



# Ischemic Stroke Lesion Segmentation

[www.isles-challenge.org](http://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](http://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

# Segmenting the ischemic penumbra: a spatial Random Forest approach with automatic threshold finding

Richard McKinley<sup>1</sup>, Levin Häni<sup>1</sup>, Roland Wiest<sup>1</sup>, Mauricio Reyes<sup>2</sup>

<sup>1</sup> Department of Diagnostic and Interventional Neuroradiology, Inselspital, Bern University Hospital, Switzerland

<sup>2</sup> Institute for Surgical Technology and Biomechanics, University of Bern, Switzerland

**Abstract.** We propose a fully automatic method for segmenting the ischemic penumbra, using image texture and spatial features and a modified Random Forest algorithm, which we call Segmentation Forests, which has been designed to adapt the original Random Forests algorithm of Breiman to the segmentation of medical images. The method is fast, taking approximately six minutes to segment a new case, and has yields convincing results (An out-of-sample average Dice coefficient of 0.85, with a standard deviation of 0.06).

## 1 Introduction

In patients presenting with acute stroke, it is important to be able to quickly identify hypoperfused tissue-at-risk, in order to assess the suitability of intra-arterial therapy. Thresholding maps derived from perfusion-weighted imaging provides a usable but crude assessment of this volume of tissue: the technique is prone to artifacts in the processing of the perfusion maps, leading to, for example, identification of tissue as at-risk on the contralateral side of the brain. Fast, automatic methods for identifying the tissue at risk which improve on thresholding are therefore needed.

In this paper we introduce such a method, based on a modification of the standard Random Forest approach [1].

Random Forests are a popular machine learning algorithm in medical imaging applications, but standard implementations of the algorithm are not optimal for medical imaging data, which is a) highly correlated at the patient level, and b) unbalanced, with the target class often having a prevalence of 1% or less. Segmentation forests avoid these problems by bootstrapping training data first at the patient level, and the second by using out-of-sample patients to empirically discover a threshold at which the Dice coefficient of the segmentation is maximized, avoiding the need for holding out training data to tune the classifier. Preliminary results of applying this technique to the ISLES acute stroke dataset are reported.

## 2 Method

Our segmentation algorithm uses pre-processing and texture and spatial features inspired by the features used in the BraTumIA brain tumour segmentation tool [2] and a previous pilot study on stroke segmentation [3]

### 2.1 Standardization and feature processing

Prior to model construction, features were extracted from the multimodal imaging volumes using the Insight Segmentation and Registration Toolik, available from [itk.org](http://itk.org).

Before feature extraction, the structural image modalities (T2 and T1 contrast) are smoothed (using the GradientAnisotropicDiffusionImageFilter from ITK version 7.4.2) and a window filter is applied to the TMax map (with minimum value 0 and maximum value 100) to suppress abnormally high values. All feature maps were then rescaled with ITK version 7.4.2 to lie within the range [0,256]. The T1c image is coregistered to an atlas to allow extraction of atlas coordinates and the location of the mid-sagittal plane.

We then extract, for each voxel of the volume, and each image modality, a feature vector, consisting of the following features: - local texture features, extracted over both 3-by-3-by-3 and 5-by-5-by-5 voxel volumes - mean intensity - intensity variance, skewness and kurtosis - signal to noise ratio, entropy and energy - local intensity percentiles - local image gradient features (gradient magnitude computed using GradientMagnitudeRecursiveGaussianImageFilter from ITK version 7.4.2, with a sigma of 1.0) - point intensity of the gradient magnitude - mean of the gradient magnitude over 3-by-3-by-3 and 5-by-5-by-5 volumes - variance of the gradient magnitude over 3-by-3-by-3 and 5-by-5-by-5 volumes - a symmetry feature computed using a corresponding point on the contralateral side of the brain (found using the previously computed atlas coordinates): the difference between the voxel intensity and a smoothed intensity (computed using a SmoothGaussFilterType from ITK version 7.4.2) from the contralateral side.

