

Ischemic Stroke Lesion Segmentation

www.isles-challenge.org

Proceedings
5th October 2015
Munich, Germany

Preface

Stroke is the second most frequent cause of death and a major cause of disability in industrial countries. In patients who survive, stroke is generally associated with high socioeconomic costs due to persistent disability. Its most frequent manifestation is the ischemic stroke, whose diagnosis often involves the acquisition of brain magnetic resonance (MR) scans to assess the stroke lesion's presence, location, extent, evolution and other factors. An automated method to locate, segment and quantify the lesion area would support clinicians and researchers alike, rendering their findings more robust and reproducible.

New methods for stroke segmentation are regularly proposed. But, more often than desirable, it is difficult to compare their fitness, as the reported results are obtained on private datasets. Challenges aim to overcome these shortcomings by providing (1) a public dataset that reflects the diversity of the problem and (2) a platform for a fair and direct comparison of methods with suitable evaluation measures. Thus, the scientific progress is promoted.

With ISLES, we provide such a challenge covering ischemic stroke lesion segmentation in multi-spectral MRI data. The task is backed by a well established clinical and research motivation and a large number of already existing methods. Each team may participate in either one or both of two sub-tasks:

SISS Automatic segmentation of ischemic stroke lesion volumes from multi-spectral MRI sequences acquired in the sub-acute stroke development stage.

SPES Automatic segmentation of acute ischemic stroke lesion volumes from multi-spectral MRI sequences for stroke outcome prediction.

The participants downloaded a set of training cases with associated expert segmentations of the stroke lesions to train and evaluate their approach, then submitted a short paper describing their method. After reviewing by the organizers, a total of 17 articles were accepted and compiled into this volume. At the day of the challenge, each teams' results as obtained on an independent test set of cases will be revealed and a ranking of methods established.

For the final ranking and more information, visit WWW.ISLES-CHALLENGE.ORG.

Oskar Maier, Universität zu Lübeck

Mauricio Reyes, University of Bern

Björn Menze, TU Munich

August 2015

Organizers

Oskar Maier, Universität zu Lübeck, Germany

Mauricio Reyes, University of Bern, Switzerland

Björn Menze, TU Munich, Germany

Sponsoring Institutions

Institute of Medical Informatics, Universität zu Lübeck, Germany

Institute for Surgical Technology & Biomechanics, University of Bern, Switzerland

Computer Science, TU Munich, Germany

ISLES Challenge 2015: Automated Model-Based Segmentation of Ischemic Stroke in MR Images

Tom Haeck^{1,2}, Frederik Maes^{1,2}, and Paul Suetens^{1,2}

¹ KU Leuven, Dept. of Electrical Engineering, ESAT/PSI

² UZ Leuven, Medical Imaging Research Center

Abstract. We present a novel fully-automated generative ischemic stroke lesion segmentation method that can be applied to individual patient images without need for a training data set. An Expectation Maximization-approach is used for estimating intensity models for both normal and pathological tissue. The segmentation is represented by a level-set that is iteratively updated to label voxels as either normal or pathological, based on which intensity model explains the voxels' intensity the best. A convex level-set formulation is adopted, that eliminates the need for manual initialization of the the level-set. The performance of the method for segmenting the ischemic stroke is summarized by an average Dice score of 0.78 and 0.51 for the SPES and SISS 2015 training set respectively.

1 Introduction

The MICCAI Ischemic Stroke Lesion Segmentation (ISLES) challenge comprises the automatic segmentation of ischemic stroke lesions acquired in the sub-acute stroke development stage (SISS) and automatic segmentation of acute ischemic stroke lesions for stroke outcome prediction (SPES).

Discriminative segmentation methods require a set of manually annotated training images from which the appearance of the brain structures of interest is implicitly learned by the algorithm. Generative models on the other hand do not require a set of annotated training images. Explicit prior knowledge of anatomy or intensity appearance is directly incorporated into the algorithm [1]. In clinical practice the availability of annotated training data may be limited or non-existent, such that a generative method that does not rely on training data may be preferred. We present a novel fully-automated generative ischemic stroke segmentation method that only makes use of a probabilistic atlas of white matter (WM), grey matter (GM) and cerebrospinal fluid (CSF) and for which no manual initialization is needed. The probabilistic prior guides the global search for voxel outliers that cannot be explained by the normal tissue model. The lesion boundary is represented as a level-set that spatially regularizes the segmentation.

