

Post-DAE: Anatomically Plausible Segmentation via Post-Processing with Denoising Autoencoders

Agostina J Larrazabal, César Martínez, Ben Glocker and Enzo Ferrante

Abstract—We introduce Post-DAE, a post-processing method based on denoising autoencoders (DAE) to improve the anatomical plausibility of arbitrary biomedical image segmentation algorithms. Some of the most popular segmentation methods (e.g. based on convolutional neural networks or random forest classifiers) incorporate additional post-processing steps to ensure that the resulting masks fulfill expected connectivity constraints. These methods operate under the hypothesis that contiguous pixels with similar aspect should belong to the same class. Even if valid in general, this assumption does not consider more complex priors like topological restrictions or convexity, which cannot be easily incorporated into these methods.

Post-DAE leverages the latest developments in manifold learning via denoising autoencoders. First, we learn a compact and non-linear embedding that represents the space of anatomically plausible segmentations. Then, given a segmentation mask obtained with an arbitrary method, we reconstruct its anatomically plausible version by projecting it onto the learnt manifold. The proposed method is trained using unpaired segmentation mask, what makes it independent of intensity information and image modality. We performed experiments in binary and multi-label segmentation of chest X-ray and cardiac magnetic resonance images. We show how erroneous and noisy segmentation masks can be improved using Post-DAE. With almost no additional computation cost, our method brings erroneous segmentations back to a feasible space.

Index Terms—anatomical segmentation, autoencoders, convolutional neural networks, learning representations, post-processing

I. INTRODUCTION

ANATOMICAL segmentation is a fundamental task in medical image computing, which consists in associating pixels of a medical image with a given organ or anatomical structure. It constitutes an essential step in many imaging pipelines such as computer assisted diagnosis, morphometric analysis for population studies and radiotherapy planning. The correctness and anatomical plausibility of these results is thus of paramount importance, since it will directly influence the overall quality of subsequent analyses.

Article accepted for publication in IEEE Transactions on Medical Imaging. Copyright (c) 2020 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

A.J. Larrazabal, C. Martínez and E. Ferrante are with the Institute for Signals, Systems and Computational Intelligence, sinc(i) CONICET-UNL, Santa Fe, Argentina. (e-mails: alarrazabal@sinc.unl.edu.ar - cmartinez@sinc.unl.edu.ar, eferrante@sinc.unl.edu.ar). B. Glocker is with the Biomedical Image Analysis Group, Imperial College London, London, UK. (e-mail: b.glocker@imperial.ac.uk)

E. Ferrante is beneficiary of an AXA Research Fund grant. The authors gratefully acknowledge NVIDIA Corporation with the donation of the GPUs used for this research, and the support of UNL (CAID-PIC-50420150100098LI, CAID-PIC-50220140100084LI) and ANPCyT (PICT 2016-0651, PICT 2018-03907).

Convolutional neural networks (CNNs) proved to perform biomedical image segmentation in a highly accurate way [1]–[3]. CNNs constitute a particular type of neural network specially suited for regularly structured data, like 2D or 3D images, where hierarchical representations of the input are learned using stacked convolutional layers. At every layer, shared parameters (also referred as weights or kernel) are used to learn new representations of the input image. This sharing scheme reduces the number of parameters that should be learnt and allows the use of CNNs in large images. Thanks to the inherently regular structure of the images, these parameters are successively convoluted with the input data resulting in more abstract representations. This trick is particularly helpful for tasks in which invariance to translation is an expected property, such as image classification. However, in medical images, where the location of the anatomical structures is often highly regular, this property may lead to incorrect predictions in regions with similar intensities when insufficient contextual information is considered. The organs observed in anatomical images tend to preserve shape and topology across patients. Nonetheless, the pixel-level predictions of most CNN architectures do not account for such higher-order topological properties, as discussed in [4].

Before the emergence of CNNs, other learning-based methods were popular for biomedical image segmentation (e.g. Random Forest (RF) [5]). When the amount of annotated data is small and insufficient for training deep CNNs, some of these classical methods are still in use. A popular strategy is to adopt patch-based methods, where handcrafted features are generated from image patches and then used to train a classifier. Such classifier will then make pixel-level predictions considering only the image area around the central pixel of the patch. This results in methods which are also agnostic to the global shape and topology of the anatomical structures.

In this work we introduce Post-DAE, a post-processing method which improves pixel-level predictions generated with arbitrary classifiers by incorporating shape and topological priors. We employ denoising autoencoders (DAEs) to learn compact and non-linear representations of anatomical structures, using only segmentation masks. The DAE is then used to bring potentially erroneous segmentation masks into an anatomically plausible space (see Figure 1).

Contributions. A preliminary version of this work was published in MICCAI 2019 [6]. In this extended version we provide additional experiments in the context of multi-class lung and heart segmentation of X-ray images, and left ventricle delineation in cardiac magnetic resonance (CMR) images.




Fig. 1: Post-DAE workflow: the method is implemented as a post-processing step which maps arbitrary segmentation masks in patch-based image segmentation, and show experimental results about the gain in performance when adding it explicitly to the description. In a more recent work [11], the authors propose the use of a spectral location parametrization specially adapted to brain volume coordinates to improve CNN-based image segmentation. However, albeit the fact that these strategies increase the accuracy of the resulting segmentation masks by incorporating contextual information, they do not incorporate explicit priors about shape and topology.

We also include a more complete and updated state-of-the-art section, a deeper analysis of how our method behaves in images with gross abnormalities and out-of-distribution cases, together with additional illustrations and extended discussion.

Our contributions can be summarized as follows: (i) we show that denoising autoencoders used as a post-processing step can improve the anatomical plausibility of unfeasible segmentation masks; (ii) we present results in the context of binary and multi-label segmentation of chest X-ray and CMR images, bench-marking with other classical post-processing method and showing the robustness of Post-DAE by improving segmentation masks coming from both, CNN and RF-based classifiers and (iii) we analyze the behaviour and limitations of our method when post-processing abnormal and out-of-distribution anatomical segmentation masks.

