



Unsupervised Domain Adaptation with Semantic Consistency Across Heterogeneous Modalities for MRI Prostate Lesion Segmentation

Eleni Chiou^{1,2(✉)}, Francesco Giganti^{3,4}, Shonit Punwani⁵, Iasonas Kokkinos², and Eleftheria Panagiotaki^{1,2}

¹ Centre for Medical Image Computing, UCL, London, UK
eleni.chiou.17@ucl.ac.uk

² Department of Computer Science, UCL, London, UK

³ Department of Radiology, UCLH NHS Foundation Trust, London, UK

⁴ Division of Surgery and Interventional Science, UCL, London, UK

⁵ Centre for Medical Imaging, Division of Medicine, UCL, London, UK

Abstract. Any novel medical imaging modality that differs from previous protocols e.g. in the number of imaging channels, introduces a new domain that is heterogeneous from previous ones. This common medical imaging scenario is rarely considered in the domain adaptation literature, which handles shifts across domains of the same dimensionality. In our work we rely on stochastic generative modeling to translate across two heterogeneous domains at pixel space and introduce two new loss functions that promote semantic consistency. Firstly, we introduce a semantic cycle-consistency loss in the source domain to ensure that the translation preserves the semantics. Secondly, we introduce a pseudo-labelling loss, where we translate target data to source, label them by a source-domain network, and use the generated pseudo-labels to supervise the target-domain network. Our results show that this allows us to extract systematically better representations for the target domain. In particular, we address the challenge of enhancing performance on VERDICT-MRI, an advanced diffusion-weighted imaging technique, by exploiting labeled mp-MRI data. When compared to several unsupervised domain adaptation approaches, our approach yields substantial improvements, that consistently carry over to the semi-supervised and supervised learning settings.

Keywords: Unsupervised domain adaptation · Pseudo-labeling · Entropy minimization · Lesion segmentation · VERDICT-MRI · mp-MRI

1 Introduction

Domain adaptation transfers knowledge from a label-rich ‘source’ domain to a label-scarce or unlabeled ‘target’ domain. This allows us to train robust




Fig. 1. Example of heterogeneous domains: standard DW-MRI (source domain) consists of 5 input channels (4 b-values and the ADC map) while VERDICT-MRI consists of 15 input channels (5 b-values in 3 orthogonal directions).

models targeted towards novel medical imaging modalities or acquisition protocols with scarce supervision, if any. Recent unsupervised domain adaptation methods typically achieve this either by minimizing the discrepancy of the feature and/or output space for the two domains or by learning a mapping between the two domains at the raw pixel space. Either way, these approaches usually consider moderate domain shifts where the dimensionality of the input feature-space between the source and the target domain is identical.

In this work we address the challenge of adapting across two heterogeneous domains where both the distribution and the dimensionality of the input features are different (Fig. 1). Inspired by [9], we rely on stochastic translation [12] to align the two domains at pixel-level; [9] shows that stochastic translation yields clear improvements in heterogeneous domain adaptation tasks compared to deterministic, CycleGAN-based [32] translation approaches. However, these improvements have been obtained with semi-supervised learning, where a few labeled target-domain images are available, whereas our goal is unsupervised domain adaptation. To this end we introduce a *semantic cycle-consistency* loss on the cycle-reconstructed source images; if a source image is translated to the target domain and then back to the source domain, we require that critical structures are preserved. We also introduce a *pseudo-labeling* loss that allows us to use the unlabeled target data to supervise the target-domain segmentation network. In particular we translate the target data to the source domain, predict their labels according to a pre-trained source-domain segmentation network and use the generated pseudo-labels to supervise the target-domain segmentation network. This allows us to use exclusively target-domain statistics and train highly discriminative models.

We demonstrate the effectiveness of our approach in prostate lesion segmentation and an advanced diffusion-weighted MRI (DW-MRI) method called VERDICT-MRI (Vascular, Extracellular and Restricted Diffusion for Cytometry in Tumours) [20, 21]. Compared to the naive DW-MRI from multiparametric

(mp)-MRI acquisitions, VERDICT-MRI has a richer acquisition protocol to probe the underlying microstructure and reveal changes in tissue features similar to histology. VERDICT-MRI has shown promising results in clinical settings, discriminating normal from malignant tissue [7, 20, 26] and identifying specific Gleason grades [15]. However, the limited amount of available labeled training data does not allow us to utilize data-driven approaches which could directly exploit the information in the raw VERDICT-MRI [6, 8]. On the other hand, labeled, large scale clinical mp-MRI datasets exist [1, 18]. In this work we exploit labeled mp-MRI data to train a segmentation network that performs well on VERDICT-MRI.

