

# Auxiliary Manifold Embedding for Fully Convolutional Networks

Christoph Baur<sup>1</sup>, Shadi Albarqouni<sup>1\*</sup>, and Nassir Navab<sup>1,2</sup>

<sup>1</sup> Computer Aided Medical Procedures (CAMP), Technical University of Munich, Munich, Germany

<sup>2</sup> Whiting School of Engineering, Johns Hopkins University, Baltimore, United States

**Abstract.** Deep learning usually requires large amounts of labeled training data, but annotating data is costly and tedious. The framework of semi-supervised learning provides the means to use both labeled data and arbitrary amounts of unlabeled data for training. Recently, semi-supervised deep learning has been intensively studied for standard CNN architectures. However, Fully Convolutional Networks (FCNs) set the state-of-the-art for many image segmentation tasks. To the best of our knowledge, there is no existing semi-supervised learning method for such FCNs yet. We lift the concept of auxiliary manifold embedding for semi-supervised learning to FCNs with the help of *Random Feature Embedding*. In our experiments on the challenging task of MS Lesion Segmentation, we leverage the proposed framework for the purpose of domain adaptation and report substantial improvements over the baseline model.

## 1 Introduction

In order to train deep neural networks, usually huge amounts of labeled data are necessary. In the medical field, however, labeled data is scarce as manual annotation is time-consuming and tedious. At the same time, when training models using a limited amount of labeled data, there is no guarantee that these models will generalize well on unseen data that is distributed slightly different. A prominent example in this context is Multiple Sclerosis lesion segmentation in MR images, which suffers from both a lack of ground-truth and distribution-shift across images from different devices [2]. However, vast amounts of unlabeled data can often be provided comparably easy. Semi-supervised learning provides the means to leverage both a limited amount of labeled data and arbitrary amounts of unlabeled data for training deep networks [11]. In recent years, various frameworks for semi-supervised deep learning have been proposed: In 2012, Weston et al. [11] presented a framework for artificial neural networks and shallow CNNs based on auxiliary manifold embedding. Using an additional embedding loss function attached to arbitrary hidden layers and graph adjacency among input samples, they forced the feature representations of neighbouring labeled and unlabeled samples to become more similar, leading to improved generalization. In

---

\* C. Baur and S. Albarqouni contributed equally to this work.

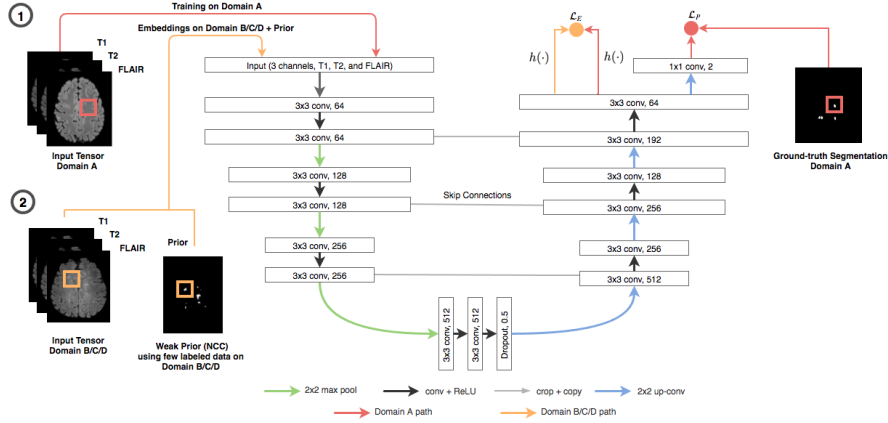


Fig. 1. Illustration of the semi-supervised deep learning framework.

2013, Lee et al. [5] also reported improved generalization when fine-tuning a model from predictions on unlabeled data using an entropy regularization loss. More recently, in 2015, Rasmus et al. [8] introduced the ladder network architecture for semi-supervised deep learning. And lately, Yang et al. [12] presented a framework also based on graph embeddings, with both transductive and inductive variants for shallow neural networks.

Even though these methods show promising results, all of them are tailored to classic CNN architectures and are often only examined on small scale computer vision datasets. In challenging problems such as biomedical image segmentation, Fully Convolutional Networks (FCNs) [6] are preferable as they are efficient and show the ability to learn context [7,4]. As far as we know, there is no existing semi-supervised learning method for such FCNs yet. In this paper, we lift the concept of auxiliary manifold embedding to FCNs with a, to the best of our knowledge, novel strategy called “Random Feature Embedding”. Subsequently, we successfully perform semi-supervised fine-tuning of FCNs for domain adaptation with our proposed embedding technique on the challenging task of Multiple Sclerosis (MS) lesion segmentation.

