



HAL
open science

Leveraging 3D information in unsupervised brain MRI segmentation

Benjamin Lambert, Maxime Louis, Senan Doyle, Florence Forbes, Michel Dojat, Alan Tucholka

► **To cite this version:**

Benjamin Lambert, Maxime Louis, Senan Doyle, Florence Forbes, Michel Dojat, et al.. Leveraging 3D information in unsupervised brain MRI segmentation. ISBI 2021 - 18th International Symposium on Biomedical Imaging, Apr 2021, Nice / Virtual, France. pp.1-4, 10.1109/ISBI48211.2021.9433894 . hal-03273953

HAL Id: hal-03273953

<https://hal.science/hal-03273953>

Submitted on 29 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LEVERAGING 3D INFORMATION IN UNSUPERVISED BRAIN MRI SEGMENTATION

*Benjamin Lambert*¹, *Maxime Louis*¹, *Senan Doyle*¹, *Florence Forbes*²,
*Michel Dojat*³, *Alan Tucholka*¹, for the Alzheimer’s Disease Neuroimaging Initiative*

¹ Pixyl, Research and Development Laboratory, 38000 Grenoble, France

² Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

³ Univ. Grenoble Alpes, Inserm, U1216, Grenoble Institut Neurosciences, GIN, 38000 Grenoble, France

ABSTRACT

Automatic segmentation of brain abnormalities is challenging, as they vary considerably from one pathology to another. Current methods are supervised and require numerous annotated images for each pathology, a strenuous task. To tackle anatomical variability, Unsupervised Anomaly Detection (UAD) methods are proposed, detecting anomalies as outliers of a healthy model learned using a Variational Autoencoder (VAE). Previous work on UAD adopted 2D approaches, meaning that MRIs are processed as a collection of independent slices. Yet, it does not fully exploit the spatial information contained in MRI. Here, we propose to perform UAD in a 3D fashion and compare 2D and 3D VAEs. As a side contribution, we present a new loss function guarantying a robust training. Learning is performed using a multicentric dataset of healthy brain MRIs, and segmentation performances are estimated on White-Matter Hyperintensities and tumors lesions. Experiments demonstrate the interest of 3D methods which outperform their 2D counterparts.

Index Terms— Deep Learning, Variational Autoencoder, Anomaly Detection, Medical Imaging

1. INTRODUCTION

Brain anomalies are widely used as indicators suggesting the presence or progress of many neurological disorders. Magnetic Resonance Imaging (MRI) is today an essential modality to reveal these markers. In recent years, a multitude of algorithms for the automatic detection of anomalies in brain MRI has been proposed, with promising results achieved by Deep Learning (DL) approaches.

*Data used in preparation of this article were partially obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf

Anomalies can manifest in a wide range of shape, location and intensities. To automatically detect them, most state-of-the-art algorithms use a supervised approach, trained on an as large as possible manually annotated dataset [1]. Acquisition of such dataset can be both time-consuming and expensive. Besides, these algorithms are specific to the task they are trained for and will therefore perform poorly when applied on images presenting unseen type of anomalies.

To overcome these limitations, Unsupervised Anomaly Detection (UAD) methods are proposed, consisting in the construction of a manifold of healthy scans. Recent DL methods propose to learn it by computing a low dimensional latent data representation. At inference, anomalous images are projected to the learned healthy manifold, allowing to reveal abnormalities. An anomaly map is then computed using the voxel-wise difference from its healthy projection. Although failing for the time being to meet the state-of-the-art results of supervised methods, UAD approaches have several advantages: they do not require any labeled data and can be adapted with very limited adjustments to different types of pathologies.

Two types of neural networks are particularly investigated in UAD : Generative Adversarial Networks (GANs) [2] and Variational Autoencoders (VAEs) [3]. GANs use an adversarial training approach to construct the healthy manifold [4]. VAEs rather encode images into latent distributions and use a regularization term in their loss function to ensure consistency in the latent space [5]. Yet, training of VAEs is known to be unstable [6]. As an attempt to alleviate this issue, we propose a robust loss function to ensure a stable training.

To the best of our knowledge, all previous research related to UAD tackle the problem in a 2D fashion, meaning that the MRI volume is processed as a collection of independent 2D slices. Despite being computationally efficient, this approach has several limitations. First, it requires the selection of homogeneous slices. Top and bottom slices, which contain no or little brain tissue, are excluded, as they may compromise the training of 2D networks. It is an impairing step because it requires hand-tuning and prevents the detection of anomalies in the entire MRI. Second, the processing of 2D data does

not fully exploit the spatial information available in the scans. This restriction to 2D may arise from the significant difficulty of 3D training, which suffers from the high dimensionality of data and the reduction of the number of training samples. Another limitation of the state of the art is the restriction of most previous work on UAD to monocentric datasets for training [5, 7, 8], which fails to ensure a clinically realistic setup.

