Stroke Lesion Segmentation using a Probabilistic Atlas of Cerebral Vascular Territories

Alexandra Derntl[1], Bjoern Menze[2], and Claudia Plant[1]

[1] Helmholtz Zentrum München, TU München, Germany
[2] Institute for Advanced Study and Department of Computer Science, TU München, Germany

Abstract. The accurate segmentation of lesions in magnetic resonance images of stroke patients is important, for example, for comparing the location of the lesion with functional areas and for determining the optimal strategy for patient treatment. Manual labeling of each lesion turns out to be time-intensive and costly, making an automated method desirable. Standard approaches for brain parcellation make use of spatial atlases that represent prior information about the spatial distribution of different tissue types and of anatomical structures of interest. Different from healthy tissue, however, the spatial distribution of a stroke lesion varies considerably, limiting the use of such brain image segmentation approaches for stroke lesion analysis, and for integrating brain parcellation with stroke lesion segmentation.
In this study, we propose to amend the standard atlas-based generative image segmentation model by a spatial atlas of stroke lesion occurrence by making use of information about the vascular territories. As the territories of the major arterial trees often coincide with the location and extensions of large stroke lesions, we use 3D maps of the vascular territories to form patient-specific atlases combined with outlier information from an initial run, following an iterative procedure. We find our approach to perform comparable to (or better than) standard approaches that amend the tissue atlas with a flat lesion prior or that treat lesion as outliers, and to outperform both for large heterogeneous lesions.

1 Introduction

The accurate segmentation of anatomical structures and of lesions that are visible in magnetic resonance image (MRI) of stroke patients has been a somewhat neglected topic in the development of automated brain image segmentation algorithms until very recently [1]. Most algorithms for segmenting structures of the brain in MRI use prior knowledge on the location and the appearance of white matter, gray matter, etc. On the one hand, there are discriminative approaches using, for example, random forests together with local image features [2, 3], that often have a high accuracy, but that can only be applied to images acquired with the exact same MR imaging sequences as the training data. Generative models, on the other hand, describe the intensity distribution in a more informative

and flexible fashion: Seghier et al. [4] proposed a method that constructs a lesion atlas using fuzzy clustering. However, in the clinical workflow often other modalities have to be taken into account. Dalca et al. [5] proposed a method based on the intensity distribution which differentiates between stroke pathologies and leukariosis lesions. However, the stroke segmentation bases only on the intensity model and ignores spatial information according to the cerebral vascular territories [6]. Some approaches consider lesions as clearly distinct outliers of a Gaussian Mixture Model (GMM) whose parameters μ and Σ are optimized by an Expectation-Maximization (EM) algorithm. An early attempt for automated model outlier detection using a GMM was proposed by Van Leemput et al. [7] for multiple sclerosis lesions which are in most cases rather small and therefore more likely to be homogeneous in their intensity. Probabilistic atlases of healthy tissue classes provide a mapping from location to intensity. However, in particular for extensive stroke lesions this is not the case: lesion and healthy tissue intensities might overlap, leading to an improper separation of those classes.

We propose a fully automated method for stroke lesion segmentation in MR images that is using GMMs as a generative model by taking into account both cerebral vascular territories (CVT) [6] and model outlier information [7]. Similar efforts have been undertaken by [8] for brain tumor segmentation which used pre/post T1-weighted contrast images to calculate a patient-specific lesion prior. By contrast, our method does not require such specific modalities, making it more flexible to available data. Further, [9] proposed a latent atlas which is inferred from the given data through an alternating optimization procedure.

In the following we describe the overall model and the resulting iterative approach (Sec. 2), we present experiments (Sec. 3), and offer conclusions (Sec. 4). More specifically, our paper provides as contributions the usage of cerebral vascular territories as additional prior spatial information (Sec. 2.3) for iterative lesion atlas construction (Sec. 2.4).

