

Segmentation of Cerebrovascular Pathologies in Stroke Patients with Spatial and Shape Priors

Adrian V. Dalca¹, Ramesh Sridharan¹, Lisa Cloonan², Kaitlin M. Fitzpatrick²,
Allison Kanakis², Karen L. Furie³, Jonathan Rosand², Ona Wu²,
Mert Sabuncu⁴, Natalia S. Rost², and Polina Golland¹

¹ Computer Science and Artificial Intelligence Lab, MIT

² Department of Neurology, Massachusetts General Hospital, Harvard Medical School

³ Department of Neurology, Rhode Island Hospital, Alpert Medical School

⁴ Martinos Center for Biomedical Imaging, Harvard Medical School

Abstract. We propose and demonstrate an inference algorithm for the automatic segmentation of cerebrovascular pathologies in clinical MR images of the brain. Identifying and differentiating pathologies is important for understanding the underlying mechanisms and clinical outcomes of cerebral ischemia. Manual delineation of separate pathologies is infeasible in large studies of stroke that include thousands of patients. Unlike normal brain tissues and structures, the location and shape of the lesions vary across patients, presenting serious challenges for prior-driven segmentation. Our generative model captures spatial patterns and intensity properties associated with different cerebrovascular pathologies in stroke patients. We demonstrate the resulting segmentation algorithm on clinical images of a stroke patient cohort.

1 Introduction

Identifying and differentiating cerebrovascular pathologies in brain MRI is critical for understanding cerebral ischemia (insufficient blood flow to the brain). Unfortunately, different lesion types, such as leukoaraiosis (small-vessel disease) and stroke, cannot be distinguished purely based on intensities or location. Clinicians use anatomical and other medical knowledge to categorize and delineate pathology. We model intensity, shape, and spatial distribution of pathologies to capture this anatomical knowledge of variability of pathology in order to successfully annotate clinical brain scans in stroke patients.

Our work is motivated by imaging studies of stroke patients that acquire multimodal brain scans within 48 hours of stroke onset. To understand susceptibility to cerebral ischemia and associated risk factors, clinicians manually outline and analyze vascular pathologies, focusing on leukoaraiosis and separating it from stroke lesions. Using this approach, leukoaraiosis burden has been shown to be lower in patients with transient ischemic attacks compared to patients with more damaging cerebral infarcts [11]. Manual delineation of leukoaraiosis and stroke takes up to 30 minutes per patient, and large population studies contain hundreds to thousands of patients. Automatic segmentation is therefore necessary. Here, we focus on segmenting leukoaraiosis and separating it from stroke lesions.




Fig. 1. Left: T2-FLAIR axial slice. Stroke (blue outline) can appear anywhere in the brain, can vary dramatically in shape, and is hyperintense. Leukoaraiosis (yellow outline) is generally peri-ventricular, has a more predictable spatial distribution than stroke lesions, and is usually roughly symmetric across hemispheres. We outline imaging artifacts (red) that can interfere with automatic segmentation of pathologies. Right: Examples of various stroke shapes and sizes, manually outlined in green.

Variability in shape and location of lesions is one of the main challenges in automatic segmentation of stroke scans. Leukoaraiosis appears hyperintense in T2-FLAIR, is found peri-ventricularly, has a widely variable extent, and is roughly bilaterally symmetric. While also hyperintense, strokes can happen nearly anywhere in the brain and vary dramatically in size and shape. While acute stroke (stroke that occurred in the last 48 hours) is visible on diffusion weighted MR (DWI), the same is not true for chronic stroke (stroke that occurred a long time before imaging). Additionally, DWI is often not available [17]. In this paper, we concentrate on the more difficult task of separating leukoaraiosis from stroke, both acute and chronic, in T2-FLAIR. Another challenge is the low quality of images in the clinical setting due to the extremely limited scanning time. This results in thick slices (5-7mm) and bright artifacts, which hinder registration and intensity equalization of clinical images and further complicate automatic segmentation. Representative images and segmentations are shown in Figure 1, illustrating our challenge.