II. RELATED WORK

Multiple alternatives have been proposed to incorporate prior knowledge in medical image segmentation (see [7] for a complete review). One popular strategy to integrate priors about shape and topology into learning based segmentation methods is to modify the loss function used to train the model. A topology aware loss function which incorporates high-order regularization was proposed in [4]. In this case, a manually defined topological validity table specifies the relation between the structures of interest. This constitutes a disadvantage since such loss function must be constructed ad-hoc for every dataset. More similar to our method are those by [8], [9], where compact anatomical representations are learnt by means of autoencoders. Such global representation is incorporated into the loss function and used to encourage anatomical plausibility into the predicted segmentation masks. Differently from our method designed to improve arbitrary segmentations, the main disadvantage of [8], [9] is that they are specifically tailored to be used when training a CNN model. Therefore, they cannot be used to improve results obtained with other segmentation approaches like RF or even level-sets methods, which do not rely on an explicit training phase.

An alternative simple but effective approach is to increase the receptive field of the network, i.e. the area of the input image that influences a single prediction. Even if this strategy does not incorporate an explicit shape prior, it allows the network to consider high-order interactions between distant image regions, learning to encode certain global features about shape and topology. In this regard, [2] proposed to increase the receptive field of the CNN by means of a dual path focusing on a wider low resolution area of the input image. This increases the contextual information provided to the network, but also

augments the complexity of the segmentation model itself. Another approach to deal with the lack of spatial context in patch-based convolutional architectures is to augment the model including information about pixel location. In [10] the authors suggest that location information is a crucial discriminator

Following a similar idea, a fully connected conditional random field (CRF) is used in [2] as a post-processing step. However, as stated by [2], it is hard to find a global set of CRF parameters which can consistently improve the segmentation of all structures of interest. Moreover, there is no shape or topological prior incorporated in these models. Instead, these methods only operate under the hypothesis that pixels which are contiguous and exhibit similar aspect should belong to the same class. Even if valid in general, this assumption does not consider more complex priors like topological restrictions or convexity, which can be easily encoded in our post-processing methods.

Similar to the work of [13], our model can be trained using segmentation-only datasets which are not paired with image data (or data of the same image modality as the problem at hand). Our model is agnostic to image intensity (and thus insensitive to domain shift) and its formulation is much simpler than the one introduced in [13].

III. ANATOMICALLY PLAUSIBLE SEGMENTATION VIA POST-DAE

Given a dataset of anatomical segmentation masks (without paired intensity images) $\mathcal{D}_A = \{S_i^A\}_{0 \leq i \leq |\mathcal{D}_A|}$ we intend to learn a model that can bring segmentations $\mathcal{D}_P = \{S_i^P\}_{0 \leq i \leq |\mathcal{D}_P|}$ predicted by different classifiers P into an anatomically feasible space. We stress the fact that our method works as a post-processing step in the space of segmentations, making it independent of the predictor, image intensities and modality.

Denoising autoencoders (DAE) are neural networks designed to reconstruct a clean input from a corrupted version of it [14]. In this work, we propose to employ DAEs to recover anatomically plausible segmentation masks from corrupted or incorrect ones. The standard architecture for an autoencoder follows an encoder-decoder scheme (see the Supplementary Material for a detailed description of the architecture used in this work). The encoder $f_{enc}(S_i)$ is a mapping function




Fig. 2: Predictions obtained with segmentation methods of several qualities: random forest and UNet trained for different number of epochs. We include examples for both binary (white images) and multi-class (color images) segmentation. for: (i) X-Ray or CMR image; (ii) segmentation mask predicted by each baseline method; (iii) segmentation mask after post-processing with a CRF; (iv) segmentation mask after post-processing with our Post-DAE; and (v) ground-truth expert segmentations.

that turns the input into a lower dimensional hidden encoding h . In our implementation, $f_{enc}(S_i)$ is composed of stacked convolutions, non-linearities and pooling layers. At the end, a fully connected layer concentrates all information into a low dimensional code h . Then, the decoder $f_{dec}(h)$ maps this code back to the original input dimensions by means of successive non-linearities and up-convolutions.

The model is called *denoising* autoencoder due to the fact that it is trained with noisy segmentations $\hat{S}_i = \phi(S_i)$, which are obtained by degrading the ground-truth segmentation masks with a degradation function ϕ . To minimize the reconstruction error of the predicted segmentations with respect to the ground-truth, we train the DAE using a loss function based on the Dice coefficient (DSC) (for an exhaustive description of the Dice loss please see [15]):

$$\mathcal{L}_{DAE}(S_i) = DSC(S_i, f_{dec}(f_{enc}(\phi(S_i)))) \quad (1)$$

The learnt encoding $h = f_{enc}(S_i)$ is forced to retain as much information as possible about the input. This is due to the bottleneck effect produced by the reduced dimensionality of h . In this context, minimizing the reconstruction loss amounts to maximizing a lower bound on the mutual information between input S_i and the learnt representation h [14].

A. Mask degradation strategy.

We simulate corrupted segmentations to train the DAE by artificially degrading the ground truth segmentation masks

S_i with the following random degradation functions $\phi(S_i)$ ¹: (i) We simulate over and under segmentation by adding and removing random geometric shapes (including polygons, lines and ellipses); (ii) erosion, dilation and other morphological operations with variable kernels are applied to perform minor mask alterations; (iii) mask borders are modified by random swapping of foreground-background labels in the pixels close to the organ boundaries. In addition, data augmentation was performed by randomly resizing the original masks.

