Extending Supervoxel-based Abnormal Brain Asymmetry Detection to the Native Image Space

Samuel B. Martins[1,2,3], Alexandru C. Telea[3], and Alexandre X. Falcão[1]

Abstract—Most neurological diseases are associated with abnormal brain asymmetries. Recent advances in automatic unsupervised techniques model normal brain asymmetries from healthy subjects only and treat anomalies as *outliers*. Outlier detection is usually done in a common *standard* coordinate space that limits its usability. To alleviate the problem, we extend a recent fully unsupervised supervoxel-based approach (SAAD) for abnormal asymmetry detection in the *native* image space of MR brain images. Experimental results using our new method, called N-SAAD, show that it can achieve higher accuracy in detection with considerably less false positives than a method based on unsupervised deep learning for a large set of MR-T1 images.

Fig. 1. **(a)** Normal asymmetries (white circles) for a healthy subject. **(b-c)** Subtle and large abnormal asymmetries (yellow circles) for stroke patients. Dashed lines show the mid-sagittal planes (MSPs). Arrows show regions in which MSPs do not symmetrically separate the hemispheres.

I. INTRODUCTION

Several normal hemispheric asymmetries can be found in many brain regions. Studies have shown that many neurological diseases present deviations from the *normal* asymmetry pattern, caused by morphological changes in (sub)cortical structures [1], with shape and size alterations in one or both hemispheres (Fig. 1). Brain anomalies are usually detected by visual slice-by-slice inspection by one or multiple specialists. This is subjective, time-consuming, and prone to error. This detection can be done in a *standard* image space (*i.e.,* the coordinate space of a template) or in the own Native Image Space (NIS). The former is chosen when considering a group of reference images from healthy subjects and/or patients for comparison during the analysis. All images are registered in the same coordinate, ensuring spatial correlation (position and shape) among same structures of interest across all images. NIS is commonly used in clinical routine to provide diagnosis, quantification of disease severity, and treatment planning.

Discriminative models (supervised learning) aim to automatically delineate anomalies (*e.g.,* tumors), trying to solve detection and segmentation simultaneously, resulting in segmentation rates still very low on MR-T1 images [2] but higher on CT, T2, FLAIR, and multimodal images [3], [4]. Discriminative models make the method specific for detecting anomalies related to disease(s) present in the training set. Their success on new images is limited by the absence of large, high-quality, annotated training sets, which motivates
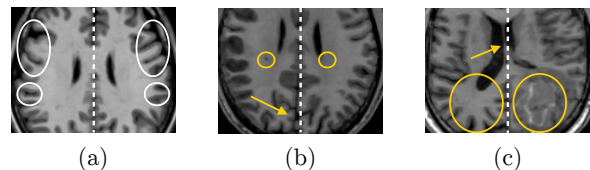
research on *unsupervised* approaches [2], [4], [5], [6], [7]. Unsupervised methods learn a model from control images of healthy subjects only and treat anomalies as *outliers*.

Recently, we proposed an unsupervised supervoxel-based approach, named SAAD, for abnormal asymmetry detection in MR brain images [7]. By having all images registered to the same symmetric template, SAAD computes asymmetries between hemispheres by using their mid-sagittal plane (MSP) as reference. It then extracts pairs of symmetric supervoxels from the left and right hemispheres for each test image, so that each pair generates a one-class classifier. This classifier is next trained on control images to find supervoxels with abnormal asymmetries. Despite having high detection accuracy, SAAD can be *only* used in the *coordinate space* of the symmetric template. However, as said earlier, specialists typically use the NIS for their clinical routine work.

We extend SAAD to perform asymmetry detection in the *native* image space of 3D MR-T1 brain images of 3T. The key challenge is finding corresponding regions between hemispheres, as these differ in shape, size, and positioning in NIS. Thus, we cannot use MSPs to symmetrically separate such structures, especially when severe morphological deformations exist (Fig. 1). Our approach, *Native Supervoxel-based Abnormal Asymmetry Detection* (N-SAAD), automatically segments hemispheres of the test image in its NIS. Next, it flips and registers one hemisphere to the other to guarantee spatial correlation between them. Asymmetries are computed for the correlated hemispheres, and symmetric supervoxels are extracted and mapped to the same coordinate space of the control images to train the one-class classifiers.