Related Work: Most domain adaptation methods align the two domains either by extracting domain-invariant features or by aligning the two domains at the raw pixel space. Ren et al. [23] and Kamnitsas et al. [16], rely on adversarial training to align the feature distributions between the source and the target domain for medical image classification and segmentation respectively. Pixel-level approaches [3, 9, 13, 30, 31], use GAN-based methods [12, 32] to align the source and the target domains at pixel level. Chen et al. [4] align simultaneously the two domains at pixel- and feature-level by utilizing adversarial training. Ouyang et al. [19] combine a variational autoencoder (VAE)-based feature prior matching and pixel-level adversarial training to learn a domain-invariant latent space which is exploited during segmentation. Similarly, [28] perform pixel-level adversarial training to extract content-only images and use them to train a segmentation model that operates well in both domains. Other studies exploit unlabeled target domain data during the discriminative training. Bateson et al. [2] and Guodong et al. [29] use entropy minimization on the prediction of target domain as an extra regularization while [17] propose a teacher-student framework to train a model using labeled and unlabeled target data as well as labeled source data.

2 Method

2.1 Problem Formulation

We consider the problem of domain adaptation in prostate lesion segmentation. Let $\mathcal{X}_S \subset \mathbb{R}^{H \times W \times C_S}$ be a set of N_S source images and $\mathcal{Y}_S \subset \{0, 1\}^{H, W}$ their segmentation masks. The sample $X_S \in \mathcal{X}_S$ is a $H \times W \times C_S$ image and the entry $Y_S^{(h, w)}$ of the mask Y_S provides the label of voxel (h, w) as a one-hot vector. Let also $\mathcal{X}_T \subset \mathbb{R}^{H \times W \times C_T}$ be a set of N_T unlabeled target images. Sample $X_T \in \mathcal{X}_T$ is an $H \times W \times C_T$ image.

Stochastic Translation with Semantic Cycle-Consistency Regularization

We rely on stochastic translation [9, 12] to learn the mapping between the two domains and introduce a semantic cycle-consistency loss to enforce the cross-domain mapping to preserve critical structures.

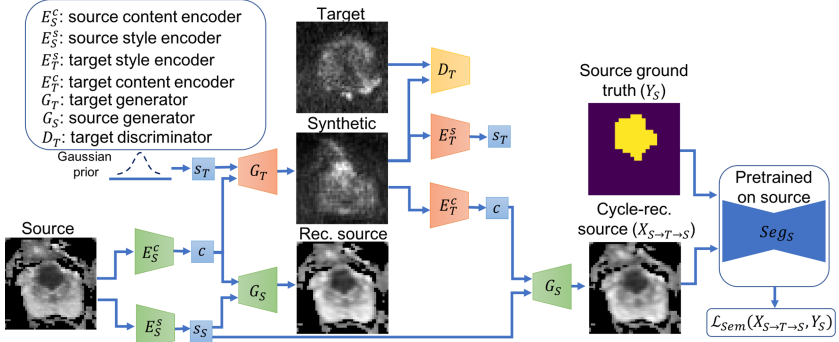


Fig. 2. We force a network for stochastic translation across domains to preserve semantics through a semantic segmentation-based loss. The image-to-image translation network translates source-domain images to the style of the target domain by combining a domain-invariant content code c with a random code s_T . We introduce a semantic cycle-consistency loss, \mathcal{L}_{Sem} , on the cycle-reconstructed images that ensures that the prostate lesions are successfully preserved.