2 Method

Classification is based on an Expectation Maximization (EM)-estimation of normal and pathological intensity models. An evolving level-set determines which of both intensity models applies to what regions in the image (Fig. 1).

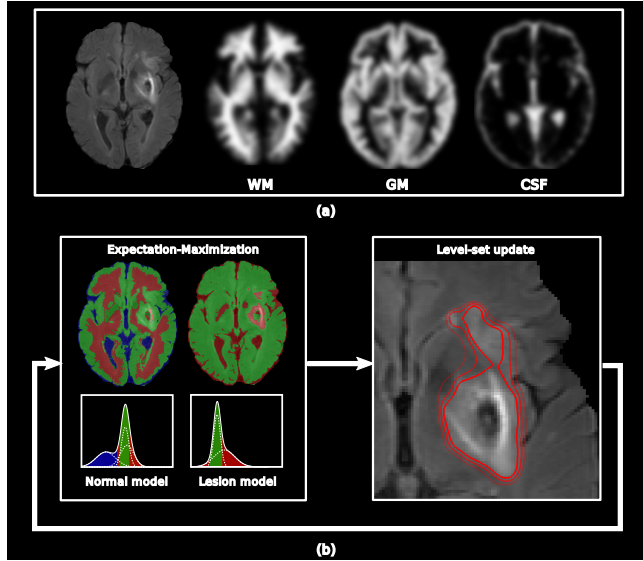


Fig. 1. (a) Spatial priors are non-rigidly registered to the patient image. (b) A full EM-estimation of the normal and pathological intensity models is done, after which a level-set is updated. This process is repeated until convergence.

Prior Registration Spatial priors of WM, GM and CSF are non-rigidly registered to the patient image. Although registration of a healthy atlas to a patient image is still an active field of research, this problem is ignored for now and standard non-rigid registration methods are used. The prior information is relaxed by smoothing the spatial priors with a Gaussian kernel.

Intensity models and the Expectation-Maximization algorithm Normal and pathological tissue intensities are modeled separately. Let G_{Σ_j} be a zero-mean multivariate Gaussian with covariance matrix Σ_j , then normal and pathological tissue are both modeled by a Gaussian mixture model

$$p(\mathbf{y}_i|\theta) = \sum_j^K G_{\Sigma_j}(\mathbf{y}_i - \mu_j)p(\Gamma_i = j), \quad (1)$$

with $\mathbf{y}_i = (y_{i_1}, \dots, y_{i_N})$ the intensity of voxel i and $\Gamma_i = \{j|j = 1 \dots K\}$ the tissue class. The intensity model parameters $\theta = \{(\mu_j, \Sigma_j)|j \in 1 \dots K\}$ are iteratively updated using an EM-approach [1]. For normal tissue, $K = 3$ and $p(\Gamma = j) = \pi_j$ are the spatial priors for WM, GM and CSF.

Convex level-set formulation The image I is subdivided into regions labeled Ω_{in} (pathological tissue) and Ω_{out} (normal tissue) for which the intensities are modeled by the probability distributions described in the previous paragraph

[2]. The regions are separated by a boundary $\partial\Omega$ that is implicitly represented by a level-set function. The boundary and intensity model parameters are found by minimizing the energy functional

$$\underset{\theta_{in}, \theta_{out}, \partial\Omega}{\operatorname{argmin}} \quad \lambda_1 \int_{\Omega_{in}} -\log p_{in}(I|\Omega_{in}, \theta_{in}) d\mathbf{x} + \lambda_2 \int_{\Omega_{out}} -\log p_{out}(I|\Omega_{out}, \theta_{out}) d\mathbf{x} + \kappa L(\partial\Omega), \quad (2)$$

where $L(\cdot)$ is the length. The first two terms penalize the negative loglikelihood of the image I evaluated in respectively the pathological and normal region. The third term penalizes the length of the boundary. Parameters λ_1 , λ_2 and κ determine the relative importance of the energy terms. For each iteration to update the level-set, a full EM-estimation of the parameters θ_{in} and θ_{out} is done.

