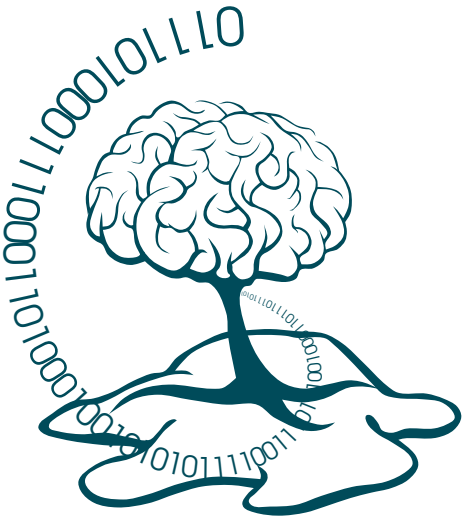

Ischemic Stroke Lesion Segmentation

www.isles-challenge.org

Proceedings

5th October 2015

Munich, Germany

UNIVERSITÄT ZU LÜBECK

Technische Universität München

UNIVERSITÄT BERN

MICCAI'15 MUNICH

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

Version: 2015-09-30

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

Version: 2015-09-30

Hierarchical Segmentation of Normal and Lesional Structures Combining an Ensemble of Probabilistic Local Classifiers and Regional Random Forest Classification

Andrew Jesson and Tal Arbel

Centre for Intelligent Machines, McGill University

1 Overview

We present a hierarchical framework for the simultaneous segmentation of normal and pathological structures in brain MRI. The framework starts with the ensemble decision made by a number of probabilistic local classifiers distributed thoughout a fixed reference space (EPLC). The EPLC provides consistent smooth segmentations for both normal and pathological structures which are then grouped into regions and passed on to a high-level regional random forest classifier (RRF).

2 Implementation

The class of a given voxel in medical image segmentation problems depends on its location within the image reference space, its intensity value, neighborhood context in both the intensity and class label domain, and regional context between continuously labeled structures. Here we model location by an esemble of local classifiers populated throughout a common reference space. Each local classifier models the class posterior probability of the class given voxel and neighbourhood context in the intensity domain by taking a filtered patch of the image as input. Neighbourhood context in the label domain is modeled by a global Markov Random Field (MRF). Finally, regional context is captured by a random forest classifier.

2.1 Ensemble of Probabilistic Local Classifiers

Reference space. The EPLC requires a common reference space and so we use the reference space defined by the MNI Linear ICBM Average Brain Stereotaxic Registration Model [2,3]. Each subject is normalized into this space using the rigid transform determined by the antsRegistration tool [4] with subject/template T1 image pair as input. The spatial centres of each local classifier are distributed throughout this space in a hexagonal close-packed lattice and are given by

$$\mathbf{x}_c = \begin{bmatrix} x_c \\ y_c \\ z_c \end{bmatrix} = \begin{bmatrix} 2i + ((j+k) \mod 2) \\ \sqrt{3}\left(j + \frac{1}{3}(k \mod 2)\right) \\ \frac{2\sqrt{6}}{3}k \end{bmatrix} r, \tag{1}$$

where i, j, k are the voxel indices and r is chosen to be 8.5mm.

Local classifiers. In total there are $N = 554$ local classifiers as we retain only those with spatial centres that overlap with the reference space model's brain mask. Each local classifier has a radially decaying spatial responsibility in the reference space given by

$$w_n(\mathbf{x}) = \frac{1}{Z} \exp\left(-c\frac{(\mathbf{x} - \mathbf{x}_c)^2}{r^2}\right), \tag{2}$$

where c is a constant controlling the decay of the spatial responsibility chosen here to be 0.693 and Z is a normalization constant defined by the sum of all spatial responsibilities at $\mathbf{x}$.

Each classifier builds distributions for $K = 11$ classes in total, namely background, cerebrospinal fluid, lateral ventricles, other ventricles, deep gray matter, cortical gray matter, cerebellar gray matter, cerebral white matter, cerebellar white matter, brain stem, and lesion. Given that ground truth is only available for lesion, we generate a set of atlases from the training data using the Multi-Atlas Label Fusion approach that was developed in previous work in the context of MS lesion segmentation to get training samples for normal structures [1].

To model neighbourhood context in the intensity domain each classifier takes a 15x7x3x3 patch centred around a given voxel as input. The first three dimensions of the patch are spatial and the fourth dimension is comprised of the T1, FLAIR, and DWI contrasts. To spatially decorrelate the input patch and reduce dimensionality we filter the patch with principal component analysis (PCA) determined kernels. Mixtures of gaussians (GMM) then model the distribution of each class given the filtered input.