We compare two versions of N-SAAD with a deep learning autoencoder (AE) approach derived from [5] and [6] to detect stroke lesions of the ATLAS dataset [8]. Experimental results show that N-SAAD has significantly higher detection rates and considerably less false positives compared to the AE approach.
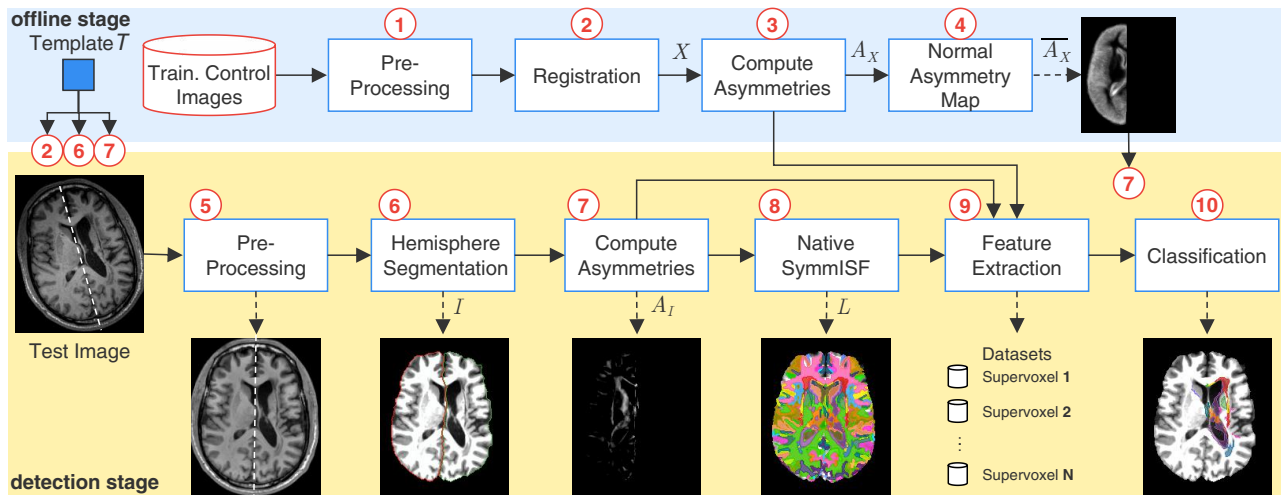
Fig. 2. Pipeline of proposed N-SAAD method. Steps 1 to 4 (blue part) are computed offline. Steps 5 to 10 (yellow part) are computed for each test image (detection stage). The template (reference image) is used in both stages (Steps 2, 6, and 7). Dashed lines on brain images show their mid-sagittal planes The output of Step 8 shows the pairs of symmetrical supervoxels in both hemispheres only for illustration purposes, since N-SAAD uses just one hemisphere for feature extraction.

II. DESCRIPTION OF METHOD

We next describe our approach (see also Fig. 2). All images in this work are 3D (volumes).

A. Preprocessing, Registration, and Segmentation

First, we preprocess the training control image set and the test image (Steps 1 and 5). We apply bias field correction, followed by median filtering for noise reduction, and linear intensity normalization within $[0, 4095]$. We then align images by their mid-sagittal planes (MSPs) which are extracted by [9]. For control images, we perform registration (Step 2) to place them in the coordinate space of the ICBM 2009c Nonlinear Symmetric template [10], whose hemisphere masks and MSP are well defined.

To cope with differences in brightness and contrast among images, we use histogram matching on the segmented hemispheres, as follows. For training images, we use the predefined segmentation mask that comes with the ICBM template; histogram matching is done after Step 2. For test images, we use SOSM-A [11], an effective and efficient probabilistic atlas-based approach for medical image segmentation; histogram matching is done after Step 6.

B. Asymmetry Computation

Let X be the set of registered training images, I the test image after preprocessing and segmentation, and T the template. We obtain the set of asymmetry maps A_X for all X by computing the voxel-wise absolute differences between left and right hemispheres with respective to their MSPs (Step 3). We then average the absolute difference values to create a *normal asymmetry map* $\overline{A_X}$ (Step 4) which is used to attenuate the detection of false positive asymmetries in I in brain regions which are commonly asymmetric (*e.g.,* cortex). For this, we register T on I and then map $\overline{A_X}$ on I by applying the resulting deformation fields (Step 7). Registration is discussed in detail in Sec. III.

Since hemispheres in the NIS can be very different in shape, size, and positioning (Figs. 1-2), we cannot *solely* rely on the MSP to compute asymmetries for I, as performed by SAAD [7]. After flipping the left segmented hemisphere to the right one (using the MSP), we next register both hemispheres (two registrations), and histogram-match them to guarantee spatial correlation between them. We compute voxel-wise absolute differences between the correlated hemispheres and subtract the mapped $\overline{A_X}$ from them. Resulting positive values form a final asymmetry map A_I for the test image I (Step 7).