B. Post-processing with denoising autoencoders.

The proposed method is rooted in the so-called manifold assumption [17], which states that natural high dimensional data (like anatomical segmentation masks) concentrate close to a non-linear low-dimensional manifold. We use the DAE to learn such anatomically plausible manifold. Then, given a segmentation mask S_i^P produced with an arbitrary predictor P (e.g. CNN or RF), we project it onto that manifold using f_{enc} and reconstruct the corresponding anatomically feasible mask with f_{dec} . Different from other approaches like [8], [9] which incorporate the anatomical priors during the segmentation network training, our method is agnostic to the training process of the original predictor. Since it is conceived as a post-processing step, segmentation masks produced with

¹Our code associated to Post-DAE, the degradation function and the UNet model is publicly available at the following Colab Python Notebook: https://colab.research.google.com/drive/1wLZLo81cINR_c-UJTpBV4fh-BMByAU8?usp=sharing

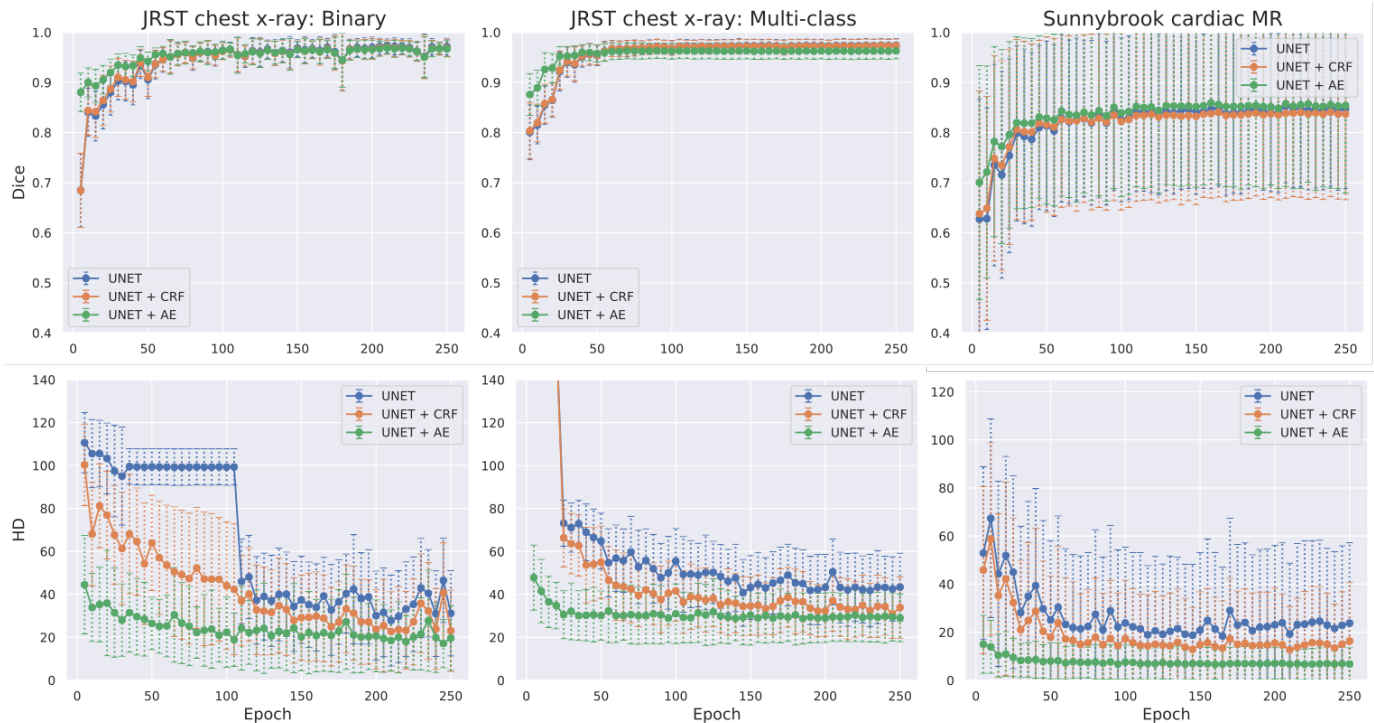


Fig. 3: Quantitative comparison between post-processing with Post-DAE and CRF [16]. We show mean and standard deviation for post-processing UNet predictions on the test fold at different training stages (from 5 epochs to convergence). First column shows the results for binary lung segmentation in JRST, the second one for multi-class (lung, heart) in JSRT and the third one for binary LV masks in Sunnybrook. We use Dice coefficient and Hausdorff distance (HD) to measure the segmentation quality. Note that initial low quality segmentations are improved both in terms of Dice and HD. For better initial segmentations, significant improvements in terms of HD are still obtained.

arbitrary methods can be improved using Post-DAE. Recent studies [18], [19] show that different autoencoders trained with healthy brain images (operating in the intensity domain) can be used to perform anomaly detection on pathological brain images, by just looking at the differences between the original pathological image and the one processed by the autoencoder. In the same spirit, our hypothesis (empirically validated with the experiments presented in the next section) is that those masks which are far from the anatomical space, will be mapped to a similar, but anatomically plausible segmentation. Meanwhile, masks which are anatomically correct, will suffer almost no modification, being mapped to themselves.

IV. EXPERIMENTAL SETTING

A. Database description.

We benchmark the proposed method in two different anatomical segmentation scenarios, including chest X-ray and cardiac magnetic resonance (CMR) images.

Chest X-ray dataset: in the context of lung and heart segmentation in X-Ray images, we used the Japanese Society of Radiological Technology (JSRT) database [20]. This is a public dataset with expert annotations composed of 247 PA thoracic X-ray images (2048x2048 pixels and spacing of 0.175mm x 0.175mm), which are downsampled to 1024x1024 in our experiments. Lungs and heart present high inter-subject variability, what makes the representation learning

task especially challenging. We divide the database in 3 folds considering 70% for training, 10% for validation and 20% for testing. The same folds were used to train the U-Net, random forest and Post-DAE methods. We did not apply image alignment for pre-processing.

Cardiac MR dataset: We used images from a version of the Sunnybrook Cardiac Dataset (SCD) [21] publicly available at <https://github.com/mshunshin/SegNetCMR>. It includes 45 cine-MR images (every image composed of 6 to 12 short-axis (SAX) 2D slices) captured at end-systole (ES) and end-diastole (ED) time points, with corresponding segmentation masks of the left ventricle (LV). The image resolution is 256x256, covering a field of view of 320 mm x 320 mm. We partitioned the dataset using the originally suggested train/test partition scheme (taking 35 images from the training fold for validation).