## 2 Methodology

Semi-supervised learning is based on the concept of training a model,  $f(\cdot)$ , using both labeled and unlabeled data  $\mathbf{X}_L$  and  $\mathbf{X}_U$ , respectively. In our framework for FCNs, training data has the form of  $\mathbf{D} = \{\mathbf{X}, \mathbf{Y}\}$ , where  $\mathbf{X} = \{\mathbf{X}_L \cup \mathbf{X}_U\} = \{\mathbf{x}_1, \dots, \mathbf{x}_{N_L}, \mathbf{x}_{N_L+1}, \dots, \mathbf{x}_{N_L+U}\} \in \mathbb{R}^{H \times W \times D \times N_{L+U}}$  are  $D$ -channel images, and  $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_{N_L}\} \in \mathbb{R}^{H \times W \times 1 \times N_L}$  are the corresponding label maps which are only available for the labeled data. Since we deal with FCNs, both images and label maps have the same dimensions,  $H \times W$ , allowing a distinct label per pixel.

## 2.1 Auxiliary Manifold Embedding

In our framework, to model  $f(\cdot)$ , we employ a modified version of the U-Net architecture [9] that processes images of arbitrary sizes and outputs a label map of the same size as the input (Fig. 1). We train the network to minimize the primary objective  $\mathcal{L}_P$ , i.e. the Dice-Loss [7], for our segmentation task (see Fig. 1) from labeled data only (Fig. 1, step 1). Simultaneously, to leverage the unlabeled data, we employ an auxiliary manifold embedding loss  $\mathcal{L}_E$  on the latent feature representations  $h(\cdot)$  of both  $\mathbf{X}_L$  and  $\mathbf{X}_U$  to minimize the discrepancy between similar inputs in the latent space. Thereby, similarity among  $h(\cdot)$  of unlabeled data is given by prior knowledge (Fig. 1, step 2). The overall objective function can be written using Lagrangian multipliers as:

$$\mathcal{L} = \mathcal{L}_P + \sum_l \lambda_l \cdot \mathcal{L}_{E_l} \quad (1)$$

where  $\lambda_l$  is the regularization parameter associated with the embedding loss  $E_l$  at hidden layer  $l$ . Typically, this embedding loss function aims at minimizing the distance among latent representations of similar  $h^l(\mathbf{x}_i)$  and  $h^l(\mathbf{x}_j)$  of neighboring data samples  $\mathbf{x}_i$  and  $\mathbf{x}_j$ , and otherwise tries to push them apart if their distance is within a margin  $m$ :

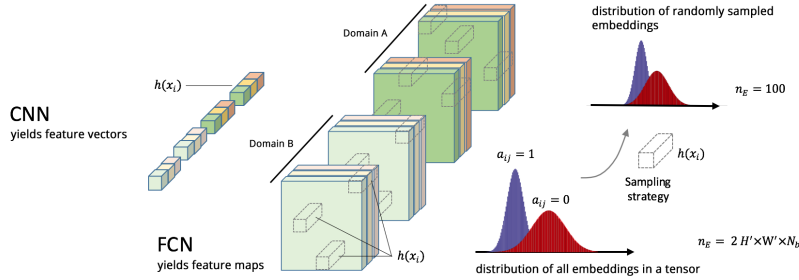
$$\mathcal{L}_{E_l}(\mathbf{X}, \mathbf{A}) = \sum_i^{n_E} \sum_j^{n_E} \begin{cases} d(h^l(\mathbf{x}_i), h^l(\mathbf{x}_j)), & \text{if } a_{ij} = 1 \\ \max(0, m - d(h^l(\mathbf{x}_i), h^l(\mathbf{x}_j))), & \text{if } a_{ij} = 0 \end{cases} \quad (2)$$

Thereby,  $\mathbf{A} \in \mathbb{R}^{n_E \times n_E}$  is an adjacency matrix between all embedding samples  $n_E$  within a training batch, and  $d(\cdot, \cdot) \in \mathbb{R}^1$  is an arbitrary distance metric measuring the distance between two latent representations. Unlike the typical  $\ell_2$ -norm distance employed in [11], we opt for the angular cosine distance (ACD) for two reasons; first, it is naturally bounded between  $[0, 1]$ , hence limits the searching area for the marginal distance parameter  $m$ , and second, it shows superior performance on high-dimensional feature representations in deep architectures [3].

