Anatomical Priors for Image Segmentation via Post-Processing with Denoising Autoencoders

Anonymous Author(s)
Affiliation
Address
email

1 Research Problem

Segmentation of anatomical structures is a fundamental task for biomedical image analysis. During the last years, convolutional neural networks (CNNs) proved to be highly accurate to perform segmentation in biomedical images [1, 2, 3]. One of the distinctive features of CNNs is the use of parameter sharing to reduce model complexity and introduce translation invariance. This is specially useful for tasks like image classification, where invariance to translation is a desired property. However, in case of anatomical structures in medical images where their location tend to be highly regular, this property leads to incorrect predictions in areas with similar intensities when enough contextual information is not considered. This issue can be alleviated by introducing prior knowledge about shape ant topology.

One popular strategy to incorporate such priors into medical image segmentation using CNNs is to modify the loss used to train the model. The work of [4] incorporates high-order regularization through a topology aware loss function. In [5, 6], an autoencoder (AE) is used to define a loss term that imposes anatomical constraints during training. The main disadvantage of these approaches is that they can only be used during training of CNN architectures.

On the other hand, post-processing methods (e.g. Conditional Random Fields [2]) to incorporate connectivity constraints into the resulting masks have also been considered in the literature. These methods are based on the assumption that objects are usually continuous and therefore nearby pixels should be assigned the same object label. Even if it is a valid assumption in general, they do not offer a straightforward way to incorporate more complex priors like convexity or shape restrictions.

In this work, we introduce Post-DAE [7] (post-processing with denoising AE), a post-processing method which produces anatomically plausible segmentations by improving pixel-level predictions coming from arbitrary classifiers, incorporating shape and topological priors. The proposed method is rooted in the so-called manifold assumption [8], which states that natural high dimensional data (like anatomical segmentation masks) concentrate close to a non-linear low-dimensional manifold. We learn such low-dimensional anatomically plausible manifold using the aforementioned DAE. Then, given a segmentation mask S_i^P obtained with an arbitrary predictor P (e.g. CNN or RF), we project it into that manifold using f_{enc} and reconstruct the corresponding anatomically feasible mask with f_{dec}.

2 Experiments

We benchmark the proposed method in the context of lung segmentation in X-Ray images, using the Japanese Society of Radiological Technology (JSRT) database [9]. We divide the database in 3 folds considering 70% for training, 10% for validation and 20% for testing. We train Post-DAE using Adam Optimizer with a loss function based on the Dice coefficient; a learning rate of 0.0001; batch size of 15 and 150 epochs. It receives 1024x1024 noisy binary segmentations as input. A degradation function ϕ is used to degrade the ground-truth segmentation masks.

Submitted to 33rd Conference on Neural Information Processing Systems (NeurIPS 2019). Do not distribute.

We compare Post-DAE with a post-processing method based on a fully connected CRF [10]. Differently from our method which uses only binary segmentations for post-processing, the CRF incorporates intensity information from the original images. As *baseline segmentation methods*, we train two different models which produce segmentation masks of various qualities to benchmark our post-processing method. The first model is a CNN based on UNet architecture [1]. To evaluate the effect of Post-DAE in different masks, we save the UNet model every 5 epochs during training, and predict segmentation masks for the test fold using all these models. The second method is a RF classifier trained using intensity and texture features.

3 Results and discussion.

Figure 1 shows some visual examples and the quantitative results. Both figures show the consistent improvement that can be obtained using Post-DAE as a post-processing step, specially in low quality segmentation masks like those obtained by the RF model and the UNet trained for only 5 epochs. In these cases, substantial improvements are obtained in terms of Dice coefficient and Hausdorff distance, by bringing the erroneous segmentation masks into an anatomically feasible space. In case of segmentations that are already of good quality (like the UNet trained until convergence), the post-processing significantly improves the Hausdorff distance, by erasing spurious segmentations (holes in the lung and small isolated blobs) that remain even in well trained models. When compared with CRF post-processing, Post-DAE significantly outperforms the baseline in the context of anatomical segmentation.