In this work, we present a 3D VAE framework for UAD. We train 2D and 3D models using our proposed robust loss function and compare them on multiple pathologies acquired in real clinical conditions. Experiments show that exploiting the full images in a 3D fashion leads to better performances than the usual 2D approaches.

2. METHODS

2.1. UAD using Variational Autoencoders

VAEs are composed of an encoder followed by a decoder. The encoder compresses the input image X into a distribution over the latent space, which is regularized to be close to a prior distribution conventionally set to be a Gaussian multivariate distribution $N(z; 0, I)$. A point z is then sampled from this latent distribution and presented to the decoder which produces a reconstruction \hat{X} of the input.

The bottleneck separating the encoder and the decoder distinguishes spatial and dense VAEs. In the spatial configuration, VAEs are fully-convolutional, meaning that the latent vector is a multi-dimensional tensor $z \in R^{N \times h \times w}$ in 2D and $z \in R^{N \times d \times h \times w}$ in 3D with N the latent space dimension, h the height, w the width and d the depth. In this configuration, spatial information is preserved during the encoding-decoding process. In the dense configuration, the image is encoded into a 1-dimensional tensor $z \in R^N$ with the use of fully-connected layers, and as a consequence spatial information is discarded.

After training on healthy subjects, we perform UAD using a reconstruction-based approach. Reconstruction is the classical approach for DL-based UAD. It consists in the computation of anomaly maps by subtracting its reconstruction \hat{X} to the input image X . As a result of the training, anomalies are badly reconstructed once projected to the healthy manifold, and thus yield to high values in the anomaly map. By binarizing them with an appropriate threshold, anomalies segmentations are obtained.

2.2. Collapsing-robust loss function

Training of a VAE is performed by minimizing a two-terms loss function: the evidence lower-bound ($ELBO(X, \hat{X})$):

$$\mathcal{L} = |X - \hat{X}|_d + D_{KL}(q_\phi(z|X)|P(z)), d \in \{1, 2\} \quad (1)$$

The left-hand side term of this equation is the reconstruction term, which computes the distance between the image X

and its reconstruction \hat{X} . The second term is the Kullback-Leibler Divergence (KL) between the posterior $q_\phi(z|X)$ and the defined prior $P(z)$, which acts as a regularization term in the latent space. Training of VAE using equation (1) is prone to *posterior collapse*, meaning that the network may learn to neglect a subset of latent variables to match the prior, and thus the generative power of the network degrades [6]. Solutions have been proposed to tackle the problem, including the addition of a hyperparameter β to balance both terms of the equation [9]. Following the same motivation, KL annealing schedules have been presented, including the KL cyclical annealing schedule described in [10]. We build on this work and propose a custom loss function defined as follows:

$$\mathcal{L}_T = \frac{|X - \hat{X}|_1}{\Sigma} + \beta(t)D_{KL}(q_\phi(z|X)|P(z)) \quad (2)$$

with $\beta(t) = \begin{cases} \frac{2t}{T} & \text{for } t \in [0, \frac{T}{2}] \\ 1 & \text{else} \end{cases}$ and $\beta(t+T) = \beta(t)$

where Σ is the moving mean computed on the L last values of the reconstruction term, t is the current iteration and T the cycle duration. T and L are two hyperparameters that we set to 50 and 10 respectively in our experiments. We choose the ℓ_1 norm as reconstruction term to reduce blurring of the reconstructed images. Normalizing this term by its moving mean maintains the reconstruction term near one, which encourages the network to learn throughout the entire training stage even when reconstruction becomes satisfying. Coupled with the KL cyclical annealing schedule, we obtain a loss robust to collapsing which ensures a stable training of VAEs. We use this loss without any additional parameter tuning in all our experiments.

2.3. Architecture Details

Our encoders are composed of 6 layers: 4 convolutional layers compressing the input image, and 2 additional layers to encode the result into a normal distribution by computing its parameters μ and σ . These two last layers differ, depending on the model architecture. They are convolutional in Spatial VAEs and fully-connected in Dense VAEs. A latent vector z is then sampled from this distribution and passed to the decoder.

Decoders adopt a symmetric architecture with 5 layers decompressing z in order to reconstruct the input image. Similarly, the first layer of the decoder is convolutional in the spatial configuration, and fully-connected otherwise. Spatial 3D and Dense 3D VAEs are obtained by replacing 2D convolutions by their 3D counterparts.

3. EXPERIMENTS

We focus on brain MRI FLAIR. To train our networks, we gather 79 healthy scans from several opensource datasets :

the ADNI dataset, IBC ¹[11] and Kirby [12]. Evaluation of our networks is performed on 2 different datasets presenting various brain lesions. We use a collection of 196 scans with annotated White-Matter Hyperintensities (WMH) gathered from several opensource datasets : the MICCAI MSSEG Challenge [13], the ISBI MS Lesion Challenge [14], and the WMH Challenge [15]. Additionally, we select 100 brain tumors scans from the BraTS 2018 dataset [16]. These datasets are multi-centric and comprise various scanners and clinical protocols, hence assuring a clinically realistic setup.

