conference on Computer Vision - ICCV, IEEE. pp. 1–8.

Corso, J.J., Sharon, E., Dube, S., El-Saden, S., Sinha, U., Yuille, A., 2008. Efficient multilevel brain tumor segmentation with integrated bayesian model classification. IEEE Transactions on Medical Imaging 27, 629–640.

Cuadra, M., Pollo, C., Bardera, A., Cuisenaire, O., Villemure, J., Thiran, J.P., 2004. Atlas-based segmentation of pathological mr brain images using a model of lesion growth. IEEE Transactions on Medical Imaging 23, 1301–1314.

Fletcher-Heath, L.M., Hall, L.O., Goldgof, D.B., Murtagh, F.R., et al., 2001. Automatic segmentation of non-enhancing brain tumors in magnetic resonance images. Artificial Intelligence in Medicine 21, 43–63.

Fonov, V.S., Evans, A.C., McKinstry, R.C., Almli, C.R., Collins, D.L., 2009. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. NeuroImage 47, S102. URL: `http://dx.doi.org/10.1016/s1053-8119(09)70884-5`, doi:10.1016/s1053-8119(09)70884-5.

Friedman, J., Hastie, T., Tibshirani, R., 2000. Additive Logistic Regression: a Statistical View of Boosting. The Annals of Statistics 38, 337–407.

García, C., Moreno, J.A., 2004. Kernel based method for segmentation and modeling of magnetic resonance images, in: Advances in Artificial Intelligence–IBERAMIA 2004. Springer, pp. 636–645.

Gering, D.T., Grimson, W.E.L., Kikinis, R., 2002. Recognizing deviations from normalcy for brain tumor segmentation, in: Dohi, T., Kikinis, R. (Eds.), Medical Image Computing and Computer-Assisted Intervention - MICCAI, Springer. pp. 388–395.

Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N., 2008a. Dense image registration through MRFs and efficient linear programming. Medical Image Analysis 12, 731–741.

Glocker, B., Paragios, N., Komodakis, N., Tziritas, G., Navab, N., 2008b. Optical flow estimation with uncertainties through dynamic mrfs., in: CVPR.

Glocker, B., Sotiras, A., Komodakis, N., Paragios, N., 2011. Deformable medical image registration: Setting the state of the art with discrete methods*. Annual review of biomedical engineering 13, 219–244.

Gooya, A., Pohl, K.M., Bilello, M., Biros, G., Davatzikos, C., 2011. Joint Segmentation and Deformable Registration of Brain Scans Guided by a Tumor Growth Model, in: G. Fichtinger, A.Martel, and T. Peters: MICCAI 2011, Part II, LNCS 6892, pp. 532-540.

Görlitz, L., Menze, B., Weber, M.A., Kelm, B., Hamprecht, F., 2007. Semi-supervised tumor detection in magnetic resonance spectroscopic images using discriminative random fields, in: Pattern Recognition. Springer, pp. 224–233.

Ho, S., Bullitt, E., Gerig, G., 2002. Level-set evolution with region competition: automatic 3-d segmentation of brain tumors, in: 16th International Conference on Pattern Recognition, IEEE. pp. 532–535.

Kaus, M., Warfield, S., Nabavi, A., Black, P., Jolesz, F., Kikinis, R., 2001. Automated segmentation of mr images of brain tumors. Radiology 218, 586–591.

Kohli, P., Torr, P.H.S., 2008. Measuring uncertainty in graph cut solutions. Computer Vision and Image Understanding 112, 30–38.

Komodakis, N., 2011. Efficient training for pairwise or higher order crfs via dual decomposition, in: CVPR, pp. 1841–1848.

Komodakis, N., Tziritas, G., Paragios, N., 2008. Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal-dual strategies. Computer Vision and Image Understanding 112, 14–29.

Kyriacou, S.K., Davatzikos, C., Zinreich, S.J., Bryan, R.N., 1999. Nonlinear elastic registration of brain images with tumor pathology using a biomechanical model [mri]. IEEE Transactions on Medical Imaging 18, 580–592.

Lee, C.H., Wang, S., Murtha, A., Brown, M.R., Greiner, R., 2008. Segmenting brain tumors using pseudo–conditional random fields, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI. Springer, pp. 359–366.

Mahapatra, D., Sun, Y., 2012. Integrating segmentation information for improved mrf-based elastic image registration. IEEE Transactions on Medical Imaging 21, 170–183.