We introduce a generative probabilistic model of the effects of the cerebrovascular disease on the brain. The model integrates important aspects of each pathology, leading to an effective inference algorithm for segmentation and separation of different tissues in stroke patients. Specifically, we learn the spatial distribution and intensity profile of leukoaraiosis, as well as the intensity profile of stroke. We train the model on an expert-labeled dataset and demonstrate that our modeling choices capture notions used by clinicians, such as symmetry and covariation of intensity patterns. To the best of our knowledge, this is the first comprehensive segmentation approach for different cerebrovascular pathologies.

Our model incorporates several approaches previously proposed for segmentation of healthy anatomy that is consistent across individuals [3,15,16]. We combine these methods to accurately model pathology. Intensity-based lesion segmentation algorithms utilize tissue intensities to segment pathology [1,7]. Spatial priors are sometimes added in a form of Markov Random Fields or spatial dis-




Fig. 2. A graphical representation of our generative model. Circles indicate random variable and rounded squares represent parameters. Shaded circles represent observed quantities and the plates indicate replication. I is the acquired image. The image intensities are generated from a normal distribution parametrized by μ_c and σ_c for each tissue class c in the label map \mathbf{C} . Priors for the tissue classes are controlled by the weights $\{\alpha_k\}$ of the K -component PCA shape model and spatial parameters β_x that define the prior probability of stroke in non-leukoaraiosis tissue.

tributions [4,12,15]. These methods are successful in delineating structures that are hyper- or hypointense compared to their surroundings, such as MS lesions or tumors. Unfortunately, these methods are not designed to differentiate between multiple hyperintense structures, such as leukoaraiosis, stroke, and certain artifacts, which share an intensity profile and can co-occur spatially. Clinicians use spatial features, such as the bilateral symmetry of leukoaraiosis, to tell them apart. Shape-based methods generally model the shape of a structure, either via an explicit [3,8,14] or implicit [5,9,10] representation. We utilize a shape model to capture the variability in spatial distribution of leukoaraiosis, which develops in a consistent pattern peri-ventricularly. In contrast, stroke can happen at random locations almost anywhere in the brain, and has no obvious shape or location profile (Figure 1). We demonstrate that combining intensity and spatial context for stroke and spatial distribution models for leukoaraiosis produces accurate segmentation. We validate the method on over 100 stroke subjects with various pathologies and artifacts.

2 Generative Model

We use a generative model to describe the spatial distribution, shape and appearance of healthy tissue and cerebrovascular pathology. Figure 2 provides a graphical representation of our model.

We let Ω be the set of all spatial locations (voxels) in an image, and $I = \{I_x\}_{x \in \Omega}$ be the acquired image. We assume image I is generated from a spatially varying label map $\mathbf{C} = \{\mathbf{C}_x\}_{x \in \Omega}$ that represents tissue classes. For each voxel x , \mathbf{C}_x is a length-3 binary indicator vector that encodes three tissue labels – leukoaraiosis (L), stroke (S) and healthy tissue (H). We use notation $C_x(c) = 1$ to mean that the tissue class at voxel x is c , for $c \in \{L, S, H\}$. Otherwise, $C_x(c) = 0$.

Given the label map \mathbf{C} , the intensity observations I_x are generated independently from a Gaussian distribution:

$$P(\mathbf{I}|\mathbf{C}; \boldsymbol{\mu}, \boldsymbol{\sigma}) = \prod_{x \in \Omega} \prod_{c \in \mathcal{C}} \mathcal{N}(I_x; \mu_c, \sigma_c)^{C_x(c)}, \quad (1)$$

where $\mathcal{N}(\cdot; \mu, \sigma)$ is the normal distribution parametrized by mean μ and variance σ^2 , $\mathcal{C} = \{L, S, H\}$, $\boldsymbol{\mu} = \{\mu_L, \mu_S, \mu_H\}$ and $\boldsymbol{\sigma} = \{\sigma_L, \sigma_H, \sigma_S\}$.