The definition of  $\mathbf{A}$  is left to the user and can be arbitrary. A useful notion of adjacency is for instance given by neighbouring frames in a video stream [11]. We either compute the coefficients of  $\mathbf{A}$  based on the label maps, or based on results from template matching with NCC, similar to [1].

## 2.2 Random Feature Embedding

In standard CNNs, the embedding loss is commonly attached to the fully connected layers, where a latent feature vector  $h(\mathbf{x}_i)$  represents a single input image  $\mathbf{x}_i$ . In the case of FCNs, however, an input of arbitrary size can produce a multi-channel feature map of arbitrary size. Since FCNs make predictions at the pixel level, meaningful embeddings can be obtained per pixel along the channel dimension (Fig. 2). However, sampling and comparing all  $h(\cdot)$  of all pixels in large



**Fig. 2.** For FCNs, we sample embeddings  $h(\mathbf{x}_i)$  of single pixels from feature maps along the channel dimension. Randomly sampled embeddings should be representative for the entire population.

images is computationally infeasible: a  $W' \times H' \times D \times N_b$  feature map tensor for a batch of size  $N_b$  will yield  $n_E = W' \cdot H' \cdot N_b$  and therefore  $n_E^2$  comparisons. This quickly becomes intractable to compute. Instead, we suggest to do *Random Feature Embedding* (RFE), where we randomly sample a limited number  $n_E$  of pixels from the feature maps of the current batch according to some sampling strategy discussed in the next section to limit the number of comparisons. The loss remains valid since we propagate back only the gradients of selected pixels.

*Sampling strategy.* Ideally, the distribution of randomly sampled embeddings should mimic the one of all embeddings (Fig. 2), while at the same time paying attention to the class distribution such that unwanted bias is not introduced to the model. Therefore, we investigate the following sampling strategies:

- *50/50 RFE:* For each class in the prior the same amount of embeddings is randomly extracted to represent  $h(\cdot)$  from different classes equally. For unbalanced classes, this might lead to oversampling of one class.
- *Distribution-Aware RFE:* Embeddings are sampled from the given training batch according to the ratio of negative and positive classes in the prior to preserve the actual class distribution. When classes are unbalanced and  $n_E$  is too small, this might lead to undersampling of one class.
- *80/20 RFE:* As a trade-off, embeddings can be randomly sampled from a predefined ratio of 80% background and 20% foreground pixels.

### 3 Experiments and Results

Our experiments on MS lesion segmentation are motivated by the fact that existing automatic segmentation methods often fail to generalize well to MRI data acquired with different devices [2]. In this context, we leverage our semi-supervised learning framework for domain adaptation, i.e. we try to improve generalization of a baseline model by fine-tuning it with unlabeled data from other domains. Therefore, using an optimal prior for the adjacency matrix  $\mathbf{A}$ ,

**Table 1.** Overview of our MS lesion data and our training / testing split

Dataset	Domain	Patients (Train / Test)	Resolution	Scanner
MSSEG	A	3 / 2	144x512x512	3T Philips Ingenia
	B	3 / 2	144x512x512	1.5T Siemens Aera
	C	3 / 2	144x512x512	3T Siemens Verio
MSKRI	D	3 / 10	300x256x256	3T Philips Achieva

we first assess different sampling strategies, the impact of different numbers of embeddings as well as different distance measures for RFE. In succession, we utilize the most promising sampling strategy and distance measure together with a real prior for domain adaptation.

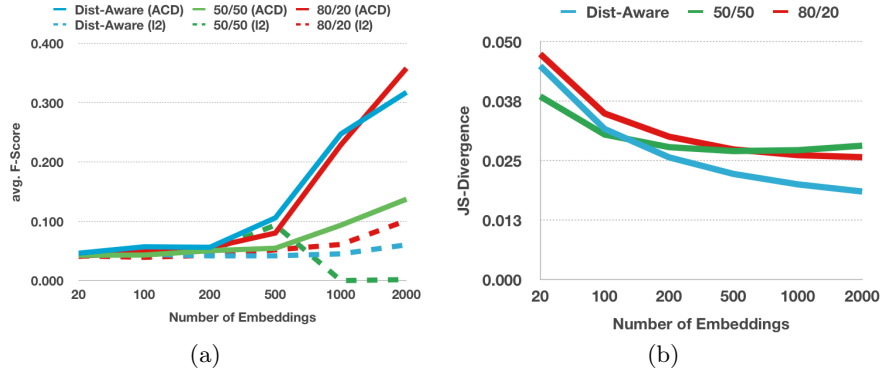
*Dataset.* Our MRI MS Lesion data is a combination of the publicly available MSSEG<sup>3</sup> training dataset and the non-publicly available MSKRI dataset (c.f. Table 1) acquired at the Neuroradiology department of Klinikum Rechts der Isar. For every patient there are coregistered T1, T2 and FLAIR volumes as well as a corresponding ground-truth segmentation volume. These are grouped into domains A, B, C and D based on the respective scanner they have been acquired with. For training and validation, we randomly crop patches around lesions from corresponding T1, T2 and FLAIR axial slices, resulting in 128x128x3 sized input tensors and corresponding binary label maps. Per domain, we thereby approx. crop 6000 patches. These are randomly split into training & validation sets according to a ratio of 7:3. Actual testing is performed on full slices of the testing volumes.

