

Shape-aware Meta-learning for Generalizing Prostate MRI Segmentation to Unseen Domains

Quande Liu¹, Qi Dou^{1,2}✉, and PhengAnn Heng^{1,3}

¹ Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong SAR, China

{qdlou, qdou, pheng}@cse.cuhk.edu.hk

² T Stone Robotics Institute, The Chinese University of Hong Kong, Hong Kong SAR, China

³ Guangdong Provincial Key Laboratory of Computer Vision and Virtual Reality Technology, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China

Abstract. Model generalization capacity at domain shift (*e.g.*, various imaging protocols and scanners) is crucial for deep learning methods in real-world clinical deployment. This paper tackles the challenging problem of domain generalization, *i.e.*, learning a model from multi-domain source data such that it can directly generalize to an unseen target domain. We present a novel shape-aware meta-learning scheme to improve the model generalization in prostate MRI segmentation. Our learning scheme roots in the gradient-based meta-learning, by explicitly simulating domain shift with virtual meta-train and meta-test during training. Importantly, considering the deficiencies encountered when applying a segmentation model to unseen domains (*i.e.*, incomplete shape and ambiguous boundary of the prediction masks), we further introduce two complementary loss objectives to enhance the meta-optimization, by particularly encouraging the *shape compactness* and *shape smoothness* of the segmentations under simulated domain shift. We evaluate our method on prostate MRI data from six different institutions with distribution shifts acquired from public datasets. Experimental results show that our approach outperforms many state-of-the-art generalization methods consistently across all six settings of unseen domains⁴.

Keywords: Domain generalization · Meta-learning · Prostate MRI segmentation.

1 Introduction

Deep learning methods have shown remarkable achievement in automated medical image segmentation [12,22,30]. However, the clinical deployment of existing models still suffer from the performance degradation under the distribution shifts across different clinical sites using various imaging protocols or scanner vendors.

⁴ Code and dataset are available at <https://github.com/liuquande/SAML>

Recently, many domain adaptation [5,13] and transfer learning methods [11,14] have been proposed to address this issue, while all of them require images from the target domain (labelled or unlabelled) for model re-training to some extent. In real-world situations, it would be time-consuming even impractical to collect data from each coming new target domain to adapt the model before deployment. Instead, learning a model from multiple source domains in a way such that it can directly generalize to an unseen target domain is of significant practical value. This challenging problem setting is *domain generalization (DG)*, in which no prior knowledge from the unseen target domain is available during training.

Among previous efforts towards the generalization problem [11,21,27], a naive practice of aggregating data from all source domains for training a deep model (called ‘DeepAll’ method) can already produce decent results serving as a strong baseline. It has also been widely used and validated in existing literature [4,8,26]. On top of DeepAll training, several studies added data augmentation techniques to improve the model generalization capability [29,24], assuming that the domain shift could be simulated by conducting extensive transformations to data of source domains. Performance improvements have been obtained on tasks of cardiac [4], prostate [29] and brain [24] MRI image segmentations, yet the choices of augmentation schemes tend to be tedious with task-dependence. Some other approaches have developed new network architectures to handle domain discrepancy [15,25]. Kour *et al.* [15] developed an unsupervised bayesian model to interpret the tissue information prior for the generalization in brain tissue segmentation. A set of approaches [1,23] also tried to learn domain invariant representations with feature space regularization by developing adversarial neural networks. Although achieving promising progress, these methods rely on network designs, which introduces extra parameters thus complicating the pure task model.

Model-agnostic meta-learning [10] is a recently proposed method for fast deep model adaptation, which has been successfully applied to address the domain generalization problem [2,7,17]. The meta-learning strategy is flexible with independence from the base network, as it fully makes use of the gradient descent process. However, existing DG methods mainly tackle image-level classification tasks with natural images, which are not suitable for the image segmentation task that requires pixel-wise dense predictions. An outstanding issue remaining to be explored is how to incorporate the shape-based regularization for the segmentation mask during learning, which is a distinctive point for medical image segmentation. In this regard, we aim to build on the advantages of gradient-based meta-learning, while further integrate shape-relevant characteristics to advance model generalization performance on unseen domains.