The prior for the tissue classes captures our knowledge about spatial distributions and shape of pathology. We assume that the spatial extent of leukoaraiosis depends on a spatial distribution $\mathbf{M} = \{M_x\}_{x \in \Omega}$, where M_x is a prior for leukoaraiosis for voxel x . As we describe later on, \mathbf{M} will be parametrized by parameter α . If voxel x is not assigned to leukoaraiosis, it is assigned to be stroke with spatially varying probability β_x , and to be healthy tissue with probability $(1 - \beta_x)$. To encourage spatial contiguity, we incorporate a Markov Random Field (MRF) as a spatial prior. Formally,

$$P(\mathbf{C}|\alpha, \beta) = \prod_x \prod_c \pi_x(c)^{C_x(c)} \prod_{y \in N(x)} \exp\left(\mathbf{C}_x^T A \mathbf{C}_y\right), \quad (2)$$

where

$$\pi_x = [M_x(\alpha), (1 - M_x(\alpha))\beta_x, (1 - M_x(\alpha))(1 - \beta_x)]^T \quad (3)$$

is a length-3 vector of prior probabilities for the three tissue classes as described above, $N(x)$ is the set of voxel locations neighboring x , and the 3×3 matrix A is chosen to encourage neighboring voxels to share the same tissue label. In our implementation, the MRF term penalizes interactions between stroke and other tissues more than leukoaraiosis bordering healthy tissue, as we find that stroke is generally more spatially contiguous than leukoaraiosis, which is more diffuse.

Using (1), (2) and (3), we form the posterior distribution for the tissue classes:

$$\begin{aligned} P(\mathbf{C} \mid \mathbf{I}; \mu, \sigma, \alpha, \beta) &\propto P(\mathbf{I}, \mathbf{C}; \mu, \sigma, \alpha, \beta) = P(\mathbf{I}|\mathbf{C}; \mu, \sigma)P(\mathbf{C}; \alpha, \beta) \\ &= \prod_{x \in \Omega} \left(\prod_{c \in \mathcal{C}} [\pi_x(c) \mathcal{N}(I_x; \mu_c, \sigma_c)]^{C_x(c)} \prod_{y \in N(x)} \exp\left(\mathbf{C}_x^T A \mathbf{C}_y\right) \right). \end{aligned} \quad (4)$$

2.1 Spatial Distribution Prior for Leukoaraiosis

We model the spatial extent of leukoaraiosis with a probabilistic atlas constructed by applying Principal Component Analysis (PCA) to a training set of manual leukoaraiosis binary segmentation maps. We let $\bar{\mathbf{M}}$ be the mean map, $\{M_k\}_{k=1}^K$ be the principal components that correspond to the K largest eigenvalues, and α_k be the weights (or loadings):

$$P(\alpha) = \mathcal{N}(\alpha; 0, \Sigma), \quad (5)$$

where Σ is the diagonal covariance matrix containing the K largest eigenvalues. Given α , the spatial prior $\mathbf{M} = \{M_x\}_{x \in \Omega}$ is deterministically defined:

$$\mathbf{M}(\alpha) = \bar{\mathbf{M}} + \sum_k \alpha_k \mathbf{M}_k.$$

We also experimented with LogOdds shape representation [10], often used for modeling normal anatomical variability. We found that the leukoaraiosis structures are in general too thin and variable in location to be properly captured by this representation. On the other hand, a simple average probability map representation fails to capture the covariation of leukoaraiosis distribution.