For each test subject, 5000 samples are drawn randomly, with replacement, and weighted by $w_n(\mathbf{x})$. The patches are reshaped to 1x945 observation vectors $\mathbf{d}(\mathbf{x})$ to form a 5000x945 observation matrix D and each column of D is standardized. GMM inputs are given by $\hat{\mathbf{d}}_n(\mathbf{x}) = \mathbf{d}(\mathbf{x})U_n$, where U_n is a 945xp matrix of principal components determined during training. U_n is unique for each local classifier and p is determined by taking the principal components with largest explained variance such that the total retained variance is $\sim 90\%$ of the total variance from the training set.

The probability density functions $p_n(\hat{\mathbf{d}}_n(\mathbf{x}) \mid C_k)$ are estimated during training using a GMM for each class C_k and the number of components for each model are determined iteratively using Bayesian information criterion (BIC). The class prior probabilities $p_n(C_k)$ are estimated as the relative frequency of class C_k sampled by the local classifier during training. The posterior probability of observing class C_k given $\hat{\mathbf{d}}_n(\mathbf{x})$ by the nth classifier is given by

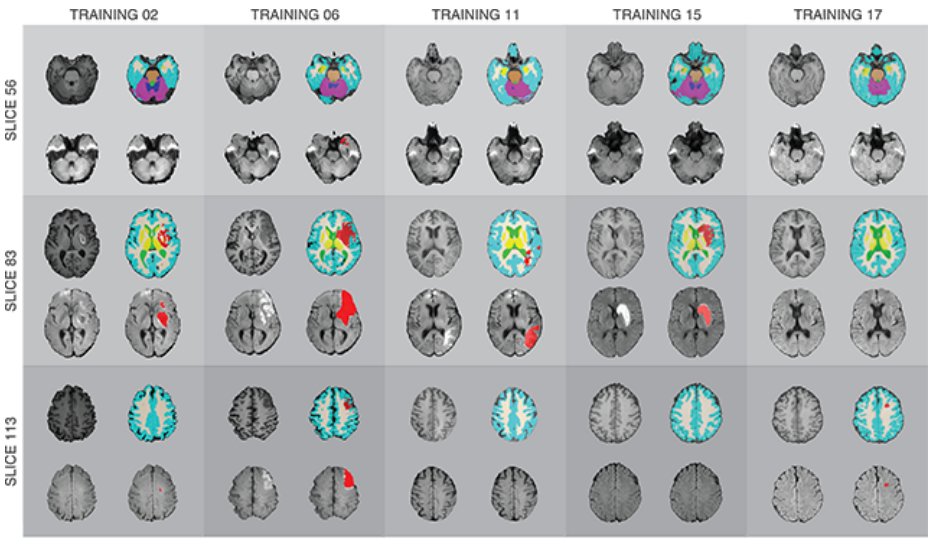

Fig. 1. Example segmentations after ensemble decision. Each square contains an axial slice from a given subject: T1, top left; Segmentation, top right; DWI, bottom left; Ground truth lesion, bottom right. Lesion, red; CSF, green; deep gray matter, yellow; cortical gray matter, cyan; cerebellar gray matter, purple; cerebral white matter, white; cerebellar white matter, blue; brain stem, beige. Results are based on 7-fold cross validation on the training data.

$$p_n\left(C_k \mid \hat{\mathbf{d}}(\mathbf{x})\right) \propto p_n\left(\hat{\mathbf{d}}_n(\mathbf{x}) \mid C_k\right) p_n\left(C_k\right). \tag{3}$$

The ensemble confidence for a given class at voxel $\mathbf{x}$ is then given by

$$f(w_n(\mathbf{x}), \hat{\mathbf{d}}_n(\mathbf{x}), C_k) = \sum_{n=1\ldots N} w_n(\mathbf{x}) p_n\left(C_k \mid \hat{\mathbf{d}}_n(\mathbf{x})\right). \tag{4}$$

Global MAP-MRF. Given that samples for each local classifier are drawn randomly with replacement there is no guarantee that all voxels will be visited by the EPLC. We use a MRF solution to yield a smooth labelling. The prior energy is given by a Potts model with $\beta = 0.1$ and the observation energy is given by $-\log(f(w_n(\mathbf{x}), \hat{\mathbf{d}}_n(\mathbf{x}), C_k))$. The optimal labelling solution is found using ICM. Figure 1 shows example segmentations from several subjects.