B. Post-processing with CRF.

The proposed method is compared with a standard post-processing strategy based on a fully connected CRF [16]. This method operates under the hypothesis that pixels which are contiguous and exhibit similar aspect should belong to the same class. We use an efficient implementation of a

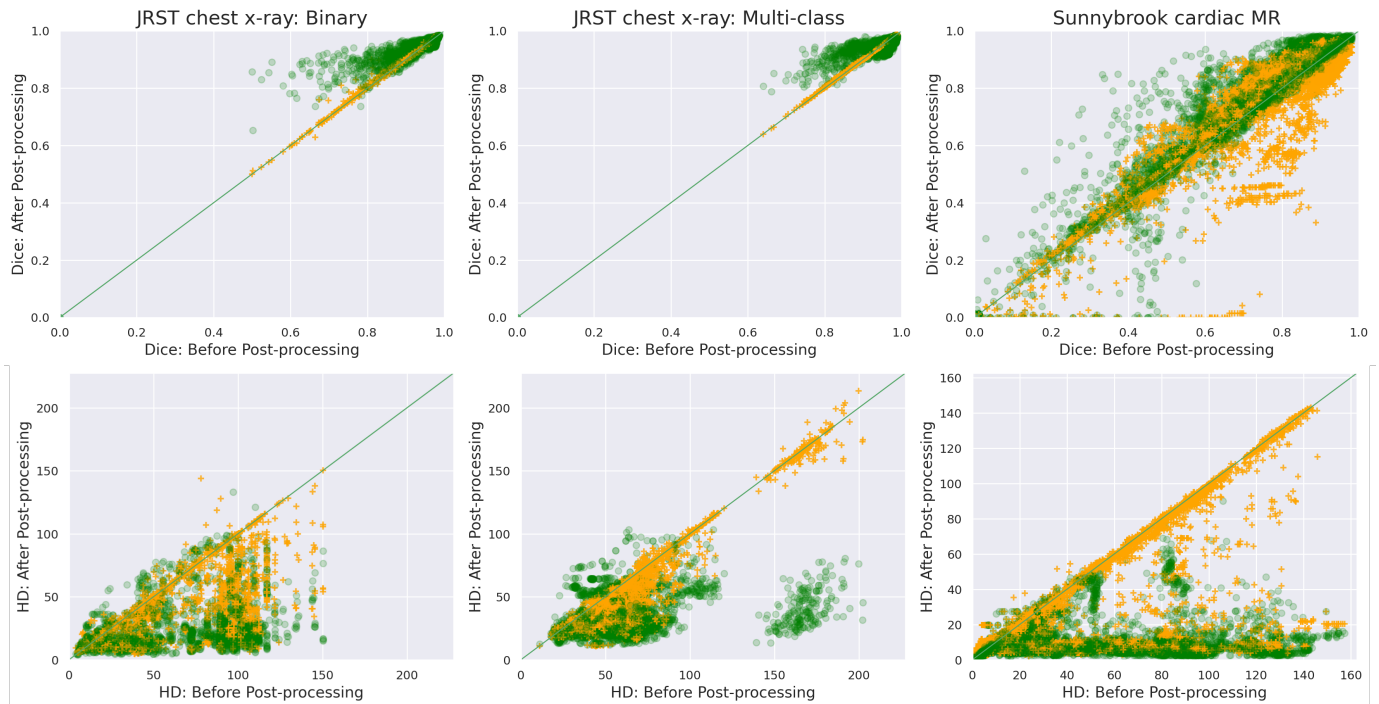


Fig. 4: Scatter plots comparing Dice (top) and HD (bottom) before and after post-processing with Post-DAE (green) and CRF (orange) for all samples in the previous study. We include segmentation masks generated with the UNet models trained from 5 to 250 epochs and random forest. First column shows the results for binary lung segmentation in JRST, the second one for multi-class (lung, heart) in JSRT and the third one for binary LV masks in Sunnybrook.

dense CRF.² Since the CRF formulation incorporates intensity information from the original images, the model parameters have to be re-adjusted whenever the image dataset is changed. In contrast, the proposed Post-DAE is agnostic to image intensity, and only needs to be trained once.

C. Training Post-DAE

Post-DAE is independent of the segmentation methods and it was trained separately from them. The model was implemented in Keras and trained for 150 epochs (the architecture and training details are included in the Supplementary Material). During training, we used the mask degradation strategy described in Section III-A. At test time, we took the segmentation masks generated by the baseline segmentation methods, and post-processed them by simply passing them through the DAE.

D. Baseline segmentation methods.

We trained binary and multi-class versions of two different segmentation models which produce masks of various qualities. For the X-ray images, we tackled binary (lungs vs background) and multi-class (lungs, heart, background)

segmentation. For the CMR images, we focus on binary LV segmentation. The RF classifier was trained using intensity and texture features. For the binary segmentation, we adopted a public implementation available online with default parameters³ which produces acceptable segmentation masks. It uses Haralick [22] features which are based on gray level co-occurrence in image patches. For the multi-label segmentation, we apply a different implementation⁴ which has a better performance for multi-label predictions and has been used in [23]. This RF variant leverages randomized offset boxes for calculating average intensity and intensity difference features efficiently via integral images. For the multi-label segmentation using RF we produced segmentations of various qualities by using different tree depths at test time. The second method is a CNN based on UNet architecture [1] (see the Supplementary Material for a detailed description of the architecture and the training parameters such as optimizer, learning rate, etc.). The UNet was implemented in Keras and trained in GPU using Dice loss function. To compare the effect of Post-DAE in different segmentation qualities, we save the UNet model every 5 epochs during training, and predict segmentation masks for the test fold using all these models. We compared post-processing with Post-DAE and CRF, reporting results for all these cases.

²We used the public implementation available at <https://github.com/lucasb-eyer/pydensecrf> with Potts compatibility function and hand-tuned parameters $\theta_\alpha = 17$, $\theta_\beta = 3$, $\theta_\gamma = 3$ for the X-ray images and $\theta_\alpha = 7$, $\theta_\beta = 3$, $\theta_\gamma = 3$ for the CMR images, chosen using the validation fold. See the website for more details about the aforementioned parameters.