Additional features were the unscaled image modalities, atlas coordinates, and an indication of whether the voxel is on the ipsi- or contralesional (inferred by comparing the means of the scaled TMax on each side of the brain.)

### 2.2 The Segmentation Forest classifier

Random forests have been successfully used in numerous medical imaging applications, either alone or together with a conditional random field regularization. Random Forests is an example of bagging, in which a number of weak classification algorithms are trained, each on a random sample of the training data: their outputs are then averaged, yielding a better classification than each individual

classifier. In the case of Random Forests, the weak classifiers are decision trees. Each tree is built on a different bootstrap sample of the training data, and at each stage of building the tree only a subset of the features are available for classification: this prevents the trees from being too closely correlated, which would spoil the benefits of ensembling the weak learners.

Bagging is less helpful when, as is the case in medical imaging, training data is stratified into correlated clusters: the training examples extracted from a single patient’s imaging data are closely correlated to one another, meaning that bagging fails to decorrelate the weak learners. In addition, the validity of out-of-bag measures of performance is diminished (since the out-of-bag samples are closely correlated to the training data, the performance on those samples will in general be much higher than on new data). To mitigate this, we introduce a classifier we call Segmentation Forests. The algorithm works by training a number of small Random forest classifiers (for example, with 50 trees), each one trained on the data given by a bootstrap sample of the training cases. The final classifier is then formed by averaging the output of each individual random forest, meaning that the final classifier is, in the end, still an ensemble of trees, as with an ordinary random forest.

Given a new case, the output of this classifier is then a score, between 0 and 1, for each voxel. To generate a segmentation from this map, we set a threshold on this map. Rather than choosing a threshold of 0.5, or calculating a theoretically correct threshold based on the relative incidence of the background and lesion classes, we instead derive an empirical threshold during training. We apply each small forest built on a bootstrap sample of the training cases to the patient cases not selected by the bootstrap sample. We calculate an estimate of the threshold which optimizes the Dice coefficient over these out-of-bag cases, and the average of these estimates is then used as the threshold for the final classifier.

### 3 Preliminary Results and Discussion

The segmentation forest classifier was implemented using the SpeedRF Random Forest implementation of the H2O machine learning package (Version 2.8.4.4), and the accompanying R package, both acquired from on CRAN or via [h2o.ai](https://h2o.ai). This implementation of Random Forests is faster and more memory-efficient than the standard R implementation, allowing for the use of all data in the training sets, without downsampling of the background class.

The initial results were generated using a segmentation forest setup, with ten bootstrap samples of the training data, each of which consisted of 50 trees, with an mtries parameter of 30 and a maximum depth of 40.

We used the out-of-bag results to generate segmentations of each test case, with an average Dice score of 0.85 (sd 0.06). The mean average symmetric surface distance (ASSD) of our method was 1.42 (sd 1.01), and the mean Hausdorff distance was

30.71 (sd 18.91). Currently our method does not involve any post-processing to remove isolated outlier voxels: we could expect the Hausdorff distance in particular to benefit from such post-processing. Time to run the algorithm on a single case, including feature extraction, was six minutes.

Further work to improve the classifier will consist in optimizing, using the out-of-bag error, the parameters of the individual random forest models, and incorporating post-processing to improve the Hausdorff distance.

## References

- [1] L. Breiman, “Random forests,” *Mach. Learn.*, pp. 5–32, 2001.
- [2] N. Porz, S. Bauer, A. Pica, P. Schucht, J. Beck, R. K. Verma, J. Slotboom, M. Reyes, and R. Wiest, “Multi-Modal Glioblastoma Segmentation: Man versus Machine,” *PLoS One*, vol. 9, no. 5, p. e96873, May 2014.
- [3] S. Bauer, P. P. Gratz, J. Gralla, M. Reyes, and R. Wiest, “Towards automatic MRI volumetry for treatment selection in acute ischemic stroke patients.” *Conf. Proc. ... Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. IEEE Eng. Med. Biol. Soc. Annu. Conf.*, vol. 2014, pp. 1521–4, Aug. 2014.