2 Methods

Our overall approach relies on a two-steps procedure: First, the algorithm setting tries to fit a *robust* GMM of the intensity distribution using healthy tissue atlases. It first identifies lesion candidates according to [7]. Then, these *outliers* provide spatial hints used as a separate lesion atlas which makes them *inliers* in a second EM run where they are an additional component of the GMM in a *standard* non-robust EM segmenter. We therefore consider the method an *outlier-inlier* approach. The calculated outlier-atlases are further enhanced by incorporating contextual knowledge using the information of CVT. Territories are weighted proportionally to the number of lesion candidates found in their spatial region. In each EM iteration, this lesion atlas gets optimized until the best model parameters with respect to the likelihood of the data are found. Further postprocessing is applied in the form of Conditional Random Fields and morphological operators to eliminate false positives resulting from the intensity-based

estimation. In this paper, we will mainly analyze ischemic stroke lesions since they occur most often in practice, accounting for up to 87% of all strokes [10].

2.1 Generative Model for Stroke Lesion Segmentation

We first revisit important concepts from [7, 11] upon which our method is based. Images are given by maps from a finite D-dimensional coordinate space to the intensity space which may be one-dimensional for gray-scale images. We denote MR images as flattened, i.e., 1-dimensional vectors $v = \{v_1, v_2, \dots, v_N\}$ where $I = \{1, \dots, N\}$ is the index set of all voxels and v_i is the intensity (as gray-value) of voxel i. The segmentation is described by labels $c = \{c_1, c_2, \dots, c_N\}$ mapping a voxel to a tissue class, possibly including stroke lesion. For a voxel i, c_i indicates to which tissue type it belongs. C is a finite set of tissue classes, e.g., $\{WM, GM, CSF\}$ or $\{WM, GM, CSF, LES\}$. The latent segmentation l has to be inferred from the observed intensities v. We optimize the model parameters $\Theta = (\mu_c, \Sigma_c)_{c \in C}$ to find the Θ^* that yields the maximal likelihood with respect to the data set v using, e.g. the EM algorithm. Recall that the PDF of a GMM with respect to C is given by:

$$p(x) = \sum_{c \in C} \alpha_c \cdot \frac{1}{(\sqrt{2\pi})^D \sqrt{|\Sigma_c|}} \exp(-\tfrac{1}{2}(x - \mu_c)^T \Sigma_c^{-1}(x - \mu_c)) \tag{1}$$

where $\sum_{c \in C} \alpha_c = 1$, $\forall c \in C$, $\alpha_c \geq 0$. The individual Gaussian PDFs corresponding to $c \in C$ are also referred to as *components* or classes. Each component c consists of a centroid (or mean) μ_c, a covariance Σ_c and a given weight α_c.

2.2 Robust Model Outlier Detection

Standard GMM model parameter optimization assumes that each data point is indeed generated by at least one class. Outliers in the data set are therefore hard to explain by only considering a GMM, in particular, if a point does not seem to fit *any* of the classes. To illustrate this, recall that in GMMs, our goal is to maximize the log-likelihood of the observed data.

$$Q(\Theta) = \sum_{i \in I} \sum_{c \in C} p(c \mid v_i) \log f_c(v_i \mid \Theta), \tag{2}$$

where $p(c \mid v_i)$ denotes the probability of i belonging to c and f_c refers to the PDF of the component c of the GMM spanned by Θ.

Maximizing Θ with respect to (2) is however not robust to significant outliers as those occurring in lesions [7]. In particular, consider a voxel i with intensity v_i that does not fit well to *any* class $c \in C$. Since a probability distribution $p(c \mid v_i)$ has to normalize to 1, the voxel cannot show small probability for *all* classes at the same time. Consequently, the algorithm has to consider very high covariances to include outliers which severely affects the results. Additionally, outliers having small probabilities strongly negatively influence the likelihood of the model in (2). This can be seen since $\log f_c(x \mid \Theta) \to -\infty$ as $f_c(x \mid \Theta) \to 0$.