C. Symmetric Supervoxel Segmentation on NIS

Directly comparing the flipped, segmented, and registered hemispheres is not helpful as it will not tell us where *small-scale* asymmetries occur. At the other extreme, comparing every voxel pair in these hemispheres is risky, since such voxels contain too little information to capture asymmetries. Hence, we propose to segment hemispheres into *supervoxels* which are large enough to capture meaningful asymmetries but small enough to capture a fine level of detail. We use SymmISF [7], a recent method based on the *Iterative Spanning Forest* framework [12] that extracts symmetrical supervoxels from left and right brain hemispheres simultaneously. It consists of three steps: (i) seed estimation; multiple iterations of (ii) connected supervoxel delineation; and (iii) seed recomputation to improve delineation.

Since SymmISF considers that brain hemispheres are equally separated by their MSP, after registration on a symmetric template, we adapt it to work in the NIS (Step 8). To build initial seeds, we select one seed per local maximum of the asymmetry map A_I (Step 6). We compute the local maxima on the foreground of a binarized A_I at 2 times Otsu's threshold, also filtering out small connected components by morphological closing with radius 1. We add a fixed number (100) of seeds in regions with irrelevant asymmetry by grid

sampling the background of the binarized image.

By stacking the resulting left registered and right hemispheres (Step 7) as the input 2-band volume for the method, we perform SymmISF only inside the right hemisphere by using its corresponding mask (Step 6). The result is a label map in which each supervoxel is assigned to a distinct label. Finally, we map these supervoxels by using the deformation fields from the corresponding hemisphere registration (Step 7) to obtain the symmetrical supervoxels in the left hemisphere, which yield the final label map L (Step 8). Note that we could proceed conversely, *i.e.*, apply SymmISF on the left hemisphere, and map the result to the right hemisphere.

D. Feature Extraction and Pattern Classification Models

For each test image I, we first need to map its symmetric supervoxels L on template T to guarantee spatial correlation between supervoxels in both coordinate spaces. For this, we register I on to T and map L by using the resulting deformation fields. This tends to be better since the computation of the inverse deformation from T to I (see Section II-B) is more complicated due to deformable registrations. Each supervoxel pair is then used to create a one-class classifier using as feature vector the normalized histogram of the asymmetry values inside the respective supervoxels (Step 9). This way, we implicitly consider the *position* of the supervoxels in the hemispheres when deciding upon their asymmetry. We use the one-class linear Support Vector Machine (OC-SVM) for this task [13]. The classifiers are trained from *control images* only and used to find supervoxels with abnormal asymmetries in the test image (Step 10).

III. EXPERIMENTS AND RESULTS

We used all 180 MR-T1 images of 3T from the CC359 dataset [14] for training. CC359 is a public dataset with images of healthy subjects acquired on three different scanners. For testing, we evaluated N-SAAD on 3T MR-T1 images of the Anatomical Tracings of Lesions After Stroke (ATLAS) public dataset Release 1.2 [8]. ATLAS is a very challenging dataset with a large variety of manually annotated lesions and images acquired from different scanners. All images have a mask with the *primary* stroke. Some images also have additional masks with other stroke lesions. Current state-of-the-art segmentation results for ATLAS are inaccurate yet [6]. We are not affected by this problem, since we aim to *detect*, and not segment, the lesions. We select 194 images which only contain lesions in the hemispheres.[1]

We compared the proposed method against the autoencoder approach (AE) used in [7], which in turn is derived from [5] and [6]. We chose to compare against AE since this is also an unsupervised method (like ours). We did not compare against supervised methods since these require hard-to-find, high-quality, large labeled training data (especially for MR-T1 brain images), and moreover such data typically

	N-SAAD		AE	
	Affine	Deformable	85^{th} perc.	90^{th} perc.
Primary Stroke	**89.89**%	86.02%	49.74%	18.65%
All Lesions	**70.50**%	64.88%	56.04%	27.13%

captures only specific lesions. In contrast, our method aims to detect any kind of (asymmetric) lesion.

AE was trained from 2D axial slices of all preprocessed training images (output of Step 1, Fig. 2). Hemispheres were automatically segmented by SOSM-A [11], then all images were cropped within their resulting masks and their axial slices were resized to 256×256, following [6]. AE contains three 2D convolutional layers with 16, 8, and 8 filters of patch size 3×3, respectively, followed by *ReLU* activation and 2D max-pooling in the *encoder*, and the corresponding operations in the *decoder*. The mean squared error between reconstructed and expected 2D axial slices was minimized by the *nadam* gradient optimizer [15]. Anomalies were found by thresholding the difference between the input image and its reconstruction to obtain a binary segmentation mask (see columns 2 and 3 in Fig. 3), similarly to [5] and [6].