Utilizing probabilistic outliers. As can be seen in figure 1 the healthy tissue segmentations are qualitatively consistent and acceptable; however, lesions appear consistently under segmented. Fortunately, by using a probabilistic model for each local classifier we can obtain a quantitative outlier measure using the Mahalanobis distance (computing CDF for each high dimension GMM is prohibitively time consuming). An outlier mask is generated for each subject by

thresholding the Mahalanobis distance corresponding to the maximum posterior class value by 2 standard deviations above the global mean measure. This mask is then added to the proposed lesion segmentation.

2.2 Regional Random Forest Classifier

Lesion segmentations are finally refined using a random forest classifier. Candidate lesions are defined by morphological 18-connected regions. Features for the random forest are the distance minimum, maximum, mean and variance from each normal tissue label excluding background; the volume, and solidity of the candidate lesion; the convex hull inertial tensor and principal moments of the candidate lesion; and 32 bin histograms from normalized T1, FLAIR, and DWI contrasts over the candidate lesion. In total there are 146 features and each candidate split in the forest randomly chooses 12 of these features. We use MATLAB's 'treebagger' class to implement the random forest. Candidate lesions for which the confidence of the random forest is greater than around 40% are retained for the final classification. Figure 2 shows the effect of the high level refinement on the lesion segmentations provided by the ensemble of local classifiers.

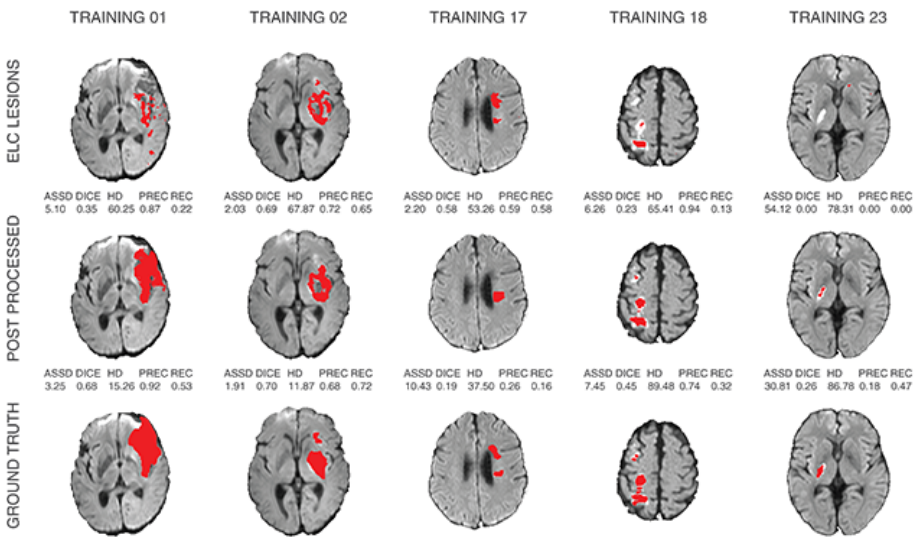

Fig. 2. Example lesions segmentations. From top to bottom: EPLC output, context based classifier output, ground truth segmentation. From left to right: different subjects in training dataset. Below rows 1 and 2 are the computed challenge metrics for the subject. Results are based on 7-fold cross validation on the training data.

References

1. Jesson, A., Arbel, T.: Hierachical MRF and Random Forest Segmentation of MS Lesions and Healthy Tissues in Brain MRI, presented at the Longitudinal Multiple Sclerosis Lesion Segmentation Challenge held in conjunction with ISBI 2015, Brooklyn, New York, April, 16, 2015. (http://www.iacl.ece.jhu.edu/w/images/7/72/Andrew_Jesson.pdf)
2. Foynov, V., Evans, A., et al: Unbiased Average Age-appropriate Atlases for Pediatric Studies. NeuroImage, vol. 54, no. 1. (2011)
3. Collins, L., Zijdenbos, A., et al: ANIMAL+INSECT: Improved Cortical Structure Segmentation. IPMI LNCS, vol. 1613/1999, 210223 (1999)
4. Avants, B., Epstein, C., et al: Symmetric Diffeomorphic Image Registration with Cross-Correlation: Evaluating Automated Labeling of Elderly and Neurodegenerative Brain. Medical Image Analysis, vol. 12, no. 1, 2641 (2008)