This problem is alleviated by means of robust statistics [11]. Instead of fitting standard Gaussian PDFs, a *contaminated* variant is proposed where each data point is either generated by a Gaussian $\mathcal{N}(\mu, \Sigma)$ with probability $1 - \varepsilon$ or by a unknown uniform outlier distribution δ with probability ε. The density function converges to a standard Gaussian density by setting $\varepsilon = 0$. The question upon seeing v_i is whether it stems from $\mathcal{N}(\mu, \Sigma)$ or from δ. A perfect classification separates data points v_i into a set G of "good" samples drawn from $\mathcal{N}(\mu, \Sigma)$ (*inliers*) and a set B of "bad" samples (*outliers*) originating from δ.

$$G = \{v_i \text{ generated by } \mathcal{N}(\mu, \Sigma) \mid i \in I\} \tag{3}$$

$$B = \{v_j \text{ not generated by } \mathcal{N}(\mu, \Sigma) \mid j \in I\} \tag{4}$$

Since G and B cannot be perfectly restored from the observed data, a practical classification method is needed. Therefore, we want to classify a data point as outlier if it exceeds a certain distance threshold κ to the distributions spanned by our model. The distance between a data point v_i and the calculated mean μ_c of one class c is estimated by the Mahalanobis-Distance $d_c^2(x_i) = (x_i - \mu_c)\Sigma_c^{(-1)}(x_i - \mu_c)$.

We classify a voxel as outlier if and only if $d_c(x_i) > \kappa$, leading to the sets G_κ and B_κ; the smalller κ the more lesion candidates are detected. Letting $\kappa \to \infty$ results in a standard (outlier-free) EM segmenter [7]:

$$G_\kappa = \{i : \exists c \in C : d_c(v_i) \leq \kappa \mid i \in \{1, \ldots, N\}\} \tag{5}$$

$$B_\kappa = \{j : \forall c \in C : d_c(v_j) > \kappa \mid j \in \{1, \ldots, N\}\} \tag{6}$$

The best values for κ are determined experimentally as done in Section 3.

2.3 Cerebral Vascular Territories

In spite of spatial information given by an atlas, we search for adequate replacements in stroke lesions. It turns out that radiologists use common patterns about cerebral vascular territories that specify the area which is covered by one of the main vessel trees to diagnose strokes since the extensions of large stroke lesions often follow the outlines of the territory the blocked artery is feeding [12].

Figure 4 shows the spatial appearance of the three territories which are covered by one of the three main vessel trees. Following [12], we particularly manually label anterior cerebral artery (ACA), posterior cerebral artery (PCA) and middle cerebral artery (MCA) territories for our evaluation. A recent study on 2213 patients [13] has shown that most of the stroke lesions appear in the MCA territory. Incidentally, large and heterogeneous lesions tend to occur in this particular territory, as we will also discuss in Section 3.

2.4 Construction of the Personalized Lesion Atlas

Based on estimated model outliers and the 3D CVT atlas we construct a new patient-specific lesion prior (see Fig. 3). We assume the set of outliers B_κ to be

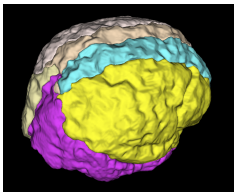

Fig. 1. Dorsal

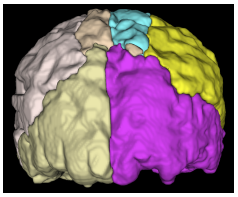

Fig. 2. Lateral

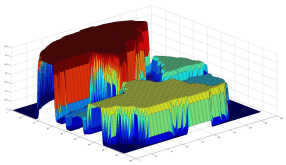

Fig. 3. Reweighting of one slice.

