Label-efficient Generalizable Deep Learning for Medical Image Segmentation

Ziyuan Zhao
Institute for Infocomm Research (I2R), A*STAR, Singapore
School of Computer Science and Engineering (SCSE), Nanyang Technological University, Singapore
zhaoz@i2r.a-star.edu.sg

Abstract

While deep learning hitherto has achieved considerable success in medical image segmentation, existing methods suffer from significant performance degradation under complex real-world clinical scenarios due to two main factors: **(1) Label scarcity:** reliance on large-scale well-annotated datasets, and **(2) Domain shift:** failure to apply well-trained models trained on old labeled datasets to new unlabeled datasets with different data distributions. In this paper, we introduce our recent works on label-efficient unsupervised domain adaptation to address both label scarcity and domain shift for cross-domain medical image analysis. We explore and advance various label-efficient learning paradigms with applications to medical image segmentation for improving model generalization and efficiency.

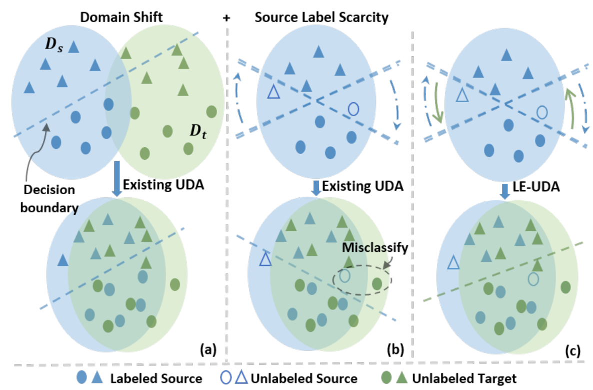

Figure 1: Schematic diagrams of label-efficient unsupervised domain adaptation.

1 Introduction

Segmentation plays a crucial role in biomedical image analysis by delineating anatomical structures and other regions of interest for various radiological tasks, such as diagnosis, treatment planning, and motion analysis [Hesamian *et al.*, 2019]. In recent years, deep learning [Zhou *et al.*, 2019; Zhou *et al.*, 2021a] has shown tremendous success in medical image segmentation [Ronneberger *et al.*, 2015; Çiçek *et al.*, 2016; Milletari *et al.*, 2016]. However, these impressive achievements always come with the price of massive pixel-accurate annotations, which are both costly and labor-intensive to obtain, thereby leading to the label scarcity problem in clinical practice [Tajbakhsh *et al.*, 2020]. To mitigate this problem, many recent efforts in medical image analysis have been devoted to developing methodologies beyond fully supervised learning techniques, such as self-supervised [Bai *et al.*, 2019; Li *et al.*, 2020a; Zhou *et al.*, 2021b; Zeng *et al.*, 2019; Zhu *et al.*, 2020], semi-supervised [Bai *et al.*, 2017; Zhao *et al.*, 2019; Zhao *et al.*, 2021b], and weakly-supervised learning [Rajchl *et al.*, 2016; Playout *et al.*, 2019; Belharbi *et al.*, 2021]. On the other hand, medical data can be collected from a variety of sources, including different medical centers, subject cohorts, imaging scanners, protocols, and even modalities, such as MRI and CT, resulting in different data distributions across domains [Gibson *et al.*, 2018; Guan and Liu, 2021]. As a result, the models trained on one domain, *i.e.*, source domain, usually fail in generalizing well to the data collected from unseen domains, *i.e.*, target domain, crippling the model performance significantly.