³The source code and a complete description of the method is publicly available online at: https://github.com/dgriffiths3/ml_segmentation

⁴Publicly available at: <https://github.com/biomedica-mira/oak2>

TABLE I: Mean and standard deviation for post-processing random forest predictions. The numbers in bold indicates that Post-DAE outperforms the other methods (no post-processing and CRF) with statistical significance according to Wilcoxon test with Bonferroni correction.

Segmentations		JRST Chest X-ray						Sunnybrook Cardiac MR
		Multi-class					Binary	Binary
		Depth: 8	Depth: 12	Depth: 16	Depth: 20	Full model	Full model	Full model
Dice	RF	0.858 (0.042)	0.913 (0.033)	0.936 (0.029)	0.949 (0.029)	0.956 (0.028)	0.781 (0.070)	0.46 (0.24)
	RF + CRF	0.860 (0.042)	0.914 (0.032)	0.937 (0.029)	0.950 (0.029)	0.956 (0.027)	0.795 (0.074)	0.44 (0.25)
	RF + DAE	0.922 (0.024)	0.943 (0.020)	0.948 (0.018)	0.951 (0.018)	0.951 (0.019)	0.865 (0.056)	0.47 (0.25)
Hausdorff Distance (HD)	RF	102.26 (11.68)	96.00 (14.31)	88.70 (14.04)	77.45 (12.98)	72.28 (14.07)	91.41 (17.52)	27.73 (9.89)
	RF+CRF	101.17 (12.94)	92.97 (14.72)	81.26 (12.73)	74.51 (13.10)	67.29 (13.53)	80.45 (22.28)	26.87 (10.03)
	RF +DAE	63.73 (11.85)	60.72 (12.20)	62.47 (15.16)	62.95 (16.84)	60.69 (14.12)	32.01 (18.44)	23.60 (9.88)

V. RESULTS AND DISCUSSION

Figure 2 shows some visual examples while Table I and Figure 3 summarize the quantitative results obtained when post-processing segmentations produced by a RF classifier and a UNet. Our best results are in line with those obtained for other deep learning based state-of-the-art methods. For JSRT, recent works [24]–[27] report average Dice values for lung and heart ranging from 0.943 [27] to 0.965 [24]. For the Sunnybrook dataset, recent works [28]–[31] report average Dice ranging from 0.88 [28] to 0.93 [30] for LV segmentation. Both figures show the consistent improvement achieved when using Post-DAE as a post-processing step, specially in low quality segmentation masks like those obtained by the binary RF model, the multiclass RF considering incomplete tree depths and the UNet trained only for a few epochs. In these cases, substantial improvements are obtained in terms of Dice coefficient and Hausdorff distance (HD), by bringing the erroneous segmentation masks into an anatomically feasible space. In case of segmentations that are already of good quality (like multi-class RF or the UNet trained until convergence), Post-DAE significantly improves the HD, by erasing spurious segmentations that remain even in well trained models, like holes in the lung or small isolated blobs. When compared with CRF post-processing, Post-DAE significantly outperforms the baseline in the context of anatomical segmentation. In terms of running time, the CRF model takes 1.3 seconds while Post-DAE takes 0.76 seconds in a Intel i7-7700 CPU.

Scatter plots in Figure 4 show the change in terms of Dice (top) and HD (bottom) between initial segmentations before post-processing (x-axis) and after post-processing (y-axis), when comparing them with the ground-truth masks. In the Dice plots we observe how the green points tend to concentrate in the upper part of the diagonal, while orange crosses stick to it, indicating that Post-DAE improves the segmentations more than CRF. HD scatter plots should be read in the opposite way, i.e. the lower the better. Green points, corresponding to

Post-DAE, concentrate in the bottom part while the orange crosses (CRF) tend to be over them, indicating that Post-DAE outperforms CRF also in terms of HD.

We include additional experiments in the Supplementary Material showing that Post-DAE can be trained with unpaired segmentation masks. These segmentation masks could be annotated on different image modalities, come from a different dataset with the same image modality or even from segmentation-only datasets. This experiment highlights the fact that our method does not require image intensity information for training/testing, making it robust to domain shift.

A. Out-of-distribution segmentation masks and limitations

In this section we analyze the behaviour and limitations of Post-DAE when post-processing masks which are out-of-distribution. In this context, out-of-distribution cases could appear mainly due to two reasons: erroneous segmentations or pathological images.

In the first case, erroneous masks may be generated by a segmentation method with low performance. See for example the masks in Figure 2, obtained with the RF model or the UNet trained for only 5 or 10 epochs. These cases are represented in the scatter plots depicted in Figure 4 by the points with low Dice or high Hausdorff before post-processing. Post-DAE clearly improves erroneous segmentations in this scenario, increasing the Dice after post processing and/or reducing its Hausdorff distance. This improvement is explained by the way Post-DAE was trained: we degraded the ground truth masks by introducing similar errors (see Section III-A) and force the DAE to reconstruct anatomically plausible segmentation masks. Since the test images are anatomically plausible as well, mapping erroneous segmentations to realistic ones improves the results.

The second scenario is related to abnormal cases. Big occlusions or deformed organs, possibly due to manifestations




Fig. 5: Analysis for out-of-distribution segmentation masks presenting moderate and extreme abnormalities for tuberculosis patients from the Shenzhen database. See Section V.A for a complete discussion about this experiment.

of a particular disease or radiological occlusions, could make these masks look entirely different from the anatomically plausible cases. To analyze the behaviour of our model in this scenario, we employed a different chest X-ray dataset containing patients diagnosed with tuberculosis. This dataset is a subset of the original Shenzhen database [32], [33] formed by 38 X-ray images exhibiting tuberculosis manifestations. Every image was annotated by two expert radiologists following different approaches to delineate the lungs as discussed in [34]. The first approach was to segment only the air cavity part of the lung field, i.e. segmenting only the dark part and ignoring lighter areas covered with fluid. We call these the *air masks*. In the second approach, the annotator delineated the expected anatomy of the lungs, including occluded areas following a comparative approach by mirroring the normal lung field onto the abnormal one [34]. We call these *anatomy masks*. Figure 5 shows examples for both types of segmentation masks (rows (ii) and (iv)). Note that those corresponding to the air approach might present moderate or extreme abnormalities (e.g. missing complete parts of the lung). We applied the Post-DAE model to the *air masks* and analyzed its effect.