Manjunath, B.S., Ma, W.Y., 1996. Texture features for browsing and retrieval of image data. Pattern Analysis and Machine Intelligence, IEEE Transactions on 18, 837–842.

Menze, B.H., Van Leemput, K., Lashkari, D., Weber, M.A., Ayache, N., Golland, P., 2010. A generative model for brain tumor segmentation in multi-modal images, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI. Springer, pp. 151–159.

Mohamed, A., Zacharaki, E.I., Shen, D., Davatzikos, C., et al., 2006. Deformable registration of brain tumor images via a statistical model of tumor-induced deformation. Medical image analysis 10, 752–763.

Ou, Y., Sotiras, A., Paragios, N., Davatzikos, C., 2011. Dramms: Deformable registration via attribute matching and mutual-saliency weighting. Medical image analysis 15, 622–639.

Ourselin, S., Roche, A., Prima, S., Ayache, N., 2000. Block matching: A general framework to improve robustness of rigid registration of medical images, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2000, Springer. pp. 557–566.

Pallud, J., Taillandier, L., Capelle, L., Fontaine, D., Peyre, M., Ducray, F., Duffau, H., Mandonnet, E., 2012. Quantitative morphological magnetic resonance imaging follow-up of low-grade glioma: a plea for systematic measurement of growth rates. Neurosurgery 71, 729–39; discussion 739–40.

Parisot, S., Duffau, H., Chemouny, S., Paragios, N., 2011. Graph Based Spatial Position Mapping of Low-Grade Gliomas, in: G. Fichtinger, A.Martel, and T. Peters: MICCAI 2011, Part II, LNCS 6892, pp. 508-515. Springer, Heidelberg.

Parisot, S., Duffau, H., Chemouny, S., Paragios, N., 2012. Joint tumor segmentation and dense deformable registration of brain mr images, in: Ayache, N., Delingette, H., Golland, P., Mori, K. (Eds.), MICCAI (2), Springer. pp. 651–658.

Pohl, K.M., Fisher, J., Grimson, W.E.L., Kikinis, R., Wells, W.M., 2006. A bayesian model for joint segmentation and registration. NeuroImage 31, 228–239.

Prastawa, M., Bullit, E., Moon, N., Leemput, K., Gerig, G., 2003. Automatic Brain Tumor Segmentation by Subject Specific Modification of Atlas Priors. Academic Radiology 10, 1341–1348.

Prastawa, M., Bullitt, E., Ho, S., Gerig, G., 2004. A brain tumor segmentation framework based on outlier detection. Medical Image Analysis 8, 275–283.

Risholm, P., Samset, E., Talos, I.F., Wells, W., 2009. A non-rigid registration framework that accommodates resection and retraction, in: Information Processing in Medical Imaging, Springer. pp. 447–458.

Rueckert, D., Sonoda, L., Hayes, I., Hill, D., Leach, M., Hawkes, D., 1999. Nonrigid registration using free-form deformations: Application to breast mr images. IEEE Transactions on Medical Imaging 18, 712–721.

Shi, W., Zhuang, X., Pizarro, L., Bai, W., Wang, H., Tung, K.P., Edwards, P.J., Rueckert, D., 2012. Registration using sparse free-form deformations., in: Ayache, N., Delingette, H., Golland, P., Mori, K. (Eds.), Medical Image Computing and Computer-Assisted Intervention–MICCAI, Springer. pp. 659–666.

Soffietti, R., Baumert, B.G., Bello, L., von Deimling, A., Duffau, H., Frénay, M., Grisold, W., Grant, R., Graus, F., Hoang-Xuan, K., Klein, M., Melin, B., Rees, J., Siegal, T., Smits, A., Stupp, R., Wick, W.a., 2010. Guidelines on management of low-grade gliomas: report of an efns-eano* task force. Eur J Neurol 17, 1124–33.

Sotiras, A., Davatzikos, C., Paragios, N., 2013. Deformable medical image registration: A survey. IEEE Transactions on Medical Imaging 32, 1153–1190.

Stefanescu, R., Commowick, O., Malandain, G., Bondiau, P.Y., Ayache, N., Pennec, X., 2004. Non-rigid atlas to subject registration with pathologies for

conformal brain radiotherapy, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI. Springer, pp. 704–711.

Taheri, S., Ong, S., Chong, V., 2010. Level-set segmentation of brain tumors using a threshold-based speed function. Image and Vision Computing 28, 26–37.

Thirion, J.P., 1998. Image matching as a diffusion process: an analogy with maxwell's demons. Medical image analysis 2, 243–260.