To improve cross-domain generalization in the existence of domain shift, one simple but effective method is to use newly annotated data from the target domain to fine-tune the model trained on the data from the source domain [Ghafoorian *et al.*, 2017]. However, labeling data from new domains inevitably involves additional time, expenses, and costs. To address the well-known domain shift phenomena in the absence of target annotations, significant progress has been made for domain adaptation and generalization [Zhu *et al.*, 2017; Russo *et al.*, 2018; Zhang *et al.*, 2018b; Bousmalis *et al.*, 2017; Zhao *et al.*, 2018; Ganin *et al.*, 2016; Tzeng *et al.*, 2017; Hoffman *et al.*, 2018; Chen *et al.*, 2019]. For unsupervised domain adaptation (UDA), most previous efforts assume that abundant source annotations are available and these methods aim to minimize domain discrepancy using utilizing labeled data from the source domain and unlabeled data from the target domain, thereby improving cross-domain adaptation performance. However, in clinical practice, such assumptions cannot always hold since the source domain may also suffer from label scarcity due to various reasons, *e.g.*, limited access, and expert knowledge. Source label scarcity can be detrimental to existing methods, causing a significant drop in performance. This is further illustrated in Fig. 1. In the absence of sufficient source annotations, the models cannot be well-trained on the source domain, resulting in a suboptimal decision boundary, which subsequently influences domain alignment. In the experiment section, we shall further demonstrate this circumstance. As such, this problem motivates us to study a practical yet challenging UDA scenario, *i.e.*, UDA under source label scarcity, where only scarce source annotations are available.

2 Problem Statement

Let there be two domains: source $\mathcal{D}_s$ and target $\mathcal{D}_t$, sharing the joint input and label space $\mathcal{X} \times \mathcal{Y}$. Source domain contains N labeled samples $\{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^N$ and M unlabeled samples $\{(\mathbf{x}_i^s)\}_{i=1}^M$, where N is much less than M under source label scarcity, while target domain includes P unlabeled samples $\{(\mathbf{x}_i^t)\}_{i=1}^P$. We aim to develop a segmentation model $\mathcal{F}_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ by leveraging available data and labels so that it can adapt well to the target domain.

3 Contributions

3.1 Label-efficient UDA

In our MICCAI 2021 work [Zhao *et al.*, 2021a] and IEEE TMI 2023 work [Zhao *et al.*, 2022a], we introduce this challenging but valuable realistic UDA setting where the source domain not only exhibits domain shift w.r.t. the target domain but also suffers from label scarcity and propose a generic label-efficient unsupervised domain adaptation framework for tackling both domain shift and source label scarcity. We first leverage the strong generative capacity of GANs [Goodfellow *et al.*, 2014] to facilitate image adaptation and generate complementary domain images to diversify the training distributions, thereby improving the model's generalization ability. To explore the rich information behind diverse domains, a dual-teacher network is constructed [Tarvainen and Valpola, 2017], in which the student model exploits intra-domain knowledge from source images and synthetic source images while also distilling inter-domain knowledge from diverse complementary domains. Furthermore, we propose a dual self-ensembling adversarial learning mechanism, in which discriminators are connected to multi-level predictions of the student and teacher networks under explicit self-ensembling consistency to derive generalizable representations implicitly. We carry out extensive studies on two challenging tasks, *i.e.*, cardiac substructure segmentation and abdominal multi-organ segmentation for bidirectional cross-modality medical image segmentation between CT and MRI. The results with various source label ratios show the effectiveness of the proposed method, receiving smaller performance drops than other UDA approaches.

3.2 Meta-hallucination

While rich synthetic data and semi-supervised learning can be leveraged for cross-modality medical image analysis, an extra domain generation step is required in advance, bringing additional computational cost, and image generation from independent networks has a limited potential to capture complex structural variations across different domains, *e.g.,* MRI and CT. In this regard, we pose a natural question: *How can we generate useful samples for quickly and reliably training a good cross-modality segmentation model with only a few source labels?* Recently, model-agnostic meta-learning ("learning to learn") [Finn *et al.*, 2017] to improve the learning model itself via the gradient descent process is flexible and independent of any model, leading to broad applications in few-shot learning [Zhang *et al.*, 2018a; Kiyasseh *et al.*, 2021] and domain generalization [Liu *et al.*, 2020; Khandelwal and Yushkevich, 2020]. Motivated by these observations, we argue that meta-learning can also enable the generator/hallucinator to "learn to hallucinate" meaningful images and obtain better segmentation models under few-shot UDA settings. In our MICCAI 2022 work [Zhao *et al.*, 2022b], we proposed a novel transformation-consistent meta-hallucination framework, in which we introduce a meta-learning episodic training strategy to optimize both the hallucination and segmentation models by explicitly simulating structural variances and domain shifts in the training process. Both the hallucination and segmentation models are trained concurrently in a collaborative manner to consistently improve few-shot cross-modality segmentation performance, where the hallucination model generates helpful samples for segmentation, whereas the segmentation model leverages transformation-consistent constraints and segmentation objectives to facilitate the hallucination process. Experimental results on cross-modality cardiac segmentation with 25% source labels in Table 1 have demonstrated the effectiveness of our method against domain shift and label scarcity.