Fig. 4. Cerebral Vascular Territories drawn in 3D with ITKSNAP from a 2D template depicted by [6]. Yellow denotes MCA, turquoise ACA and violet PCA for the right hemisphere. The left hemisphere was labeled equivalently.

determined by the first EM run using only healthy tissue classes. Suppose $I = \{1, \ldots, N\}$ let $v : I \to [0, 1]$ be an image mapping voxel indices to MR intensity values (we write v_i for the intensity at voxel i) and $\mathcal{V}$ be the set of all images. The lesion atlas is a particular image $l \in \mathcal{V}$ where l_i can be interpreted as proportional to the probability of i being part of a lesion. We write $t \in T = \{ACA, MCA, \ldots\}$ for a vascular *territory*. Each vascular territory is characterized by its included voxels $I_t \subseteq I$ (see Fig. 2.4).

For each vascular territory t, we estimate an atlas by setting voxels in $I_t \cap B_\kappa$ to 1 and smoothing this image. Formally, we first obtain images $v^{(t)} : I \to [0, 1]$ by setting $v^{(t)}(i) = 1$ if $i \in I_t \cap B_\kappa$ and 0 otherwise. We smooth this image using a Gaussian filter in 3D. We denote the smoothed image as $\tilde{v}^{(t)} = \mathit{smooth}(v^{(t)})$. Then, we estimate a normalized *voting* coefficient proportional to the probability of the lesion occurring in t by $\mathit{vCoeff}_t = \frac{\sum_{i \in I_t} \tilde{v}^{(t)}(i)}{|B_\kappa|}$.

Finally, we obtain the *voted* territory by $\hat{v}^{(t)}(i) = v^{(t)}(i) \cdot \mathit{vCoeff}_t$. Each territory consequently gets reweighted as can be seen in Fig. 3. The overall lesion atlas used for the second EM-run is finally obtained by the image $\hat{l}(i) \in \mathcal{V}$, defined by $\hat{l}(i) = \prod_{t \in T} \hat{v}^{(t)}(i)$ to achieve a multiplicative, smoothing effect in bordering regions. Care has to be taken if for a voxel i, either of $\hat{v}^{(t)}(i)$ is 0, effectively erasing all other values. We avoid this problem by substituting 0 by 1 in $\hat{l}(i)$ temporarily and replacing these artificially inserted ones by zeros later on. Fig. 2.4 and 2.4 show the main approach we used to construct the lesion prior out of the cerebral vascular territory atlas.[3]

3 Results

We applied our stroke lesion segmentation framework onto 13 different patient datasets with 152 manually annotated ground truth slices (axial, coronal and

[3] An important detail is to label left and right hemispheres individually, according to their perfusion pattern. Otherwise, e.g., the large MCA region (including weak false positives) is weighted disproportionally high.

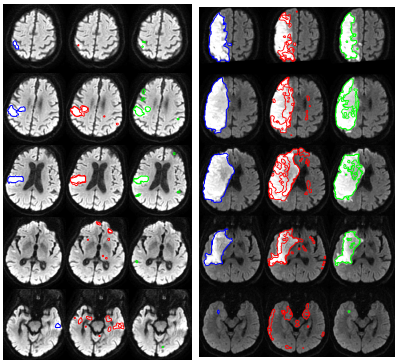

Fig. 5. Qualitative Comparison of the manual ground truth segmentation (blue line), Model Outlier Detection (red line) and our proposed method (green line) presented for FLAIR MRI of patient 1 (left) and 13 (right) using morphological operators as post-processing.

sagittal) including a variety of stroke types and shapes. All datasets where co-registered onto a T1-weighted reference image such that the input images are aligned with the tissue atlases. We drew the CVT altas with ITK-SNAP.[4] To make sure that the atlas is aligned with all the other images (i.e., MRI images and atlases), we drew the vascular territories onto the above-mentioned T1-weighted reference image. We first did a quantitative assessment of the segmentation results by computing the DICE score (also known as F1-score) with different post-processing techniques (morphological operator and Conditional Random Fields) and parameter settings (different flat priors and κ values) as denoted in Figure 6, 7, 8 and 9. It became clear that simply adding an additional class with a flat prior is not competitive to Model Outlier Detection or CVT-Outlier-Inlier (compare Fig. 7). Furthermore, with the best setting we applied a paired student t-test to compare those methods statistically in order to obtain a valid comparison we applied all the methods on the same datasets and used the best configuration we could obtain from them.