We used a Post-DAE model trained with the JSRT dataset, where the lung masks are mostly anatomically plausible since there are no big abnormalities or occlusions in the images. As expected, our method tends to map the air to the anatomy masks. However, note that when abnormalities are too extreme (see columns 3 and 4 in Figure 5) the real anatomy can not be completely reconstructed. We quantified this experiment by

measuring the Dice coefficient between the air and anatomy masks before and after post-processing the air masks with Post-DAE. The violin plots included in Figure 5 show that the post-processed air masks are significantly closer to the anatomy masks than the original ones, both in terms of Dice and Hausdorff metrics. This constitutes, at the same time, an advantage and a limitation of our approach: Post-DAE will transform the segmentation masks so that they look closer to the anatomically plausible ones used at training. These are important facts that must be considered when designing segmentation workflows which include Post-DAE. The same holds for problems different from anatomical segmentation. In scenarios like brain lesion or tumor segmentation, where shape and topology is not regular, the applicability of Post-DAE may be limited.

VI. CONCLUSIONS

In this work we have shown that denoising autoencoders can be used to render erroneous segmentations of different organs into anatomically plausible masks. Our method works as an independent post-processing step, allowing to incorporate anatomical priors into arbitrary segmentation methods. The provided experimental evaluation in the context of binary and multi-class anatomical segmentation of X-ray and CMR images indicates that our method can deal with a variety of anatomical structures in different image modalities. Moreover, Post-DAE does not use intensity information. Therefore, it can be trained with unpaired segmentation masks annotated on

different image modalities or coming from segmentation-only datasets, making the method robust to domain shift. Post-DAE can be easily implemented, is fast at inference, can cope with arbitrary shape priors and is independent of the image modality and segmentation method. In the future, we plan to explore the use of Post-DAE in the context of lesion segmentation [35], where the regions of interest are not as regular as anatomical structures.

ACKNOWLEDGMENTS

We thank Alexandros Karargyris, Sema Candemir and Stefan Jaeger for sharing the segmentation masks used in the out-of-distribution experiments.

REFERENCES

- [1] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. of MICCAI*, 2015.
- [2] K. Kamnitsas *et al.*, "Efficient multi-scale 3d CNN with fully connected CRF for accurate brain lesion segmentation," *Medical Image Analysis*, vol. 36, pp. 61–78, 2017.
- [3] M. Shakeri *et al.*, "Sub-cortical brain structure segmentation using F-CNN's," in *Proc. of ISBI*, 2016.
- [4] A. BenTaieb and G. Hamarneh, "Topology aware fully convolutional networks for histology gland segmentation," in *Proc. of MICCAI*. Springer, 2016, pp. 460–468.
- [5] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [6] A. Larrazabal, C. Martinez, and E. Ferrante, "Anatomical priors for image segmentation via post-processing with denoising autoencoders," in *Proc. of MICCAI*, 2019.
- [7] M. S. Nosrati and G. Hamarneh, "Incorporating prior knowledge in medical image segmentation: a survey," *arXiv preprint arXiv:1607.01092*, 2016.
- [8] O. Oktay *et al.*, "Anatomically constrained neural networks (ACNNs): application to cardiac image enhancement and segmentation," *IEEE TMI*, vol. 37, no. 2, pp. 384–395, 2018.
- [9] H. Ravishankar, R. Venkataramani, S. Thiruvankadam, P. Sudhakar, and V. Vaidya, "Learning and incorporating shape models for semantic segmentation," in *Proc. of MICCAI*, 2017.
- [10] C. Wachinger, M. Brennan, G. Sharp, and P. Golland, "On the importance of location and features for the patch-based segmentation of parotid glands," in *MICCAI Workshop on Image-Guided Adaptive Radiation Therapy*, 2014.
- [11] C. Wachinger, M. Reuter, and T. Klein, "Deepnat: Deep convolutional neural network for segmenting neuroanatomy," *NeuroImage*, vol. 170, pp. 434–445, 2018.
- [12] N. Paragios, E. Ferrante, B. Glocker, N. Komodakis, S. Parisot, and E. I. Zacharaki, "(Hyper)-graphical models in biomedical image analysis," *Medical Image Analysis*, jun 2016.
- [13] A. V. Dalca, J. Guttag, and M. R. Sabuncu, "Anatomical priors in convolutional networks for unsupervised biomedical segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9290–9299.
- [14] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *JMLR*, vol. 11, pp. 3371–3408, 2010.
- [15] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation," in *Proc. of Fourth International Conference on 3D Vision (3DV)*, 2016.
- [16] P. Krähenbühl and V. Koltun, "Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials," in *Proc. of NIPS*, 2011.
- [17] O. Chapelle, B. Scholkopf, and A. Zien, *Semi-supervised learning*. MIT Press, 2009.
- [18] N. Pawlowski *et al.*, "Unsupervised lesion detection in brain ct using bayesian convolutional autoencoders," in *MIDL 2018*, 2018.
- [19] H. Uzunova, S. Schultz, H. Handels, and J. Ehrhardt, "Unsupervised pathology detection in medical images using conditional variational autoencoders," *International journal of computer assisted radiology and surgery*, vol. 14, no. 3, pp. 451–461, 2019.
- [20] J. 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*, vol. 174, no. 1, pp. 71–74, 2000.
- [21] P. Radau, Y. Lu, K. Connelly, G. Paul, A. Dick, and G. Wright, "Evaluation framework for algorithms segmenting short axis cardiac mri," *The MIDAS Journal*, vol. 49, 2009.
- [22] R. M. Haralick, K. Shanmugam, and I. H. Dinstein, "Textural features for image classification," *IEEE Transactions on systems, man, and cybernetics*, no. 6, pp. 610–621, 1973.
- [23] B. Glocker, D. Zikic, E. Konukoglu, D. R. Haynor, and A. Criminisi, "Vertebrae localization in pathological spine ct via dense classification from sparse annotations," in *MICCAI*. Springer, 2013, pp. 262–270.
- [24] M. Frid-Adar, A. Ben-Cohen, R. Amer, and H. Greenspan, "Improving the segmentation of anatomical structures in chest radiographs using u-net with an imagenet pre-trained encoder," in *Image Analysis for Moving Organ, Breast, and Thoracic Images*. Springer, 2018, pp. 159–168.
- [25] W. Dai, N. Dong, Z. Wang, X. Liang, H. Zhang, and E. P. Xing, "Scan: Structure correcting adversarial network for organ segmentation in chest x-rays," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, 2018, pp. 263–273.
- [26] A. A. Novikov, D. Lenis, D. Major, J. Hladuvka, M. Wimmer, and K. Buhler, "Fully convolutional architectures for multiclass segmentation in chest radiographs," *IEEE transactions on medical imaging*, vol. 37, no. 8, pp. 1865–1876, 2018.
- [27] L. Mansilla, D. H. Milone, and E. Ferrante, "Learning deformable registration of medical images with anatomical constraints," *Neural Networks*, vol. 124, pp. 269–279, 2020.
- [28] Q. Zheng, H. Delingette, N. Duchateau, and N. Ayache, "3-d consistent and robust segmentation of cardiac images by deep learning with spatial propagation," *IEEE transactions on medical imaging*, vol. 37, no. 9, pp. 2137–2148, 2018.
- [29] A. H. Curiale, F. D. Colavecchia, and G. Mato, "Automatic quantification of the lv function and mass: a deep learning approach for cardiovascular mri," *Elsevier CMPB*, vol. 169, pp. 37–50, 2019.
- [30] M. Chen, L. Fang, and H. Liu, "Fr-net: Focal loss constrained deep residual networks for segmentation of cardiac mri," in *ISBI 2019*. IEEE, 2019, pp. 764–767.
- [31] J. V. Stough, J. DiPalma, Z. Ma, B. K. Fornwalt, and C. M. Haggerty, "Ventricular segmentation and quantitative assessment in cardiac mr using convolutional neural networks," in *Medical Imaging 2018: Biomedical Applications in Molecular, Structural, and Functional Imaging*, vol. 10578. International Society for Optics and Photonics, 2018, p. 1057826.
- [32] S. Candemir *et al.*, "Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration," *IEEE TMI*, vol. 33, no. 2, pp. 577–590, 2013.
- [33] S. Jaeger *et al.*, "Automatic tuberculosis screening using chest radiographs," *IEEE TMI*, vol. 33, no. 2, pp. 233–245, 2013.
- [34] A. Karargyris *et al.*, "Combination of texture and shape features to detect pulmonary abnormalities in digital chest x-rays," *International journal of computer assisted radiology and surgery*, vol. 11, no. 1, pp. 99–106, 2016.
- [35] N. Roulet, D. F. Slezak, and E. Ferrante, "Joint learning of brain lesion and anatomy segmentation from heterogeneous datasets," in *MIDL*, 2019, pp. 401–413.
- [36] S. Jaeger, S. Candemir, S. Antani, Y.-X. J. Wang, P.-X. Lu, and G. Thoma, "Two public chest x-ray datasets for computer-aided screening of pulmonary diseases," *Quantitative imaging in medicine and surgery*, vol. 4, no. 6, p. 475, 2014.