Table 1: Segmentation performance of different approaches.

Method	Dice (%) ↑				
	AA	LAC	LVC	MYO	Average
Supervised-only	85.0	87.1	75.2	63.0	77.6
W/o adaptation	18.9	4.6	20.7	11.6	14.0
ADDA [Tzeng *et al.*, 2017]	35.5	4.2	2.1	36.9	19.7
CycleGAN [Zhu *et al.*, 2017]	43.7	49.8	43.2	23.1	40.0
SIFA [Chen *et al.*, 2020]	42.3	61.0	46.4	42.0	47.9
MT [Tarvainen and Valpola, 2017]	59.0	59.3	45.3	35.9	49.9
TCSM [Li *et al.*, 2020b]	65.3	62.7	50.9	38.3	54.3
ISTN [Robinson *et al.*, 2020]	34.0	61.0	47.1	32.9	43.8
VoxelMorph [Balakrishnan *et al.*, 2019]	57.6	67.2	41.1	35.7	50.4
MT-UDA [Zhao *et al.*, 2021a]	67.2	80.0	72.1	56.2	68.9
LE-UDA [Zhao *et al.*, 2022a]	**76.7**	**80.9**	67.3	58.4	70.8
Meta-Hal [Zhao *et al.*, 2022b]	75.6	75.1	**82.3**	**69.6**	**75.6**

4 Conclusion

This extended abstract summarises our recent research outcomes on label-efficient generalizable deep learning for medical image segmentation. We briefly introduce a new setting in unsupervised domain adaptation and review our proposed methods, including dual-domain self-ensemble learning and meta-learning-based data generation techniques, and results. The main contributions are that we study an underexplored but valuable UDA setting, and the proposed methods are generic, which can be integrated with different models and easily extended to various segmentation tasks and wider applications beyond segmentation.

References

[Bai *et al.*, 2017] Wenjia Bai, Ozan Oktay, Matthew Sinclair, Hideaki Suzuki, Martin Rajchl, Giacomo Tarroni, Ben Glocker, Andrew King, Paul M Matthews, and Daniel Rueckert. Semi-supervised learning for network-based cardiac mr image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 253–260. Springer, 2017.

[Bai *et al.*, 2019] Wenjia Bai, Chen Chen, Giacomo Tarroni, Jinming Duan, Florian Guitton, Steffen E Petersen, Yike Guo, Paul M Matthews, and Daniel Rueckert. Self-supervised learning for cardiac mr image segmentation by anatomical position prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 541–549. Springer, 2019.

[Balakrishnan *et al.*, 2019] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019.

[Belharbi *et al.*, 2021] Soufiane Belharbi, Jérôme Rony, Jose Dolz, Ismail Ben Ayed, Luke McCaffrey, and Eric Granger. Deep interpretable classification and weakly-supervised segmentation of histology images via max-min uncertainty. *IEEE Transactions on Medical Imaging*, 41(3):702–714, 2021.

[Bousmalis *et al.*, 2017] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3722–3731, 2017.

[Chen *et al.*, 2019] Cheng Chen, Qi Dou, Hao Chen, Jing Qin, and Pheng-Ann Heng. Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, pages 865–872, 2019.

[Chen *et al.*, 2020] Cheng Chen, Qi Dou, Hao Chen, Jing Qin, and Pheng Ann Heng. Unsupervised bidirectional cross-modality adaptation via deeply synergistic image and feature alignment for medical image segmentation. *IEEE transactions on medical imaging*, 39(7):2494–2505, 2020.

[Çiçek *et al.*, 2016] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.