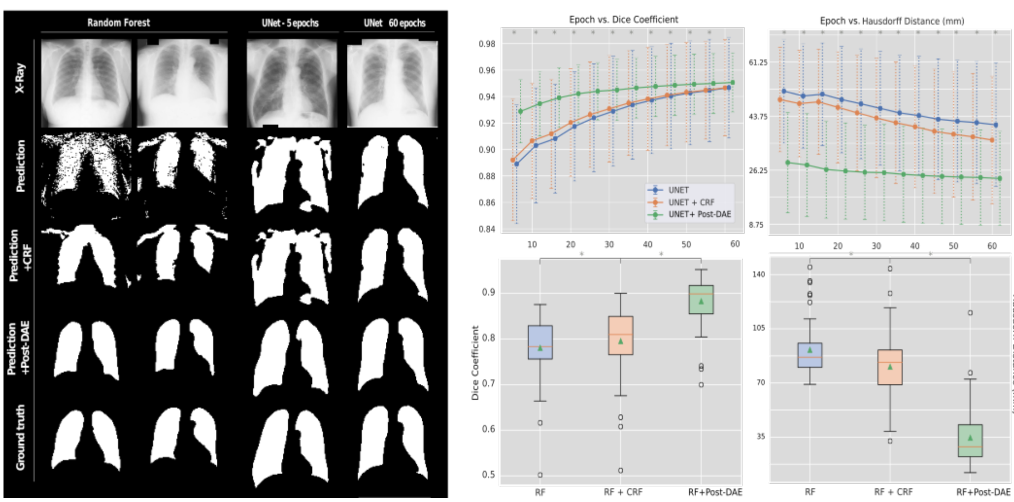

Figure 1: Left: Qualitative evaluation of the proposed method. Right: Top row shows mean and std for post-processing UNet predictions on the test fold at different training stages. Bottom row show results for post-processing the RF predictions. The symbol $*$ indicates that Post-DAE outperforms the other methods with statistical significance (p-value < 0.05 according to Wilcoxon test).

4 Conclusions.

In this work we have showed, for the first time in the MIC community, that DAE can be used as an independent post-processing step to incorporate anatomical priors into arbitrary segmentation methods. Post-DAE can be easily implemented, using segmentation-only datasets or anatomical masks coming from arbitrary image modalities, since no intensity information is required during training. We have validate Post-DAE in the context of lung segmentation in X-ray images, benchmarking with other classical post-processing method and showing its robustness by improving segmentation masks coming from both, CNN and RF-based classifiers.

References

[1] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proc. of MICCAI*, 2015.

[2] Konstantinos Kamnitsas et al. Efficient multi-scale 3d CNN with fully connected CRF for accurate brain lesion segmentation. *Medical Image Analysis*, 36:61 – 78, 2017.

[3] Mahsa Shakeri et al. Sub-cortical brain structure segmentation using F-CNN's. In *Proc. of ISBI*, 2016.

[4] Aïcha BenTaieb and Ghassan Hamarneh. Topology aware fully convolutional networks for histology gland segmentation. In *Proc. of MICCAI*, 2016.

[5] Ozan Oktay et al. Anatomically constrained neural networks (ACNNs): application to cardiac image enhancement and segmentation. *IEEE TMI*, 37(2):384–395, 2018.

[6] Hariharan Ravishankar et al. Learning and incorporating shape models for semantic segmentation. In *Proc. of MICCAI*, 2017.

[7] Agostina J Larrazabal, Cesar Martinez, and Enzo Ferrante. Anatomical priors for image segmentation via post-processing with denoising autoencoders. *MICCAI 2019*, 2019.

[8] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. *Semi-supervised learning*. MIT Press, 2009.

[9] Junji Shiraishi et al. Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules. *Am Jour of Roent*, 174(1):71–74, 2000.

[10] Philipp Krähenbühl and Vladlen Koltun. Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials. In *Proc. of Nips*, 2011.