SUPPLEMENTARY MATERIAL

A. UNet details

The UNet model (see Table II) receives a 1024x1024 gray image as input and was trained using the soft Dice loss [15], batch size of 4, Adam optimizer with learning rate 1e-5 and the other parameters as by Keras default. We used data augmentation including random rotations, shifts, zoom and shear. We also used dropout for regularization, including a dropout layer after layer L_5 with keep probability $p=0.5$. For the multi-class UNet we used categorical cross-entropy loss and changed the initial learning rate to 1e-4.

TABLE II: Detailed description of the UNet architecture used as baseline model segmentation

	Kernel	Stride	#Kernels		NonLin		
			Binary	MC	Binary	MC	
L1	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Max Pooling	(f:2,2)	(s:2,2)				
L2	Conv	(f:3,3)	(s:1,1)	(N:32)	(N:32)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:32)	(N:32)	ReLU	ReLU
	Max Pooling	(f:2,2)	(s:2,2)				
L3	Conv	(f:3,3)	(s:1,1)	(N:64)	(N:64)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:64)	(N:64)	ReLU	ReLU
	Max Pooling	(f:2,2)	(s:2,2)				
L4	Conv	(f:3,3)	(s:1,1)	(N:128)	(N:128)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:128)	(N:128)	ReLU	ReLU
	Max Pooling	(f:2,2)	(s:2,2)				
L5	Conv	(f:3,3)	(s:1,1)	(N:256)	(N:256)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:256)	(N:256)	ReLU	ReLU
L6	UpConv	(f:3,3)	(s:1,1)	(N:128)	(N:128)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:128)	(N:128)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:128)	(N:128)	ReLU	ReLU
L7	UpConv	(f:3,3)	(s:1,1)	(N:64)	(N:64)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:64)	(N:64)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:64)	(N:64)	ReLU	ReLU
L8	UpConv	(f:3,3)	(s:1,1)	(N:32)	(N:32)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:32)	(N:32)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:32)	(N:32)	ReLU	ReLU
L9	UpConv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
L10	Conv	(f:3,3)	(s:1,1)	(N:2)	(N:3)	ReLU	ReLU
	Conv	(f:1,1)	(s:1,1)	(N:1)	(N:3)	Sigmoid	SoftMax

B. Post-DAE

Post-DAE (see Table III) receives a 1024x1024 segmentation as input. The network was also trained to minimize the Dice loss function using Adam Optimizer. We used learning rate of 0.0001, batch size of 15 and 150 epochs. The multi-class Post-DAE implementation receives a one-hot encoded segmentation of size 1024x1024x3 segmentation as input. Because of memory restrictions, in this case we reduced the batch size to 8.