[Finn *et al.*, 2017] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135. PMLR, 2017.

[Ganin *et al.*, 2016] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.

[Ghafoorian *et al.*, 2017] Mohsen Ghafoorian, Alireza Mehrtash, Tina Kapur, Nico Karssemeijer, Elena Marchiori, Mehran Pesteie, Charles RG Guttmann, Frank-Erik de Leeuw, Clare M Tempany, Bram van Ginneken, et al. Transfer learning for domain adaptation in mri: Application in brain lesion segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 516–524. Springer, 2017.

[Gibson *et al.*, 2018] Eli Gibson, Yipeng Hu, Nooshin Ghavami, Hashim U Ahmed, Caroline Moore, Mark Emberton, Henkjan J Huisman, and Dean C Barratt. Inter-site variability in prostate segmentation accuracy using deep learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 506–514. Springer, 2018.

[Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[Guan and Liu, 2021] Hao Guan and Mingxia Liu. Domain adaptation for medical image analysis: a survey. *IEEE Transactions on Biomedical Engineering*, 69(3):1173–1185, 2021.

[Hesamian *et al.*, 2019] Mohammad Hesam Hesamian, Wenjing Jia, Xiangjian He, and Paul Kennedy. Deep learning techniques for medical image segmentation: achievements and challenges. *Journal of digital imaging*, 32:582–596, 2019.

[Hoffman *et al.*, 2018] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018.

[Khandelwal and Yushkevich, 2020] Pulkit Khandelwal and Paul Yushkevich. Domain generalizer: A few-shot meta learning framework for domain generalization in medical imaging. In *Domain Adaptation and Representation Transfer, and Distributed and Collaborative Learning*, pages 73–84. Springer, 2020.

[Kiyasseh *et al.*, 2021] Dani Kiyasseh, Albert Swiston, Ronghua Chen, and Antong Chen. Segmentation of left atrial mr images via self-supervised semi-supervised meta-learning. In *MICCAI*, pages 13–24. Springer, 2021.

[Li *et al.*, 2020a] Xiaomeng Li, Mengyu Jia, Md Tauhidul Islam, Lequan Yu, and Lei Xing. Self-supervised feature learning via exploiting multi-modal data for retinal disease diagnosis. *IEEE Transactions on Medical Imaging*, 39(12):4023–4033, 2020.

[Li *et al.*, 2020b] Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, Lei Xing, and Pheng-Ann Heng.

Transformation-consistent self-ensembling model for semisupervised medical image segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2):523–534, 2020.

[Liu *et al.*, 2020] Quande Liu, Qi Dou, and Pheng-Ann Heng. Shape-aware meta-learning for generalizing prostate mri segmentation to unseen domains. In *MICCAI*, pages 475–485. Springer, 2020.

[Milletari *et al.*, 2016] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016.

[Playout *et al.*, 2019] Clément Playout, Renaud Duval, and Farida Cheriet. A novel weakly supervised multitask architecture for retinal lesions segmentation on fundus images. *IEEE transactions on medical imaging*, 38(10):2434–2444, 2019.

[Rajchl *et al.*, 2016] Martin Rajchl, Matthew CH Lee, Ozan Oktay, Konstantinos Kamnitsas, Jonathan Passerat-Palmbach, Wenjia Bai, Mellisa Damodaram, Mary A Rutherford, Joseph V Hajnal, Bernhard Kainz, et al. Deepcut: Object segmentation from bounding box annotations using convolutional neural networks. *IEEE transactions on medical imaging*, 36(2):674–683, 2016.

[Robinson *et al.*, 2020] Robert Robinson, Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, Marius de Groot, Ronald M Summers, Daniel Rueckert, and Ben Glocker. Image-level harmonization of multi-site data using image-and-spatial transformer networks. In *MICCAI*, pages 710–719. Springer, 2020.

[Ronneberger *et al.*, 2015] 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.

[Russo *et al.*, 2018] Paolo Russo, Fabio M Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8099–8108, 2018.