The energy functional is non-convex and the gradient flow finds a solution that depends on a manual initialization of the level-set. This initialization typically has significant impact on the segmentation result. In this work, this problem is overcome by using a convex level-set formulation that performs a global search over the image and makes a manual initialization superfluous. A global minimum is guaranteed by replacing the gradient flow by another gradient flow with the same steady-state solution and by restricting the level-set to lie in a finite interval [3]. The problem is thus reformulated as an L_1 -minimization problem that is solved by the Split Bregman-numerical scheme [3]. It is important to note that, by using spatial priors of WM, GM and CSF, the global optimum coincides with the clinically meaningful notion of normal and pathological regions.

3 Experiments and Results

The SPES and SISS training data are already skull-stripped and registered intrapatient. No further pre-processing is done. Prior registration is based on the T1-weighted MNI-Colin27 atlas (2008) that is registered to the patient volume with a cross-correlation similarity measure (radius 4 voxels) by the Advanced Normalization Tools (ANTs) toolbox [4]. The spatial priors are relaxed by a Gaussian kernel with $\sigma = 3$ voxels. For segmentation of the SPES data, we use the T2-weighted and TTP-weighted MR images and for SISS the diffusion weighted and FLAIR-weighted MR images. For SPES, the modalities are used in a completely multivariate way, i.e. with bivariate Gaussian models. For SISS, the modalities are segmented separately and a voxel is only labeled as lesion if it is a lesion in both modalities. The number of Gaussians for modeling the lesion intensities is set to 1. The energy functional hyperparameters are $\lambda_1 = \lambda_2 = 1e1$ and $\kappa = 1e1$. Performance of the algorithm for both SPES and SISS is evaluated by means of the ASSD, Dice overlap coefficient, Hausdorff distance and precision and recall (Table 1). The median Dice scores for the SPES and SISS training sets are 0.79 and 0.60 respectively (Fig. 2).

Table 1. Performance of the presented method on the SPES and SISS training set

	ASSD		Dice		Hausdorff		Precision		Recall	
	avg	std	avg	std	avg	std	avg	std	avg	std
SPES	3.51	2.13	0.78	0.08	46.31	25.17	0.78	0.11	0.80	0.12
SISS	14.43	25.88	0.53	0.26	69.67	30.77	0.62	0.31	0.56	0.29




Fig. 2. **Left:** Boxplots for the SPES and SISS Dice scores. **Right:** T2- and TTP-weighted MR example image from SPES and FLAIR- and diffusion weighted MR example image from SISS with ground truth segmentations (red) and the resulting segmentations (green) for a typical segmentation (Dice score 0.79 and 0.50 for SPES and SISS).

4 Discussion and Conclusion

In plenty of clinical settings only a handful of patient images needs to be processed without the availability of an annotated training set. Generative methods have therefore an enormous practical value. We have presented a generative method for segmenting the ischemic stroke lesion in the SPES and SISS training set. The method is abundantly flexible to detect any intensity abnormality, and therefore also suitable to detect other lesions like tumor or MS.

References

1. K. Van Leemput, F. Maes, D. Vandermeulen, and P. Suetens. Automated Model-Based Tissue Classification of MR Images of the Brain. *IEEE Transactions on Medical Imaging*, 18:897–908, 1999.
2. M. Rousson and R. Deriche. A variational framework for active and adaptive segmentation of vector valued images. In *Proceedings of the Workshop on Motion and Video Computing, MOTION '02*. IEEE Computer Society, 2002.
3. T. Goldstein, X. Bresson, and S. Osher. Geometric applications of the split bregman method: Segmentation and surface reconstruction. *Journal of Scientific Computing*, 45(1-3):272–293, 2010.
4. B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee. A reproducible evaluation of ants similarity metric performance in brain image registration. *Neuroimage*, 54(3):2033–2044, Feb 2011.