Verma, R., Zacharaki, E.I., Ou, Y., Cai, H., Chawla, S., Lee, S.K., Melhem, E.R., Wolf, R., Davatzikos, C., 2008. Multi-parametric tissue characterization of brain neoplasms and their recurrence using pattern classification of mr images. Academic radiology 15, 966.

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.

Wels, M., Carneiro, G., Aplas, A., Huber, M., Hornegger, J., Comaniciu, D., 2008. A discriminative model-constrained graph cuts approach to fully automated pediatric brain tumor segmentation in 3-d mri, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI. Springer, pp. 67–75.

Wyatt, P.P., Noble, J.A., 2003. MAP MRF joint segmentation and registration of medical images. Medical Image Analysis 7, 539–52.

Xuan, X., Liao, Q., 2007. Statistical structure analysis in mri brain tumor segmentation, in: 4th International Conference on Image and Graphics, IEEE. pp. 421–426.

Yezzi, A., Zollei, L., Kapur, T., 2003. A variational framework for integrating segmentation and registration through active contours, in: Medical Image Analysis, pp. 171–185.

Zacharaki, E., Shen, D., Lee, S., Davatzikos, C., 2008. ORBIT: a multiresolution framework for deformable registration of brain tumor images. IEEE Transations on Medical Imaging 27, 1003–1017.

Zhan, Y., Shen, D., 2003. Automated segmentation of 3d us prostate images using statistical texture-based matching method, in: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2003. Springer, pp. 688–696.

Zhang, J., Ma, K.K., Er, M.H., Chong, V., et al., 2004. Tumor segmentation from magnetic resonance imaging by learning via one-class support vector machine, in: International Workshop on Advanced Image Technology (IWAIT'04), pp. 207–211.

Zikic, D., Glocker, B., Konukoglu, E., Criminisi, A., Demiralp, C., Shotton, J., Thomas, O., Das, T., Jena, R., Price, S., 2012. Decision forests for tissue-specific segmentation of high-grade gliomas in multi-channel mr, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI. Springer, pp. 369–376.

Zikic, D., Glocker, B., Kutter, O., Groher, M., Komodakis, N., Kamen, A., Paragios, N., Navab, N., 2010. Linear intensity-based image registration by markov random fields and discrete optimization. Medical image analysis 14, 550–562.