Scans are rigidly registered on a $170 \times 204 \times 170$ template with an isotropic resolution of $1mm$, and corrected for bias using [17]. Training and WMH scans are skull-stripped using the HD-BET algorithm [18]. Volumes are then cropped to $160 \times 192 \times 75$ to concentrate on central slices, for which brain tissue is abundant. Finally, intensity range is set to $[0, 1]$.

Segmentations are produced by binarizing anomaly maps, obtained as the difference between an image and its reconstruction. Binarization is performed using a threshold above which voxels are considered as anomalies. In order to find its optimal value, we split each testing set in half. On the first half, we compute Dice scores between ground truths and segmentations obtained with 15 different thresholds, ranged between 0 and 0.15, which we found to be a relevant range in our experiments. Then, the threshold which provided the best performance is applied on the second half.

The obtained segmentations are then multiplied by a slightly eroded brain mask, in order to remove false positives occurring near the brain contour. Finally, a median filter is applied to remove regions with less than 10 voxels.

4. RESULTS

We evaluate our networks segmentation performance using Dice, sensitivity and specificity scores. Reconstruction performance is estimated using the voxel-wise mean average error (MAE) between the input scan and its reconstruction. To fairly compare 2D and 3D networks, they are all evaluated on the same 75 central slices for each test scan. Performances of the Spatial 2D and Dense 2D, as well as their 3D adaptations are presented in Table 1. Illustrative reconstructions and segmentations can be found in Figure 1 for both WMH and tumor lesions. 3D adaptations demonstrate a boost in performance as compared to 2D. On both the WMH and Tumors datasets, the best network is the Dense 3D VAE, with dice scores of 0.463 ± 0.259 and 0.650 ± 0.190 , a specificity of 0.996 and 0.981, and a sensitivity of 0.511 and 0.711 respectively.

To compare our unsupervised method with classic supervised DL approaches, we trained 3D U-Nets in a 4-fold cross-validation setting, on each test dataset. Their performances are presented in Table 1 as Supervised 3D.

¹This data was obtained from the OpenfMRI database. Its accession number is ds000244.

5. DISCUSSION AND CONCLUSION

In this work, the interest of 3D methods for UAD in brain MRIs was evaluated. Using our proposed collapsing-robust loss function, 2D and 3D models were compared on two pathological datasets showing an overall increase in performance in the 3D setting. For spatial VAEs, this gain was more significant than for dense networks. In the dense configuration, spatial information was not preserved during the encoding-decoding scheme. Thus, the introduction of additional spatial context with the use of 3D did not significantly increase the segmentation performance. Additionally, we observed that dense networks outperformed spatial networks in both the 2D and 3D cases. They indeed operated a significantly higher compression of the scan as compared to spatial networks. Consequently, reconstructions presented very little details, and anomalies were entirely discarded, which made them more easily detectable.

While models trained with our proposed loss function achieved better reconstruction quality, they demonstrated no improvement regarding anomaly segmentation performances, as compared to models trained with (1). These results seem to indicate the existence of a trade-off between both effects.

Performances of our approaches, although being satisfying, remain lower than standard supervised methods. Although UAD methods fail at reaching these standards, they are very promising as they allow to detect anomalies in a generic fashion and without the use of annotated training samples. One interesting application of UAD approaches is the generation of raw segmentations, which can be used as a starting point for manual segmentations, saving time for the rater. In future work, advanced 3D architectures could be tested with the use of deeper networks. Extensions to multi-sequential MRI data is also of interest.

6. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available by the following sources: ADNI, OpenfMRI, NITRC, the MICCAI MSSEG Challenge, the ISBI MS Lesion Challenge, the WMH Challenge and BraTS 2018. Ethical approval was not required as confirmed by the license attached with the data.

7. ACKNOWLEDGMENTS

BL, ML, SD and AT are employees of the Pixyl Company. MD and FF serve on Pixyl advisory board.

8. REFERENCES

- [1] Z. Akkus et al., "Deep learning for brain MRI segmentation: state of the art and future directions," *J Dig Im*, vol. 30, no. 4.