We evaluated N-SAAD using both *affine* and *deformable* registration (Steps 2, 6, and 7 of Fig. 2). All registrations were performed by Elastix [16][2]. For N-SAAD, we used the same parameters as in [7]: 100 background seeds, $\alpha = 0.08$, $\beta = 3.0$, asymmetry histograms of 128 bins, and $\nu = 0.1$ for the linear OC-SVM. For AE, we follow [5] and select two thresholds as the 85^{th} and 90^{th} percentile from the histogram of reconstruction errors on the considered training set, resulting the brightness of 70 and 104, respectively.

Table I shows the detection results for *primary* and *all* stroke lesions for ATLAS, based on at least 15% overlap between supervoxels with abnormal asymmetries and lesions. Fig. 3 shows some visual results. Both affine and deformable registration instances of N-SAAD were able to detect many more lesions than AE. *Affine* registration was best, finding 89.89% of the *primary* strokes and 70.50% of *all* lesions. For both registration methods, N-SAAD presents higher accuracies for the *primary* strokes than the reported one (82.31%) by SAAD [7]. AE found far fewer lesions (its best results were 49.74% for *primary* and 56.04% for *all* lesions). It might present better results by using a considerable large training set and/or some additional post-processing – which is not considered in [5], [6].

N-SAAD consistently detects well-defined abnormal asymmetries for visual inspection, which may be related to lesions (Fig. 3, image 1). It is also able to find small abnormal asymmetries (Fig. 3, image 2). However, N-SAAD cannot detect very subtle and/or very small asymmetries (Fig. 3, image 3) and pairs of similar and symmetric lesions in the same region in both hemispheres (Fig. 3, image 4). Deformable registration N-SAAD yields less accentuated
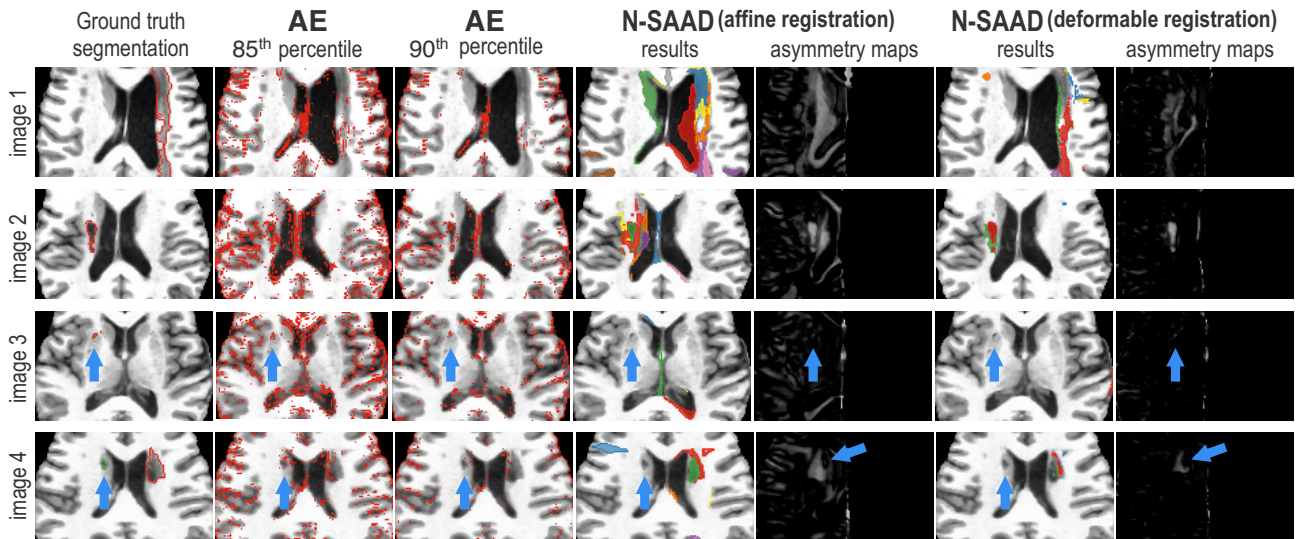
Fig. 3. Results on ATLAS dataset. Each row is a test image. First column: Ground-truth lesion segmentations. Columns 2–4 and 6: Results of AE and N-SAAD methods. Columns 5 and 7: N-SAAD asymmetry maps with affine and deformable registration respectively. Arrows indicate undetected lesions.

asymmetry maps than affine registration N-SAAD (Fig. 3, Columns 5 and 7) since it uses localized deformations to better align hemispheres. Despite yielding fewer false positive anomalies in commonly asymmetric regions (*e.g.,* the cortex), deformable registration N-SAAD is also slower and may also suppress true abnormal asymmetries (*e.g.,* the ventricles of Fig. 3, image 1).