3 Inference

To obtain the segmentations, we perform MAP inference and seek

$$\hat{C} = \arg \max_C P(C|I; \mu, \sigma, \alpha, \beta) = \arg \max_C P(C, I; \mu, \sigma, \alpha, \beta). \quad (6)$$

Since exact computations become infeasible whenever the MRF weight matrix \mathbf{A} is non-zero, we employ a variational EM approximation [6] to estimate the MAP solution. Specifically, we approximate the posterior distribution $P(C|I; \mu, \sigma, \alpha, \beta)$ with the fully factored distribution

$$q(C) = \prod_{x \in \Omega} q(C_x) = \prod_{x \in \Omega} \prod_c w_x(c)^{C_x(c)}, \quad (7)$$

where \mathbf{w}_x is a vector of probabilities for the three tissue classes at voxel x . Due to space constraints, we omit the derivations and provide the resulting updates. Because the prior for the PCA loadings $P(\alpha)$ is not conjugate to the likelihood $P(C|\alpha)$, we approximate the corresponding E-step computation with a regularized projection:

$$\begin{aligned} \alpha &\leftarrow \arg \min_{\mathbf{a}} \|\mathbf{w}(L) - U\mathbf{a}\|^2 + \lambda \mathbf{a}^T \Sigma^{-1} \mathbf{a} \\ &= (U^T U + \lambda \Sigma^{-1})^{-1} U^T \mathbf{w}(L), \end{aligned} \quad (8)$$

where $U = [M_1, \dots, M_k]$, and we use clipping to force the resulting values in $\mathbf{M}(\alpha)$ to be between 0 and 1.

In the M-step, we update the parameters of the model. The updates are intuitive. The class mean and variance estimates are computed as weighted averages:

$$\mu_c \leftarrow \frac{\sum_x w_x(c) I_x}{\sum_x w_x(c)}, \quad \sigma_c \leftarrow \frac{\sum_x w_x(c) (I_x - \mu_c)^2}{\sum_x w_x(c)}, \quad \text{for } c \in \mathcal{C} \quad (9)$$

Given large variable intensity pathologies and severe artifacts, image inhomogeneity cannot be corrected through pre-processing steps. To address image inhomogeneity for the healthy tissue, we model the intensity mean estimate as spatially varying, and introduce a low pass filter G_H to enforce spatial smoothness, similar to the original EM-segmentation formulation [16]. Specifically,

$$\mu_H \leftarrow G_H * (w_x(H) \cdot I), \quad (10)$$

where $*$ denotes spatial convolution. The healthy tissue prior β_x is a fraction of current frequency estimates for stroke and healthy tissue probabilities:

$$\beta_x \leftarrow \frac{w_x(S)}{w_x(H) + w_x(S)}. \quad (11)$$

Finally, the variational posterior parameters \mathbf{w}_x are weighted by their agreement with the neighbors:

$$w_x(c) \leftarrow \pi_x(c) \mathcal{N}(I_x; \mu_c, \sigma_c) \prod_{y \in N(x)} \exp(\mathbf{w}_x^T \mathbf{A} \mathbf{w}_y), \quad (12)$$

where $\pi_x(c)$ is defined in (3). We iterate the updates until the parameter estimates converge.

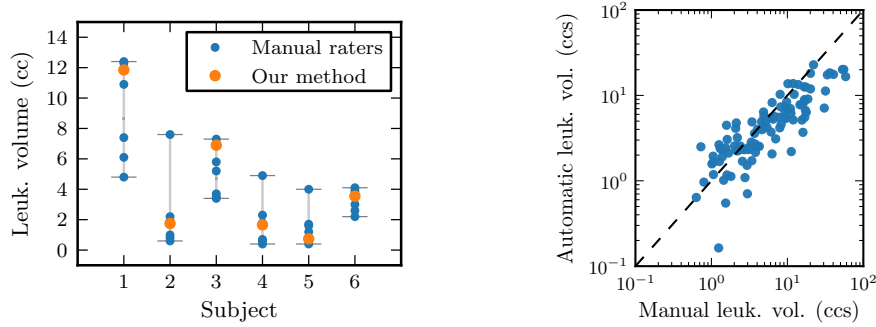


Fig. 3. Leukoaraiosis segmentation evaluation. Left: Volume measurements based on the automatic segmentation (orange) are within the range of experts (blue circles). Right: Volume estimates based on the automatic segmentation of leukoaraiosis against volume estimates based on the manual segmentations; the correlation coefficient is $r = 0.82$.