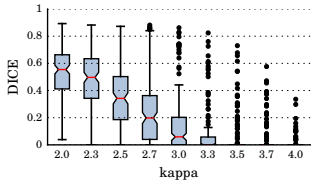

Fig. 6. Model outlier detection

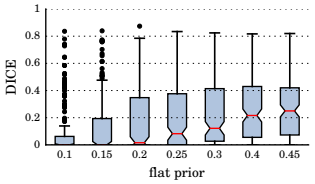

Fig. 7. Lesion class with flat prior.

Even though the Model Outlier Detection works better *on average* over all patients, we could show that our approach performs better on patients with extensive lesions at a significance level of $\alpha = 0.005$. This subgroup of patients was

[4] `http://www.itksnap.org`

selected prior to the evaluation. All results are presented in Table 1. Qualitative results from patient 1 and 13 are shown in Figure 5. Admittedly, the CVT-based approach has difficulties with very small lesions (e.g., patient 5, 8, or 11) where model outlier detection is more robust but it provides much more confidence with large lesions (e.g., patient 13 or 7).

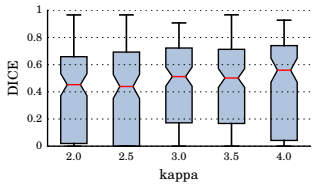

Fig. 8. Outlier-inlier with morph. oper.

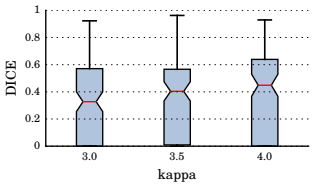

Fig. 9. Outlier-inlier with CRF

Table 1. Patient overview for different segmentation algorithms with different post-processing approaches

Patient Id	MO Open/Close	CVT Open/Close	MO CRF	CVT CRF
Patient 1	0.54	**0.61**	0.55	0.51
Patient 2	0.62	**0.67**	0.72	0.64
Patient 3	0.41	0.37	0.19	0.34
Patient 4	0.47	**0.5**	0.35	0.42
Patient 5	0.39	0.01	0.3	0.0
Patient 6	0.61	**0.62**	0.65	0.4
Patient 7	0.51	**0.72**	0.46	0.62
Patient 8	0.51	0.0	0.1	0.0
Patient 9	0.53	**0.56**	0.08	0.42
Patient 10	0.19	0.06	0.13	0.01
Patient 11	0.57	0.01	0.43	0.0
Patient 12	0.6	0.5	0.4	0.44
Patient 13	0.62	**0.81**	0.42	0.77
Average	0.51	0.42	0.37	0.35
Stdev	0.12	0.29	0.2	0.26

4 Conclusion

We investigated an automated method for stroke lesion segmentation that could prove to be useful, e.g., in the analysis of images acquired in clinical studies. Our method extends previous work in model outlier detection for lesions, first applied to multiple sclerosis patients using a GMM. Using robust statistics, outliers can be detected and classified as lesion. However, for large and heterogeneous lesions this is not enough, as a lesion spans a spectrum of intensity values which can better be captured by a dedicated Gaussian component in the mixture model. Drawing inspiration from the way radiologists perform stroke detection, we incorporated knowledge about cerebral vascular territories that is combined with outlier information to form a lesion atlas. This lesion atlas is then reweighted

proportionally to the incurred outliers of each territory in each iteration until maximum likelihood model parameters are found. Several approaches to construct this atlas were examined and compared to a flat prior as a baseline.