*Implementation* Our framework is built on top of MatConvNet [10]. All models were trained in batches of 12 and the learning rate of the primary objective fixed to 10e-6. For embedding with  $\ell_2$ , we set  $\lambda = 0.01$  and the margin parameter  $m = 1000$  (empirically chosen), for embedding with ACD we use  $\lambda = 1$  and  $m = 1$ , such that  $h(\cdot)$  from different classes become orthogonal.

### 3.1 Baseline Models

In order to measure the impact of our auxiliary manifold embedding, we first train so called lower bound and upper bound baseline models in a totally supervised fashion from labeled data only. We train the **Lower Bound Model**  $A_L$  from domain A training data for 50 epochs to obtain a model which produces decent segmentations on data from domain A, but does not generalize well to other domains. Further, we train so called **Upper Bound Models** by taking  $A_L$  after 35 epochs and fine-tuning it until epoch 50 using mixed, labeled training data from domain A and  $d \in [B, C, D]$ . We obtain three different models

<sup>3</sup> <https://portal.fli-iam.irisa.fr/msseg-challenge/overview>



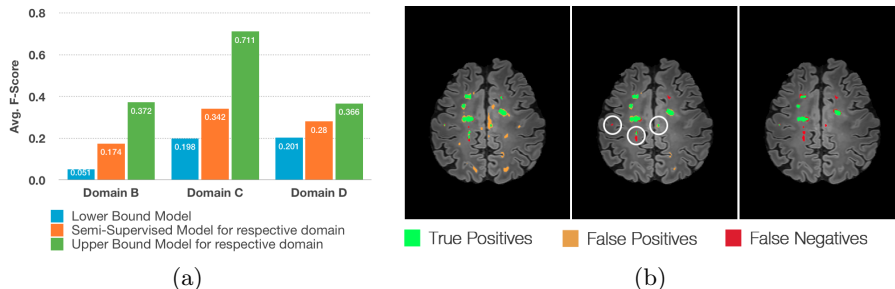
**Fig. 3.** a) Average F-Scores reported on domain B testing data for models trained with different settings and b) the impact of increasing  $n_E$  on the Jensen-Shannon divergence between randomly sampled embeddings and the space of all embeddings.

which should ideally be able to segment MS lesion volumes from domain A & B, A & C or A & D respectively.

### 3.2 Semi-Supervised Embedding

For the purpose of semi-supervised deep learning, we now assume there is labeled data from domain A, and multiple unlabeled or barely labeled MRI data from the other domains  $d \in [B, C, D]$ . All of the models trained in the following experiments originally build upon epoch 35 of the lower bound model  $A_L$ . One embedding loss is attached to the second-last convolutional layer of the network (Fig. 1). This choice is due to the fact that this particular layer produces feature maps at the same resolution as the input images, thus there is no need to downsample the prior required for embedding and thus no risk involved in losing heavily underrepresented, small lesion pixels (the lesion load within volumes is often less than 1%).

**Proof of Concept** We first assume a perfect prior and concentrate on investigating the best choice of sampling strategy, number of embeddings  $n_E$  and distance metric. We construct a perfect prior such that  $a_{i,j} = 1$  if the labels of  $h(x_i)$  and  $h(x_j)$  are equal and  $a_{i,j} = 0$  otherwise. Based on this prior, we fine-tune models for target domain B using i) 50/50, Distribution-Aware and 80/20 RFE with ii) the  $\ell_2$ -norm and the ACD as a distance metric and iii) different number of embeddings  $n_E \in \{20, 100, 200, 500, 1000, 2000\}$ , yielding a total of 36 different models. For fine-tuning in these proof-of-concept experiments, we use only a subset of 200 images from domain A and target domain B each, rather than the full training set. Our results show that, as we ramp up  $n_E$ , the distribution of randomly sampled embeddings more and more resembles