4 Results

In this section, we present experimental results on 100 test images with manually delineated leukoaraiosis, and another six test volumes, each with manually delineated leukoaraiosis by multiple experts. Leukoaraiosis is the primary phenotype in many stroke studies, and thus its segmentation and delineation from stroke is our main focus in the experiments. We only run our algorithm inside the white matter where we expect to see most of the leukoaraiosis and strokes. A future direction for our work is to include healthy gray matter and cerebrospinal fluid classes.

In our experiments, the scans include T2-FLAIR scans (1×1 mm in-plane, slice thickness 5-7mm, PROPELLER sequence sometimes used if the patient moved). Acquisition TR and TE varied depending on the image protocol. All subjects are registered to an atlas template using ANTs [2] based on T1 images acquired for each subject in the study [13].

Parameters. We trained the PCA shape model $(\{M_k\}, \Sigma)$ on binary maps of manual leukoaraiosis segmentations in 42 training scans, different from the 106 test scans. The fixed parameters λ and A were chosen manually to optimize results in a single test example, not included in the results below. In particular, we use $\lambda = 250$, $A(c, c) = 100$ for $c \in \{L, H, S\}$, $A(L, H) = 97$, $A(S, L) = 1$, and $A(S, H) = 20$. This choice discourages stroke from neighbouring leukoaraiosis more than neighbouring healthy tissue. We initialized the posterior estimates using a simple threshold classifier learned from the training subjects [13].

In Figure 3, we compare the volume of leukoaraiosis obtained by our method against expert delineations. The first graph demonstrates that the automatic approach is consistently within the range of inter-rater variability. The second graph compares the automatically computed leukoaraiosis burden to that based on manual segmentations in a test set of 100 subjects. Visual inspection of the




Fig. 4. Automatic segmentation on three subjects showing stroke segmentation (blue outlines) and leukoaraiosis (yellow outlines): the image used to tune the parameters (left), a typical result of our method (center), and an example of a difficult case (right).

outlier image reveals little to no apparent leukoaraiosis and a possible manual over-segmentation.

Figure 4 provides example segmentation results for three subjects with leukoaraiosis and stroke. The first subject was used for parameter tuning (and was not included in testing), and includes a stroke in areas where leukoaraiosis is often found near the ventricles. The second subject is a typical result, where we see that most of the stroke is accurately separated from leukoaraiosis. The third subject is an example of an outlier result, where the separation of the two hyperintense pathologies is nearly impossible to define even by a clinical expert. In both the second and third subjects, our leukoaraiosis segmentation is conservative, which is likely caused by the regularized projection. This observation is consistent with the results shown in Figure 3.

5 Conclusions

We presented an algorithm for segmentation of separate cerebrovascular pathologies in brain MRI. Our algorithm is derived from a generative probabilistic model that captures experts' knowledge of the disease. By modeling the spatial distribution of leukoaraiosis, as well as the intensities of leukoaraiosis and stroke lesions, our method automatically segments tissues that are indistinguishable based on intensity alone. We presented our method on a study of stroke patients, and showed strong agreement between our results and expert segmentation volumes.

Acknowledgements. We acknowledge the following funding sources: NSERC CGS-D, NSF GRFP, NIH NIBIB 1K25EB013649-01, BrightFocus grant (AHAF-A2012333), NIH NIBIB NAC P41EB015902, and NIH NIBIB NIMIC U54-EB005149, NIH NINDS NS082285, NIH NINDS K23NS064052, NIH NINDS U01NS069208 and the American Stroke Association-Bugher Foundation Centers for Stroke Prevention Research.