Our evaluation showed that the outlier-inlier approach on average performs comparable to the model outlier detection for an overall set of 13 patients and significantly better than a uniform prior for a lesion class. The performance was enhanced by the postprocessing methods: conditional random fields and morphological operators. Considering large stroke lesion patients alone, our method dominates the other approaches evaluated in this paper.

In future work, we consider additional features and disease patterns to improve clustering. One way would be to enhance the EM optimizing a GMM (or a similar clustering method) with a so-called minimal description length (MDL). MDL is a well-known information theoretical concept that leads to reduced overfitting. Based on this criterion, one could algorithmically decide whether to use model outlier detection or the outlier-inlier approach. The right parameter setting for κ can be done automatically with respect to MDL.

References

1. Rekik, I., Allassonnire, S., Carpenter, T.K., Wardlaw, J.M.: Medical image analysis methods in mr/ct-imaged acute-subacute ischemic stroke lesion: Segmentation, prediction and insights into dynamic evolution simulation models. a critical appraisal. NeuroImage: Clinical **1** (2012) 164 – 178
2. Maier, O., Wilms, M., von der Gablentz, J., Krmer, U.M., Mnte, T.F., Handels, H.: Extra tree forests for sub-acute ischemic stroke lesion segmentation in mri sequences. Journal of Neuroscience Methods **240** (2015) 89 – 100
3. Mitra, J., Bourgeat, P., Fripp, J., Ghose, S., Rose, S., Salvado, O., Connelly, A., Campbell, B., Palmer, S., Sharma, G., Christensen, S., Carey, L.: Lesion segmentation from multimodal mri using random forest following ischemic stroke. NeuroImage **98** (2014) 324 – 335
4. Seghier, M.L., Ramlackhansingh, A., Crinion, J., Leff, A.P., Price, C.J.: Lesion identification using unified segmentation-normalisation models and fuzzy clustering. NeuroImage **41** (2008) 1253 – 1266
5. Dalca, A., Sridharan, R., Cloonan, L., Fitzpatrick, K., Kanakis, A., Furie, K., Rosand, J., Wu, O., Sabuncu, M., Rost, N., Golland, P.: Segmentation of cerebrovascular pathologies in stroke patients with spatial and shape priors. **8674** (2014) 773–780
6. Savoiardo, M.: The vascular territories of the carotid and vertebrobasilar systems. diagrams based on CT studies of infarcts. Italian journal of neurological sciences **7** (1986) 405–409
7. Leemput, K.V., Maes, F., Vandermeulen, D., Colchester, A.C.F., Suetens, P.: Automated segmentation of multiple sclerosis lesions by model outlier detection. IEEE Trans. Med. Imaging **20** (2001) 677–688
8. Moon, N., Bullitt, E., Van Leemput, K., Gerig, G.: Automatic brain and tumor segmentation. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2002. Springer (2002) 372–379
9. Riklin-Raviv, T., Van Leemput, K., Menze, B.H., Wells, W.M., Golland, P.: Segmentation of image ensembles via latent atlases. Medical image analysis **14** (2010) 654–665

10. Association, A.S.: Ischemic strokes (clots) (2015) [Online; URL: http://www.strokeassociation.org/STROKEORG/AboutStroke/TypesofStroke/IschemicClots/Ischemic-Strokes-Clots_UCM_310939_Article.jsp; accessed 28-March-2015].
11. Zhuang, S., Huang, Y., Palaniappan, K., Zhao, Y.: Gaussian mixture density modeling, decomposition, and applications. Image Processing, IEEE Transactions on **5** (1996) 1293–1302
12. Damasio, H.: A computed tomographic guide to the identification of cerebral vascular territories. Archives of Neurology **40** (1983) 138–142
13. Ng, Y.S., Stein, J., Ning, M., Black-Schaffer, R.M.: Comparison of clinical characteristics and functional outcomes of ischemic stroke in different vascular territories. Stroke **38** (2007) 2309–2314