[Tajbakhsh *et al.*, 2020] Nima Tajbakhsh, Laura Jeyaseelan, Qian Li, Jeffrey N Chiang, Zhihao Wu, and Xiaowei Ding. Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Medical Image Analysis*, 63:101693, 2020.

[Tarvainen and Valpola, 2017] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017.

[Tzeng *et al.*, 2017] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017.

[Zeng *et al.*, 2019] Zeng Zeng, Yang Xulei, Yu Qiyun, Yao Meng, and Zhang Le. Sese-net: Self-supervised deep learning for segmentation. *Pattern Recognition Letters*, 128:23–29, 2019.

[Zhang *et al.*, 2018a] Ruixiang Zhang, Tong Che, Zoubin Ghahramani, Yoshua Bengio, and Yangqiu Song. Metagan: An adversarial approach to few-shot learning. *Advances in neural information processing systems*, 31, 2018.

[Zhang *et al.*, 2018b] Yue Zhang, Shun Miao, Tommaso Mansi, and Rui Liao. Task driven generative modeling for unsupervised domain adaptation: Application to x-ray image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 599–607. Springer, 2018.

[Zhao *et al.*, 2018] He Zhao, Huiqi Li, Sebastian Maurer-Stroh, Yuhong Guo, Qiuju Deng, and Li Cheng. Supervised segmentation of un-annotated retinal fundus images by synthesis. *IEEE transactions on medical imaging*, 38(1):46–56, 2018.

[Zhao *et al.*, 2019] Ziyuan Zhao, Xiaoman Zhang, Cen Chen, Wei Li, Songyou Peng, Jie Wang, Xulei Yang, Le Zhang, and Zeng Zeng. Semi-supervised self-taught deep learning for finger bones segmentation. In *2019 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, pages 1–4. IEEE, 2019.

[Zhao *et al.*, 2021a] Ziyuan Zhao, Kaixin Xu, Shumeng Li, Zeng Zeng, and Cuntai Guan. Mt-uda: Towards unsupervised cross-modality medical image segmentation with limited source labels. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 293–303. Springer, 2021.

[Zhao *et al.*, 2021b] Ziyuan Zhao, Zeng Zeng, Kaixin Xu, Cen Chen, and Cuntai Guan. Dsal: Deeply supervised active learning from strong and weak labelers for biomedical image segmentation. *IEEE Journal of Biomedical and Health Informatics*, 25(10):3744–3751, 2021.

[Zhao *et al.*, 2022a] Ziyuan Zhao, Fangcheng Zhou, Kaixin Xu, Zeng Zeng, Cuntai Guan, and S Kevin Zhou. Le-uda: Label-efficient unsupervised domain adaptation for medical image segmentation. *IEEE Transactions on Medical Imaging*, 2022.

[Zhao *et al.*, 2022b] Ziyuan Zhao, Fangcheng Zhou, Zeng Zeng, Cuntai Guan, and S Kevin Zhou. Meta-hallucinator: Towards few-shot cross-modality cardiac image segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part V*, pages 128–139. Springer, 2022.

[Zhou *et al.*, 2019] S Kevin Zhou, Daniel Rueckert, and Gabor Fichtinger. *Handbook of medical image computing and computer assisted intervention*. Academic Press, 2019.

[Zhou *et al.*, 2021a] S Kevin Zhou, Hayit Greenspan, Christos Davatzikos, James S Duncan, Bram Van Ginneken, Anant Madabhushi, Jerry L Prince, Daniel Rueckert, and Ronald M Summers. A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. *Proceedings of the IEEE*, 109(5):820–838, 2021.

[Zhou *et al.*, 2021b] Zongwei Zhou, Vatsal Sodha, Jiaxuan Pang, Michael B Gotway, and Jianming Liang. Models genesis. *Medical image analysis*, 67:101840, 2021.

[Zhu *et al.*, 2017] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.

[Zhu *et al.*, 2020] Jiuwen Zhu, Yuexiang Li, Yifan Hu, Kai Ma, S Kevin Zhou, and Yefeng Zheng. Rubik's cube+: A self-supervised feature learning framework for 3d medical image analysis. *Medical image analysis*, 64:101746, 2020.