References

1. Anbeek, P., Vincken, K., van Osch, M., Bisschops, R., van der Grond, J.: Probabilistic segmentation of white matter lesions in mr imaging. *NeuroImage* 21(3), 1037–1044 (2004)
2. Avants, B., Tustison, N., Song, G., Cook, P., Klein, A., Gee, J.: A reproducible evaluation of ants similarity metric performance in brain image registration. *Neuroimage* 54(3), 2033–2044 (2011)
3. Cootes, T., Hill, A., Taylor, C., Haslam, J.: Use of active shape models for locating structures in medical images. *Image and Vision Computing* 12(6), 355–365 (1994)
4. Geremia, E., Clatz, O., Menze, B., Konukoglu, E., Criminisi, A., Ayache, N.: Spatial decision forests for MS lesion segmentation in multi-channel magnetic resonance images. *NeuroImage* 57(2), 378–390 (2011)
5. Golland, P., Grimson, W., Shenton, M., Kikinis, R.: Detection and analysis of statistical differences in anatomical shape. *Medical Image Analysis* 9(1), 69–86 (2005)
6. Jaakkola, T.: Tutorial on variational approximation methods. *Advanced mean field methods: theory and practice* p. 129 (2001)
7. Johnston, B., Atkins, M., Mackiewicz, B., Anderson, M.: Segmentation of multiple sclerosis lesions in intensity corrected multispectral MRI. *Medical Imaging, IEEE Transactions on* 15(2), 154–169 (1996)
8. Kelemen, A., Szekely, G., Gerig, G.: Elastic model-based segmentation of 3-D neuroradiological data sets. *Medical Imaging, IEEE Transactions on* 18(10), 828–839 (1999)
9. Leventon, M., Grimson, W., Faugeras, O.: Statistical shape influence in geodesic active contours. In: *CVPR, 2000. Proceedings. IEEE Conference on.* vol. 1, pp. 316–323. IEEE (2000)
10. Pohl, K., Fisher, J., Bouix, S., Shenton, M., McCarley, R.W., Grimson, W., Kikinis, R., Wells, W.: Using the logarithm of odds to define a vector space on probabilistic atlases. *Medical Image Analysis* 11(5), 465–477 (2007)
11. Rost, N., Fitzpatrick, K., Biffi, A., Kanakis, A., Devan, W., Anderson, C., Cortellini, L., Furie, K., Rosand, J.: White matter hyperintensity burden and susceptibility to cerebral ischemia. *Stroke* 41(12), 2807–2811 (2010)
12. Sajja, B., Datta, S., He, R., Mehta, M., Gupta, R., Wolinsky, J., Narayana, P.: Unified approach for multiple sclerosis lesion segmentation on brain MRI. *Annals of Biomedical Engineering* 34(1), 142–151 (2006)
13. Sridharan, R., Dalca, A., Fitzpatrick, K., Cloonan, L., Kanakis, A., Wu, O., Furie, K., Rosand, J., Rost, N., Golland, P.: Quantification and analysis of large multimodal clinical image studies: Application to stroke. In: Shen, L., Liu, T., Yap, P., Huang, H., Shen, D., Westin, C. (eds.) *Multimodal Brain Image Analysis, LNCS*, vol. 8159, pp. 18–30. Springer (2013)
14. Styner, M., Lieberman, J., Pantazis, D., Gerig, G.: Boundary and medial shape analysis of the hippocampus in schizophrenia. *Medical Image Analysis* 8(3), 197–203 (2004)
15. Van Leemput, K., Maes, F., Vandermeulen, D., Colchester, A., Suetens, P.: Automated segmentation of multiple sclerosis lesions by model outlier detection. *Medical Imaging, IEEE Transactions on* 20(8), 677–688 (2001)
16. Wells III, W., Grimson, W., Kikinis, R., Jolesz, F.: Adaptive segmentation of MRI data. *Medical Imaging, IEEE Transactions on* 15(4), 429–442 (1996)
17. Wintermark, M., Albers, G., Broderick, J., Demchuk, A., Fiebach, J., other: Acute stroke imaging research roadmap II. *Stroke* 44(9), 2628–2639 (2013)