**Fig. 4.** a) Comparing the lower bound model, the semi-supervised models fine-tuned with NCC & ACD, and the upper bound models on the respective target domain; b) The semi-supervised model for domain C (middle) produces much fewer FP than the lower bound model (left image) and is only slightly inferior to the respective upper bound model (right image).

the full distribution (Fig. 3(b)), which renders the random sampling generally valid. Moreover, with increasing  $n_E$ , we notice consistent segmentation improvements on target domain B when using ACD as a distance metric (Fig. 3(a)). In repeated experiments, we always obtain similar results. The improvements are most pronounced with the 80/20 sampling strategy. However, the  $\ell_2$ -norm performs poorly and seems unstable as the number of embeddings  $n_E$  increases. We believe this is because the  $\ell_2$ -norm penalizes the magnitudes of the vectors being compared, whereas ACD is scale-invariant and only constrains their direction.

**Real Prior** Motivated by previous results, we now train models with a real, noisy prior using ACD and 80/20 RFE. For a target domain  $d \in [B, C, D]$ , we obtain the prior by selecting the first labeled FLAIR training volume  $V_1$  of  $d$  and randomly extracting  $5 \times 5 \times 5$  voxel sized 3D templates around MS lesions. Using 30 different, randomly sampled 3D templates, we perform NCC on the remaining volumes  $V_2$  and  $V_3$  of the current domain  $d$ . For thresholding the template matching output, the same matching is applied to  $V_1$  itself. Using the geometric mean of the responses and the ground-truth labels, a threshold is chosen which maximizes the Dice-Coefficient for  $V_1$ . Using this noisy prior, we fine-tune models for domain B, C and D using approx. 4000 training patches from volumes  $V_2$  and  $V_3$  and all labeled training patches from domain A. We set  $n_E = 100$  to obtain an overall number of embeddings similar to the proof of concept experiment ( $n_E = 2000$ ), but with lower computational cost. The models show consistent improvements over the lower bound model  $A_L$  for all target domains (see. Fig. 4(a)). Visual inspection (Fig. 4(b)) reveals that the semi-supervised embedding seems to dramatically reduce the number of false positives (FP). Interestingly, it detects some lesions (encircled) where the upper bound model fails, but it is not able to spot very small lesions. This is probably due to the fact that embeddings of smaller lesions are more unlikely to be sampled.

## 4 Discussion and Conclusion

In summary, we presented the concept of auxiliary manifold embedding and successfully integrated it into a semi-supervised deep learning framework for FCNs. At the heart of this framework is *Random Feature Embedding*, a simple method for sampling feature representations which serve as a statistic for non-linear embedding. Our experiments on MS lesion segmentation revealed that the method can improve generalization capabilities of existing models when using ACD as a distance metric. Yet, there is a lot of room for follow-up investigations. For instance, in future work, the effect of attaching multiple embedding objectives at different layers of an FCN could be investigated. In general, the proposed method should be applied to other problems as well to reveal its full potential.

## 5 Acknowledgement

We thank our clinical partners, in particular Dr. med. Paul Eichinger and Dr. med. Benedikt Wiestler, from the Neuroradiology department of Klinikum Rechts der Isar for providing us with their MRI MS Lesion dataset.

## References

1. Róger Bermúdez-Chacón, Carlos Becker, Mathieu Salzmann, and Pascal Fua. Scalable unsupervised domain adaptation for electron microscopy. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 326–334. Springer, 2016.
2. Daniel García-Lorenzo, Simon Francis, Sridar Narayanan, Douglas L Arnold, and D Louis Collins. Review of automatic segmentation methods of multiple sclerosis white matter lesions on conventional magnetic resonance imaging. *Medical image analysis*, 17(1):1–18, 2013.
3. Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines vinod nair. 2010.
4. Konstantinos Kamnitsas, Christian Ledig, Virginia FJ Newcombe, Joanna P Simpson, Andrew D Kane, David K Menon, Daniel Rueckert, and Ben Glocker. Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *arXiv preprint arXiv:1603.05959*, 2016.
5. Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*, volume 3, 2013.
6. Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
7. Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. *arXiv preprint arXiv:1606.04797*, 2016.
8. Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *Advances in Neural Information Processing Systems*, pages 3546–3554, 2015.



9. Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
10. A. Vedaldi and K. Lenc. Matconvnet – convolutional neural networks for matlab. In *Proceeding of the ACM Int. Conf. on Multimedia*, 2015.
11. Jason Weston, Frédéric Ratle, Hossein Mobahi, and Ronan Collobert. Deep learning via semi-supervised embedding. In *Neural Networks: Tricks of the Trade*, pages 639–655. Springer, 2012.
12. Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. *arXiv preprint arXiv:1603.08861*, 2016.