The image-to-image translation network (Fig. 2) consists of content encoders E_S^c , E_T^c , style encoders E_S^s , E_T^s , generators G_S , G_T and domain discriminators D_S , D_T for both domains. The content encoders E_S^c , E_T^c extract a domain-invariant content code $c \in \mathcal{C}$ ($E_S^c : \mathcal{X}_S \rightarrow \mathcal{C}$, $E_T^c : \mathcal{X}_T \rightarrow \mathcal{C}$) while the style encoders E_S^s , E_T^s extract domain-specific style codes $s_S \in \mathcal{S}_S$ ($E_S^s : \mathcal{X}_S \rightarrow \mathcal{S}_S$) and $s_T \in \mathcal{S}_T$ ($E_T^s : \mathcal{X}_T \rightarrow \mathcal{S}_T$). Image-to-image translation is performed by combining the content code ($c = E_S^c(X_S)$) extracted from a given input ($X_S \in \mathcal{X}_S$) and a random style code s_T sampled from the target-style space. We note that the random style-code sampled from a Gaussian distribution represents structures that cannot be accounted by a deterministic mapping and results in one-to-many translation. We train the networks with a loss function consisting of domain adversarial, self-reconstruction, latent reconstruction and semantic cycle-consistency losses.

Domain Adversarial Loss

$$\mathcal{L}_{GAN}^T = \mathbb{E}_{c_S \sim \mathcal{C}, s_T \sim \mathcal{S}_T} [\log(1 - D_T(G_T(c_S, s_T)))] + \mathbb{E}_{X_T \sim \mathcal{X}_T} [\log(D_T(X_T))].$$

Self-reconstruction Loss

$$\mathcal{L}_{recon}^S = \mathbb{E}_{X_S \sim \mathcal{X}_S} [\|G_S(E_S^c(X_S), E_S^s(X_S)) - X_S\|_1].$$

Latent Reconstruction Loss

$$\begin{aligned} \mathcal{L}_{recon}^{cS} &= \mathbb{E}_{X_S \sim \mathcal{X}_S, s_T \sim \mathcal{S}_T} [\|E_T^c(G_T(E_S^c(X_S), s_T)) - E_S^c(X_S)\|_1], \\ \mathcal{L}_{recon}^{sT} &= \mathbb{E}_{X_S \sim \mathcal{X}_S, s_T \sim \mathcal{S}_T} [\|E_T^s(G_T(E_S^c(X_S), s_T)) - s_T\|_1]. \end{aligned}$$

Cycle-Consistency Loss

$$\mathcal{L}_{cyc}^S = \mathbb{E}_{X_S \sim \mathcal{X}_S, s_T \sim \mathcal{S}_T} [\|G_S(E_T^c(G_T(E_S^c(X_S), s_T)), E_S^s(X_S)) - X_S\|_1].$$

\mathcal{L}_{GAN}^S , \mathcal{L}_{recon}^T , \mathcal{L}_{recon}^{cT} , \mathcal{L}_{recon}^{sS} , \mathcal{L}_{cyc}^T are defined in a similar way.




Fig. 3. Pseudo-labeling through translation to the source domain: We translate the target data to the source domain and predict their pseudo-labels according to a pre-trained source-domain segmentation network Seg_S .

Semantic Cycle-Consistency Loss. Recent studies [3, 9, 13, 31] enforce semantic consistency between the real source and the synthetic target images by exploiting a target-domain segmentation network trained on a few available labeled target-domain images. However, in the unsupervised scenario, where there is no supervision available for the target domain, such approach is not feasible. To this end we introduce a semantic cycle-consistency loss or lesion segmentation loss on the cycle-reconstructed source images $X_{S→T→S}$; if a source image is translated to the target domain and then back to the source domain, we require that critical structures, corresponding to lesions, are preserved. The naive cycle-consistency loss, introduced in [32], penalizes inconsistencies in the entire image and may fail to preserve small structures corresponding to lesions. In contrast our semantic cycle-consistency loss penalizes inconsistencies in the label space enforcing the translation network to preserve the lesions. The semantic cycle-consistency loss is a soft generalization of the dice score given by

$$\mathcal{L}_{Sem} = - \frac{2 \sum_{h,w} P^{(h,w,1)} Y_S^{(h,w,1)}}{\sum_{h,w} (P^{(h,w,1)} + Y_S^{(h,w,1)})}, \quad (1)$$

where $P^{(h,w,1)}$ is the predictive probability of class 1 for voxel (h, w) provided by the pre-trained source network Seg_S . The full objective is given by