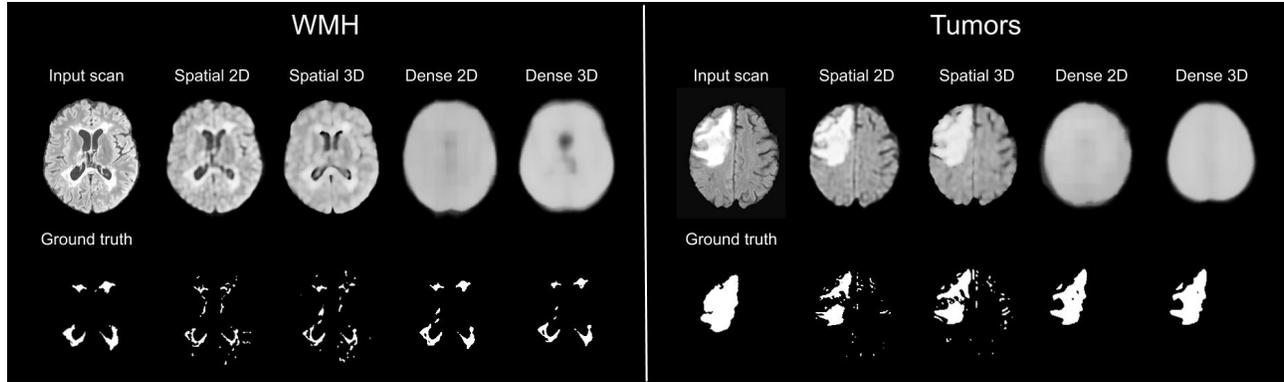


Fig. 1: Reconstructions and segmentations of our networks. Left : WMH lesions. Right : tumors lesions.

Model	WMH				Tumors			
	Dice	Spe	Sen	MAE	Dice	Spe	Sen	MAE
Spatial 2D	0.280 ± 0.174	0.986	0.547	0.049	0.260 ± 0.121	0.912	0.376	0.079
Spatial 3D	0.336 ± 0.203	0.989	0.605	0.056	0.386 ± 0.160	0.898	0.655	0.065
Dense 2D	0.460 ± 0.262	0.994	0.604	0.102	0.618 ± 0.167	0.978	0.676	0.133
Dense 3D	0.463 ± 0.259	0.996	0.511	0.104	0.650 ± 0.190	0.981	0.711	0.151
Supervised 3D	0.750 ± 0.097	0.992	0.789	—	0.740 ± 0.174	0.999	0.738	—

Table 1: Performance metrics obtained on the test datasets. Spe = Specificity, Sen = Sensitivity, MAE = Mean Average Error.

- [2] Ian Goodfellow et al., “Generative adversarial nets,” *Adv Neural Inf Proc Sys*, 2014.
- [3] D. P. Kingma et al., “Auto-encoding variational Bayes,” *Int Conf Learn Repres*, 2013.
- [4] F. Di Mattia et al., “A survey on GANs for anomaly detection,” *CoRR*, 2019.
- [5] C. Baur et al., “Autoencoders for unsupervised anomaly segmentation in brain MR images: A comparative study,” *CoRR*, 2020.
- [6] J. Lucas et al., “Understanding posterior collapse in generative latent variable models,” *Int Conf Learn Repres*, 2019.
- [7] X. Chen et al., “Unsupervised detection of lesions in brain MRI using constrained adversarial auto-encoders,” *CoRR*, 2018.
- [8] D. Zimmerer et al., “Context-encoding variational autoencoder for unsupervised anomaly detection,” *CoRR*, 2018.
- [9] I. Higgins et al., “beta-VAE: Learning basic visual concepts with a constrained variational framework,” *Int Conf Learn Repres*, 2016.
- [10] H. Fu et al., “Cyclical annealing schedule: A simple approach to mitigating KL vanishing,” *NAACL-HLT*, vol. 1, 2019.
- [11] A. L. Pinho et al., “Individual brain charting, a high-resolution fmri dataset for cognitive mapping,” *Scientific data*, vol. 5, 2018.
- [12] B. A. Landman et al., “Multi-parametric neuroimaging reproducibility: a 3T resource study,” *NeuroIm*, 2011.
- [13] Olivier Commowick et al., “Objective evaluation of multiple sclerosis lesion segmentation using a data management and processing infrastructure,” *Scientific reports*, vol. 8, no. 1, 2018.
- [14] Aaron Carass et al., “Longitudinal multiple sclerosis lesion segmentation: resource and challenge,” *NeuroImage*, vol. 148, 2017.
- [15] Hugo J. Kuijf et al., “Standardized assessment of automatic segmentation of white matter hyperintensities and results of the wmh segmentation challenge,” *IEEE transactions on medical imaging*, vol. 38, no. 11, 2019.
- [16] B. H. Menze et al., “The multimodal brain tumor image segmentation benchmark (BRATS),” *TMI*, vol. 34, no. 10, 2014.
- [17] N. J. Tustison et al., “N4itk: improved N3 bias correction,” *TMI*, vol. 29, no. 6, 2010.
- [18] F. Isensee et al., “Automated brain extraction of multi-sequence MRI using artificial neural networks,” *HBM*, vol. 40, no. 17, 2019.