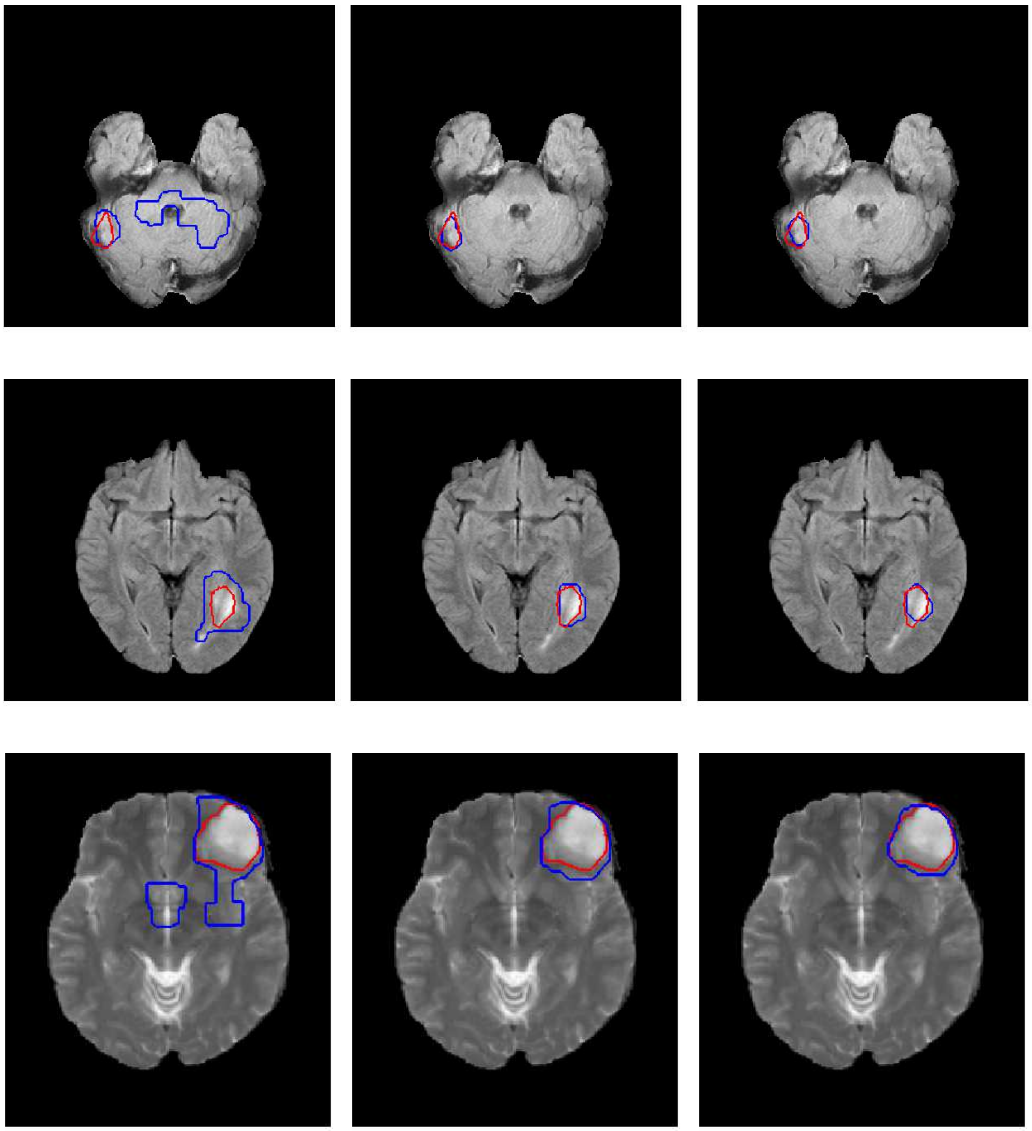

(a) (b) (c)

Figure 8: Visual Segmentation results on the FLAIR database (first two rows) and the BRATS database (bottom row, T2 volume). (a) individual framework, (b) Joint framework, high resolution, (c) Joint framework, with adaptive sampling. Automatic segmentations (blue) are compared to the manual segmentation (red)