$$\begin{aligned} \min_{E_S^c, E_S^s, E_T^c, E_T^s, G_S, G_T} \max_{D_S, D_T} & \lambda_{GAN} (\mathcal{L}_{GAN}^S + \mathcal{L}_{GAN}^T) + \lambda_x (\mathcal{L}_{recon}^S + \mathcal{L}_{recon}^T) \\ & + \lambda_c (\mathcal{L}_{recon}^{cS} + \mathcal{L}_{recon}^{cT}) + \lambda_s (\mathcal{L}_{recon}^{sS} + \mathcal{L}_{recon}^{sT}) + \lambda_{cyc} (\mathcal{L}_{cyc}^S + \mathcal{L}_{cyc}^T) \\ & + \lambda_{sem} \mathcal{L}_{sem}, \end{aligned} \quad (2)$$

where λ_{GAN} , λ_x , λ_c , λ_s , λ_{cyc} , λ_{sem} control the importance of each term.

Pseudo-Labeling Through Translation to the Source

We generate pseudo-labels for the target images by translating them to the source domain and predicting their labels according to the pre-trained source-domain segmentation network Seg_S , trained on the labeled source data (Fig. 3).

Given a synthetic source image $X_{T→S}$ and the segmentation network Seg_S we obtain a soft-segmentation map $P_{X_{T→S}} = Seg_S(X_{T→S})$, where each vector $P_{X_{T→S}}^{(h,w)}$ corresponds to a probability distribution over classes. Assuming that




Fig. 4. We use data that have exclusively target-domain statistics to train the target segmentation network (Seg_T). We translate the source data to the target domain and supervise Seg_T using the ground-truth segmentation masks. We also use target pseudo-labels to supervise Seg_T .

high-scoring pixel-wise predictions on synthetic source samples are correct, we obtain a segmentation mask \hat{Y}_T by selecting high-scoring pixels with a fixed threshold. Each entry $\hat{Y}_T^{(h,w)}$ can be either a discrete one-hot vector for high-scoring pixels or a zero-vector for low-scoring pixels. The pseudo-labeling configuration is defined as follows

$$\hat{Y}_T^{(h,w,c)} = \begin{cases} 1, & \text{if } c = \arg \max_c P_{X_{T \rightarrow S}}^{(h,w)} \text{ and } P_{X_{T \rightarrow S}}^{(h,w,c)} > \text{threshold} \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Segmentation Network

The target-domain segmentation network, Seg_T , is an encoder-decoder network [5, 24]. We supervise Seg_T using both the synthetic target images and the corresponding source labels and the real target images and their pseudo-labels (Fig. 4). Given an image X and its segmentation mask Y , the segmentation loss is defined as

$$\mathcal{L}_{Seg}(X, Y) = - \frac{2 \sum_{h,w} P^{(h,w,1)} Y^{(h,w,1)}}{\sum_{h,w} (P^{(h,w,1)} + Y^{(h,w,1)})}, \quad (4)$$

where $P^{(h,w,1)}$ is the predictive probability of class 1 for voxel (h, w) .

As in recent studies [2, 27], to further regularize the network on the target-domain data for which we have not obtained pseudo-labels, we apply entropy-based regularization. The loss \mathcal{L}_{ent} is defined as follows

$$\mathcal{L}_{Ent}(X_T) = \sum_{h,w} \frac{-1}{\log C} \sum_{c=1}^C P_{X_T}^{(h,w,c)} \log P_{X_T}^{(h,w,c)} \quad (5)$$

The full objective is given by $\min_{Seg_T} \mathcal{L}_{Seg} + \mathcal{L}_{Ent}$.

2.2 Implementation Details

We implement our framework using Pytorch [22]. Our code is publicly available at <https://github.com/elchiou/SemanticConsistUDA>.

Image-to-Image Translation Network: The content encoders consist of convolutional layers and residual blocks followed by instance normalization [25]. The style encoders consist of convolutional layers followed by fully connected layers. The decoders include residual blocks followed by upsampling and convolutional layers. The residual blocks are followed by adaptive instance normalization (AdaIN) [11] layers to adjust the style of the output image. The discriminators consist of convolutional layers. For training we use Adam optimizer, a batch size of 32, a learning rate of 0.0001 and set losses weights to $\lambda_{GAN} = 1$, $\lambda_x = 10$, $\lambda_c = 1$, $\lambda_s = 1$, $\lambda_{sem} = 10$. We train the translation network for 50000 iterations.