C. Additional experiments for segmentation-only datasets

We performed an extra experiment aiming to show that it is possible to use annotations from a different dataset to train Post-DAE. We used the DAE trained with JSRT database to post-process results obtained with the binary UNet for the Montgomery County X-ray Set [36], a different chest X-ray dataset with manual lung annotations. X-ray images in this dataset were acquired from the tuberculosis control program of the Department of Health and Human Services of Montgomery County, MD, USA. This set contains 138 posterior-anterior x-rays, which were divided in 100 for training, 14

TABLE III: Detailed architecture of the simple denoising auto encoder model used to implement the proposed Post-DAE.

	Kernel	Stride	#Kernels		NonLin		
			Binary	MC	Binary	MC	
L1	Conv	(f:3,3)	(s:2,2)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
L2	Conv	(f:3,3)	(s:2,2)	(N:32)	(N:32)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:32)	(N:32)	ReLU	ReLU
L3	Conv	(f:3,3)	(s:2,2)	(N:32)	(N:32)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:32)	(N:32)	ReLU	ReLU
L4	Conv	(f:3,3)	(s:2,2)	(N:32)	(N:32)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:32)	(N:32)	ReLU	ReLU
L5	Conv	(f:3,3)	(s:2,2)	(N:32)	(N:32)	ReLU	ReLU
L6	FC	-	-	(N:512)	(N:1024)	None	None
L6	FC	-	-	(N:1024)	(N:4096)	Relu	Relu
L8	UpConv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
L9	UpConv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
L10	UpConv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
L11	UpConv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
L12	UpConv	(f:3,3)	(s:1,1)	(N:16)	(N:16)	ReLU	ReLU
	Conv	(f:3,3)	(s:1,1)	(N:1)	(N:3)	Sigmoid	SoftMax

for validation and 24 for testing. The size of the X-rays is either 4020x4892 or 4892x4020 pixels. All images were downsampled to 1024x1024 in our experiments, padded with 0's to obtain a 1:1 aspect ratio and rigidly aligned. With this dataset, we trained two segmentation models: a Random Forest and a UNet architecture (saving its output every 5 epochs during training), predicting segmentation masks for the test fold using all these models. These masks were then post-processed using Post-DAE. Figure 6 shows the results for this experiment, where unpaired annotations coming from a different dataset are used to train Post-DAE. It can be observed how our method improves the segmentation quality even when the annotations used to train Post-DAE are coming from a different dataset.

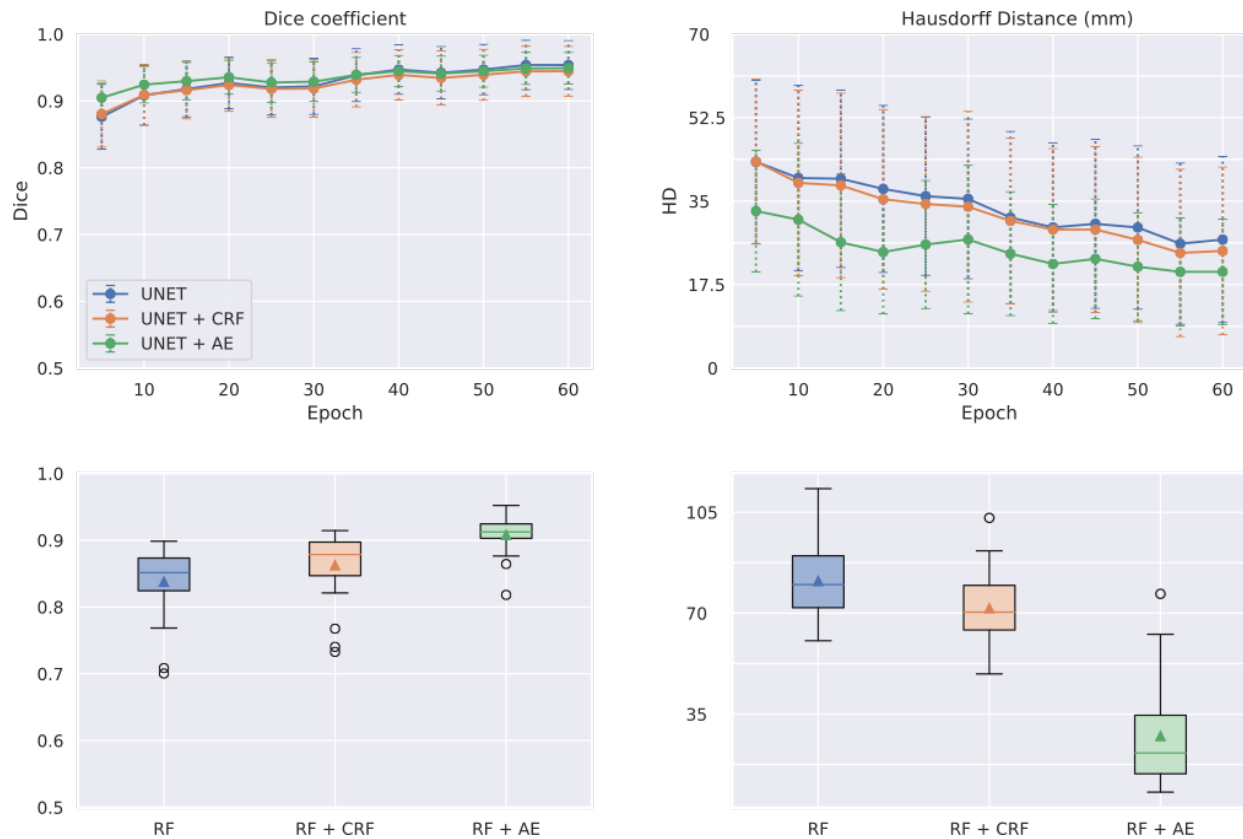


Fig. 6: Quantitative evaluation of the proposed method in a new data-set, which was not seen by the Post-DAE. Top row shows mean and standard deviation for post-processing UNet predictions on the test fold at different training stages (from 5 epochs to convergence). We use Dice coefficient and Hausdorff distance to measure the segmentation quality. Bottom row show results for post-processing the random forest predictions. The triangle in the box indicates the mean value.