In spite of detecting small and medium lesions when using a lower threshold, AE outputs drastically more false positive voxels than N-SAAD mainly when decreasing its threshold value. Indeed, its high number of false positives makes the visual inspection less effective (Fig. 3). AE seems to be more aligned and accurate to detect lesions in other medical image modalities, such as MR-T2, as originally designed in [5] and [6]. Lesions are considerably more highlighted in MR-T2 compared to MR-T1.

All experiments were executed on an Intel i7 3.60GHz PC with 64GB RAM and a NVIDIA Titan XP 12GB GPU. N-SAAD with *affine* registration takes around 3 mins and 45 s. The *deformable* version takes around 6 mins and 20 s.

IV. CONCLUSION

We extended the recent fully unsupervised supervoxel-based approach (SAAD) for abnormal asymmetry detection in the *native* image space of MR brain images. Our approach, named N-SAAD, was validated on 3T MR-T1 images of stroke patients with annotated lesions, attaining much better detection accuracy and drastically lower false positives compared to an autoencoder approach. We next intend to improve N-SAAD by refining its symmetric supervoxel segmentation to define supervoxels in subtle lesions, evaluate other techniques for feature extraction, improve the normal asymmetry map computation to reduce false positives without missing anomalies, and optimize its parameters. We also want to compare N-SAAD more extensively against a wider set of competing methods, including supervised ones.

REFERENCES

[1] L. Wang and et al., "Statistical analysis of hippocampal asymmetry in schizophrenia," *Neuroimage*, vol. 14, no. 3, pp. 531–545, 2001.

[2] D. Guo *et al.*, "Automated lesion detection on MRI scans using combined unsupervised and supervised methods," *BMC Medical Imaging*, vol. 15, no. 1, p. 50, 2015.

[3] M. Havaei and et al., "Brain tumor segmentation with deep neural networks," *Med Image Anal*, vol. 35, pp. 18–31, 2017.

[4] D. Sato *et al.*, "A primitive study on unsupervised anomaly detection with an autoencoder in emergency head ct volumes," in *SPIE Med. Imag.*, 2018, p. 105751P.

[5] C. Baur, B. Wiestler, S. Albarqouni, and N. Navab, "Deep autoencoding models for unsupervised anomaly segmentation in brain MR images," *arXiv preprint arXiv:1804.04488*, 2018.

[6] X. Chen, N. Pawlowski, M. Rajchl, B. Glocker, and E. Konukoglu, "Deep generative models in the real-world: An open challenge from medical imaging," *arXiv preprint arXiv:1806.05452*, 2018.

[7] S. B. Martins, G. Ruppert, F. Reis, C. L. Yasuda, and A. X. Falcão, "A supervoxel-based approach for unsupervised abnormal asymmetry detection in MR images of the brain," in *Proc. IEEE ISBI*, 2019, pp. 882–885.

[8] S.-L. Liew and et al., "A large, open source dataset of stroke anatomical brain images and manual lesion segmentations," *Scientific Data*, vol. 5, p. 180011, 2018.

[9] G. Ruppert and et al., "A new symmetry-based method for mid-sagittal plane extraction in neuroimages," in *Proc. IEEE ISBI*, 2011, pp. 285–288.

[10] V. S. Fonov and et al., "Unbiased nonlinear average age-appropriate brain templates from birth to adulthood," *Neuroimage*, vol. 47, p. S102, 2009.

[11] S. B. Martins and et al., "A multi-object statistical atlas adaptive for deformable registration errors in anomalous medical image segmentation," in *SPIE Med. Imag.*, 2017, pp. 101 332G–1–8.

[12] J. E. Vargas-Muñoz and et al., "An iterative spanning forest framework for superpixel segmentation," *IEEE T Image Process*, 2019, in press.

[13] L. M. Manevitz and M. Yousef, "One-class SVMs for document classification," *J Mach Learn Res*, vol. 2, pp. 139–154, Dec 2001.

[14] R. Souza and et al., "An open, multi-vendor, multi-field-strength brain MR dataset and analysis of publicly available skull stripping methods agreement," *Neuroimage*, vol. 170, pp. 482–494, 2018.

[15] T. Dozat, "Incorporating nesterov momentum into adam," Stanford University, Tech. Rep., 2015.

[16] S. Klein and et al., "Elastix: A toolbox for intensity-based medical image registration," *IEEE T Med Imag*, vol. 29, no. 1, pp. 196–205, 2010.