Segmentation Network: The encoder of the segmentation network is a standard ResNet [10] consisting of convolutional layers while the decoder consists of upsampling and convolutional layers. For training we use stochastic gradient decent and a batch size of 32. We split the training set into 80% training and 20% validation to select the learning rate, the number of iterations and the threshold to perform pseudo-labeling.

3 Datasets

VERDICT-MRI: We use VERDICT-MRI data from 60 men. The DW-MRI was acquired with pulsed-gradient spin-echo sequence (PGSE) and an optimized imaging protocol for VERDICT prostate characterization with 5 b-values (90, 500, 1500, 2000, 3000 s/mm^2) in 3 orthogonal directions [14]. The DW-MRI sequence was acquired with a voxel size of $1.25 \times 1.25 \times 5$ mm^3 , 5 mm slice thickness, 14 slices, and field of view of 220×220 mm^2 and the images were reconstructed to a 176×176 matrix size. A radiologist contoured the lesions on VERDICT-MRI using mp-MRI for guidance.

DW-MRI from mp-MRI Acquisition: We use DW-MRI data from 80 patients from the ProstateX challenge dataset [18]. Three b-values were acquired (50, 400, 800 s/mm^2), and the ADC map and a b-value image at $b = 1400$ s/mm^2 were calculated by the scanner. The DW-MRI data were acquired with a single-shot echo planar imaging sequence with a voxel size of $2 \times 2 \times 3.6$ mm^3 , 3.6 mm slice thickness. Since the ProstateX dataset provides only the position of the lesion, a radiologist manually annotated the lesions on the ADC map using as reference the provided position.

4 Results

We evaluate the performance based on the average recall, precision, dice similarity coefficient (DSC), and average precision (AP) across 5-folds.

Table 1. Average recall, precision, dice similarity coefficient (DSC), and average precision (AP) across 5 folds. The results are given in mean (\pm std) format.

Model	Recall	Precision	DSC	AP
VERDICT-MRI (Oracle)	66.2 (8.1)	70.5 (9.9)	68.9 (9.2)	72.1 (10.4)
VERDICT-MRI + Synth (MUNIT + \mathcal{L}_{Sem})	71.1 (8.9)	72.5 (10.4)	72.1 (8.7)	76.7 (9.6)
mp-MRI + EntMin (ADVENT [27])	50.8 (12.3)	48.0 (11.4)	49.8 (13.0)	51.4 (13.9)
Synth (MUNIT)	51.5 (13.3)	60.6 (11.9)	53.6 (12.7)	60.2 (13.0)
Synth (MUNIT + \mathcal{L}_{Sem})	55.1 (13.9)	62.4 (12.8)	55.3 (10.9)	62.0 (13.4)
Synth (MUNIT + \mathcal{L}_{Sem}) + EntMin (ours)	54.7 (11.5)	69.2 (10.3)	57.1 (10.8)	63.4 (12.8)
Synth (MUNIT + \mathcal{L}_{Sem}) + PsLab (ours)	59.8 (10.1)	64.8 (11.1)	61.5 (10.3)	64.9 (10.1)
Synth (MUNIT + \mathcal{L}_{Sem}) + EntMin + PsLab (ours)	61.4 (9.9)	66.9 (10.7)	62.1 (9.8)	65.6 (10.9)




Fig. 5. Lesion segmentation results for two patients - the proposed approach performs well on the target domain despite the fact that it does not utilize any manual target annotations during training.




Fig. 6. Performance as we vary the percentage of labeled target data used for training. We observe that our method improves with more supervision and the improvements introduced by our method over the baseline of target-only training carry over all the way to the fully-supervised regime.