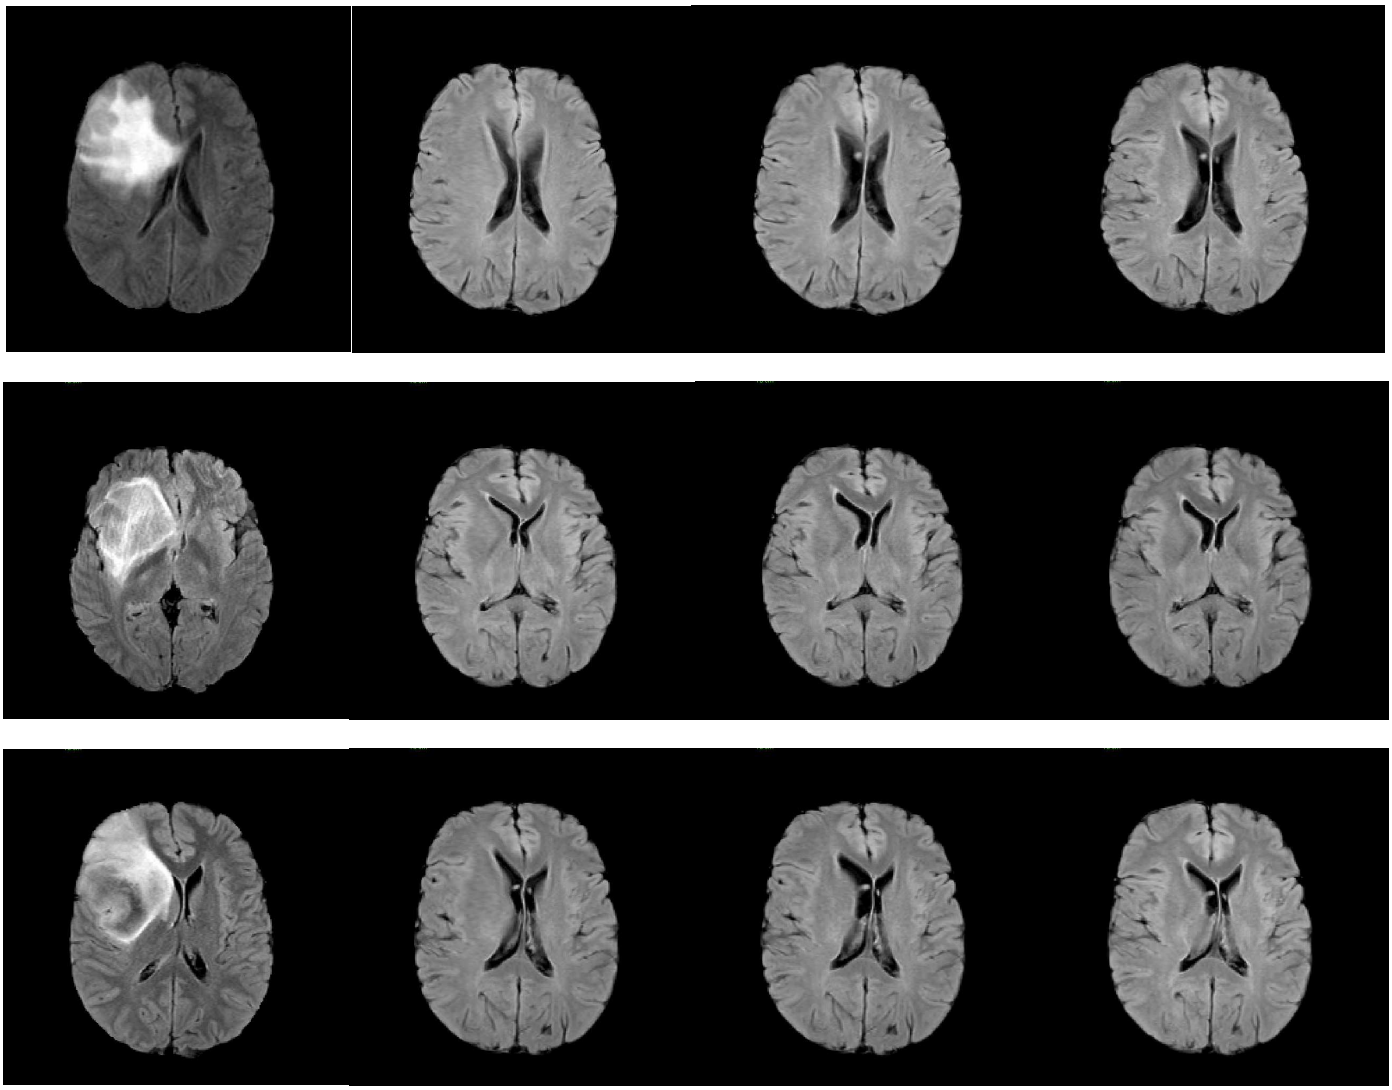

(a) (b) (c) (d)

Figure 9: Visual Registration results, FLAIR dataset. (a) Target image, (b) individual framework, (c) Joint framework, high resolution, (d) Joint framework, with adaptive sampling.