We compare our approach to several baselines. i) VERDICT-MRI: train using VERDICT-MRI only. ii) VERDICT-MRI + Synth (MUNIT + \mathcal{L}_{Sem}): train using real VERDICT-MRI and the synthetic VERDICT-MRI obtained from MUNIT with semantic cycle-consistency loss. iii) mp-MRI + EntMin (ADVENT [27]): train the model by minimizing the segmentation loss, $\mathcal{L}_{Seg}(X_S, Y_S)$, on the raw mp-MRI and the entropy loss, $\mathcal{L}_{Ent}(X_T)$, on VERDICT-MRI, an approach proposed in [2, 27, 29]. iv) Synth (MUNIT): use the naive MUNIT to map from source to target and train only on the synthetic data. v) Synth (MUNIT + \mathcal{L}_{Sem}): use MUNIT with semantic cycle-consistency loss to translate from source to target. vi) Synth (MUNIT + \mathcal{L}_{Sem}) + EntMin: use (v) and entropy-based regularization on VERDICT-MRI data. vii) Synth (MUNIT + \mathcal{L}_{Sem}) + PsLab: use (v) and pseudo-labels to train the segmentation network on real VERDICT-MRI. viii) Synth (MUNIT + \mathcal{L}_{Sem}) + EntMin

+ PsLab: use (vi) and pseudo-labels to train the segmentation network on real VERDICT-MRI.

We report the results in Table 1. We observe that the performance is poor when the segmentation network is trained on the mp-MRI and VERDICT-MRI data (mp-MRI + EntMin (ADVENT [27])). However, we observe that when we train the network with synthetic VERDICT-MRI and real VERDICT-MRI (Synth (MUNIT + \mathcal{L}_{Sem}) + EntMin) the performance improves. This justifies our assumption that pixel-level alignment is beneficial in cases where there is a large distribution shift. The performance further improves when we use pseudo-labels obtained from confident predictions (Synth (MUNIT + \mathcal{L}_{Sem}) + EntMin + PsLab). We also observe that compared to the naive MUNIT without the semantic cycle-consistency loss (Synth (MUNIT)) our approach (Synth (MUNIT + \mathcal{L}_{Sem})) performs better since it successfully preserves the lesions. When combining real and synthetic data (VERDICT-MRI + Synth (MUNIT + \mathcal{L}_{Sem})) to train the network in a fully-supervised manner we get better results compared to the oracle, where we use only the real VERDICT-MRI. In Fig. 5 we present lesion segmentation results produced by the different models for two patients. The results indicate that the proposed approach performs well despite the fact that it does not use any manual annotations during training.

So far we have considered only the unsupervised case. However, our approach can also be used in a semi-supervised setting. To evaluate the performance of our method when labeled target data is available, we perform additional experiments varying the percentage of labeled data; we use the pseudo-labels (PsLab) and entropy minimization (EntMin) for the unlabeled data. Figure 6 shows that the performance of our method improves as the percentage of real data increases and always outperforms the baseline that is trained only on the target domain.

5 Conclusion

In this work we propose a domain adaptation approach for lesion segmentation. Our approach relies on appearance alignment along with pseudo-labeling to train a target domain classifier using exclusively target domain statistics. We demonstrate the effectiveness of our approach for lesion segmentation on VERDICT-MRI which is an advanced imaging technique for cancer characterization. However, the proposed work is a general method that can be extended to other applications where there is a large distribution gap between the source and the target domain.

References

1. Ahmed, H.U., et al.: Diagnostic accuracy of multi-parametric MRI and TRUS biopsy in prostate cancer (PROMIS): a paired validating confirmatory study. *Lancet* **389**, 815–822 (2017)