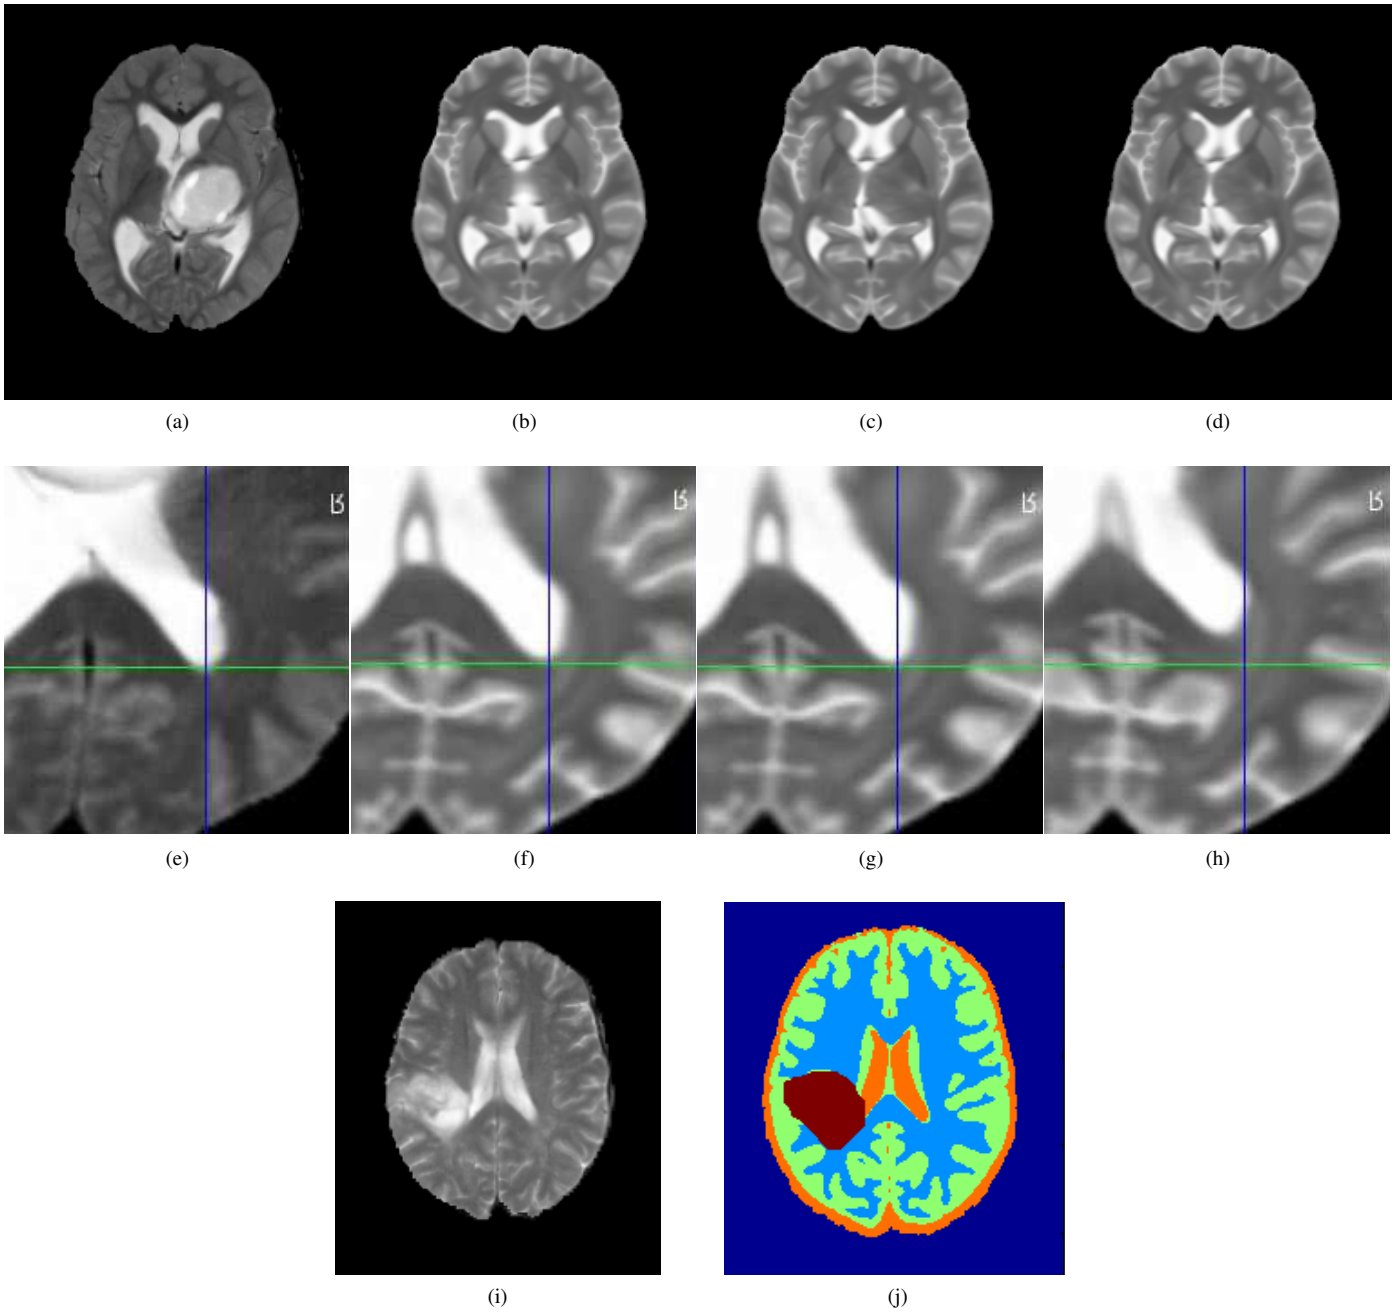

Figure 10: Visual Registration results, BRATS dataset. First row: comparison with the individual registration scheme. (a) Target image, (b) individual framework, (c) Joint framework, high resolution, (d) Joint framework, with adaptive sampling. Second row: close up comparison with low uniform resolution. (e) Target image, (f) Joint framework, high resolution, (g) Joint framework, with adaptive sampling, (h) Joint framework, low uniform resolution. Errors in registration can be seen with respect to the cross lines. Bottom row: example segmentation using the registered MNI-ICBM probabilities. (i) Target image, (j) Segmented image.