2. Bateson, M., Kervadec, H., Dolz, J., Lombaert, H., Ben Ayed, I.: Source-relaxed domain adaptation for image segmentation. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 490–499. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_48
3. Cai, J., Zhang, Z., Cui, L., Zheng, Y., Yang, L.: Towards cross-modal organ translation and segmentation: a cycle and shape consistent generative adversarial network. *MedIA* **52**, 174–184 (2019)
4. Chen, C., Dou, Q., Chen, H., Qin, J., Heng, P.A.: Synergistic image and feature adaptation: towards cross-modality domain adaptation for medical image segmentation. In: AAAI (2019)
5. Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11211, pp. 833–851. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01234-2_49
6. Chiou, E., Giganti, F., Bonet-Carne, E., Punwani, S., Kokkinos, I., Panagiotaki, E.: Prostate cancer classification on VERDICT DW-MRI using convolutional neural networks. In: MLMI (2018)
7. Chiou, E., Giganti, F., Punwani, S., Kokkinos, I., Panagiotaki, E.: Automatic classification of benign and malignant prostate lesions: a comparison using VERDICT DW-MRI and ADC maps. In: ISMRM (2019)
8. Chiou, E., Giganti, F., Punwani, S., Kokkinos, I., Panagiotaki, E.: Domain adaptation for prostate lesion segmentation on VERDICT-MRI. In: ISMRM (2020)
9. Chiou, E., Giganti, F., Punwani, S., Kokkinos, I., Panagiotaki, E.: Harnessing uncertainty in domain adaptation for MRI prostate lesion segmentation. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 510–520. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_50
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
11. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV (2017)
12. Huang, X., Liu, M.-Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11207, pp. 179–196. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01219-9_11
13. Jiang, J., et al.: Tumor-aware, adversarial domain adaptation from CT to MRI for lung cancer segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11071, pp. 777–785. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_86
14. Johnston, E., Chan, R.W., Stevens, N., Atkinson, D., Punwani, S., Hawkes, D.J., Alexander, D.C.: Optimised VERDICT MRI protocol for prostate cancer characterisation. In: ISMRM (2015)
15. Johnston, E.W., et al.: VERDICT-MRI for prostate cancer: intracellular volume fraction versus apparent diffusion coefficient. *Radiology* **291**, 391–397 (2019)
16. Kamnitsas, K., et al.: Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In: IPMI (2017)
17. Li, K., Wang, S., Yu, L., Heng, P.-A.: Dual-teacher: integrating intra-domain and inter-domain teachers for annotation-efficient cardiac segmentation. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 418–427. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_41
18. Litjens, G., Debats, O., Barentsz, J., Karssemeijer, N., Huisman, H.: Computer-aided detection of prostate cancer in MRI. *TMI* **33**, 1083–1092 (2014)

19. Ouyang, C., Kamnitsas, K., Biffi, C., Duan, J., Rueckert, D.: Data efficient unsupervised domain adaptation for cross-modality image segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 669–677. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_74
20. Panagiotaki, E., et al.: Microstructural characterization of normal and malignant human prostate tissue with vascular, extracellular, and restricted diffusion for cytometry in tumours magnetic resonance imaging. *Investigate Radiol.* **50**, 218–227 (2015)
21. Panagiotaki, E., et al.: Noninvasive quantification of solid tumor microstructure using VERDICT MRI. *Cancer Res.* **74**, 1902–1912 (2014)
22. Paszke, A., et al.: Automatic differentiation in PyTorch. In: Autodiff Workshop, NIPS (2017)
23. Ren, J., Hacihaliloglu, I., Singer, E.A., Foran, D.J., Qi, X.: Adversarial domain adaptation for classification of prostate histopathology whole-slide images. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11071, pp. 201–209. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_23
24. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
25. Ulyanov, D., Vedaldi, A., Lempitsky, V.S.: Improved texture networks: maximizing quality and diversity in feed-forward stylization and texture synthesis. In: CVPR (2017)
26. Valindria, V., Palombo, M., Chiou, E., Singh, S., Punwani, S., Panagiotaki, E.: Synthetic q-space learning with deep regression networks for prostate cancer characterisation with verdict. In: ISBI (2021)
27. Vu, T.H., Jain, H., Bucher, M., Cord, M., Perez, P.: ADVENT: adversarial entropy minimization for domain adaptation in semantic segmentation. In: CVPR (2019)
28. Yang, J., Dvornik, N.C., Zhang, F., Chapiro, J., Lin, M.D., Duncan, J.S.: Unsupervised domain adaptation via disentangled representations: application to cross-modality liver segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 255–263. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_29
29. Zeng, G., et al.: Entropy guided unsupervised domain adaptation for cross-center hip cartilage segmentation from MRI. In: Martel, A.L., et al. (eds.) Medical Image Computing and Computer Assisted Intervention. MICCAI 2020. LNCS, vol. 12261, pp. 447–456. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_44
30. Zhang, Y., Miao, S., Mansi, T., Liao, R.: Task driven generative modeling for unsupervised domain adaptation: application to X-ray image segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11071, pp. 599–607. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_67
31. Zhang, Z., Yang, L., Zheng, Y.: Translating and segmenting multimodal medical volumes with cycle and shape consistency generative adversarial network. In: CVPR (2018)
32. Zhu, J., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)