We present a novel **shape-aware meta-learning (SAML)** scheme for domain generalization on medical image segmentation. Our method roots in the meta-learning episodic training strategy, to promote robust optimization by simulating the domain shift with meta-train and meta-test sets during model training. Importantly, to address the specific deficiencies encountered when applying a learned segmentation model to unseen domains (i.e., incomplete shape and am-

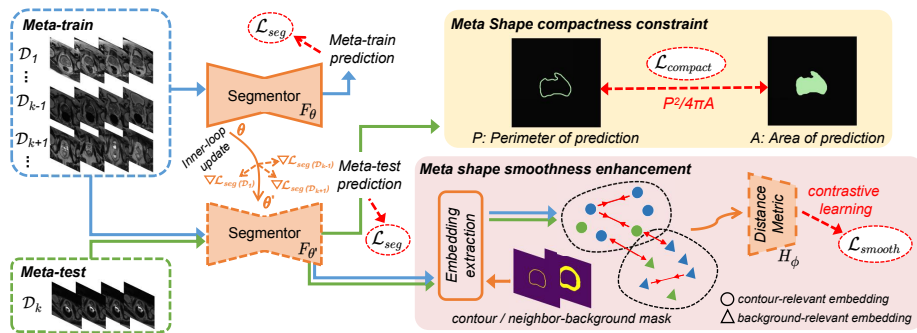


Fig. 1. Overview of our shape-aware meta-learning scheme. The source domains are randomly split into meta-train and meta-test to simulate the domain shift (Sec. 2.1). In meta-optimization: (1) we constrain the shape compactness in meta-test to encourage segmentations with complete shape (Sec. 2.2); (2) we promote the intra-class cohesion and inter-class separation between the contour and background embeddings regardless of domains, to enhance domain-invariance for robust boundary delineation (Sec. 2.3).

ambiguous boundary of the predictions), we further propose two complementary shape-aware loss functions to regularize the meta optimization process. First, we regularize the *shape compactness* of predictions for meta-test data, enforcing the model to well preserve the complete shape of segmentation masks in unseen domains. Second, we enhance the *shape smoothness* at boundary under domain shift, for which we design a novel objective to encourage domain-invariant contour embeddings in the latent space. We have extensively evaluated our method with the application of prostate MRI segmentation, using public data acquired from six different institutions with various imaging scanners and protocols. Experimental results validate that our approach outperforms many state-of-the-art methods on the challenging problem of domain generalization, as well as achieving consistent improvements for the prostate segmentation performance across all the six settings of unseen domains.

2 Method

Let $(\mathcal{X}, \mathcal{Y})$ denote the joint input and label space in a segmentation task, $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K\}$ be the set of K source domains. Each domain \mathcal{D}_k contains image-label pairs $\{(x_n^{(k)}, y_n^{(k)})\}_{n=1}^{N_k}$ sampled from domain distributions $(\mathcal{X}_k, \mathcal{Y})$, where N_k is the number of samples in the k -th domain. Our goal is to learn a segmentation model $F_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ using all source domains \mathcal{D} in a way such that it generalizes well to an unseen target domain \mathcal{D}_{tg} . Fig. 1 gives an overview of our proposed shape-aware meta-learning scheme, which we will detail in this section.

2.1 Gradient-based Meta-learning Scheme

The foundation of our learning scheme is the gradient-based meta-learning algorithm [17], to promote robust optimization by simulating the real-world domain shifts in the training process. Specifically, at each iteration, the source domains \mathcal{D} are randomly split into the meta-train \mathcal{D}_{tr} and meta-test \mathcal{D}_{te} sets of domains. The meta-learning can be divided into two steps. First, the model parameters θ are updated on data from meta-train \mathcal{D}_{tr} , using Dice segmentation loss \mathcal{L}_{seg} :

$$\theta' = \theta - \alpha \nabla_{\theta} \mathcal{L}_{seg}(\mathcal{D}_{tr}; \theta), \quad (1)$$

where α is the learning-rate for this inner-loop update. Second, we apply a meta-learning step, aiming to enforce the learning on meta-train \mathcal{D}_{tr} to further exhibit certain properties that we desire on unseen meta-test \mathcal{D}_{te} . Crucially, the meta-objective \mathcal{L}_{meta} to quantify these properties is computed with the updated parameters θ' , but optimized towards the original parameters θ . Intuitively, besides learning the segmentation task on meta-train \mathcal{D}_{tr} , such a training scheme further learns how to generalize at the simulated domain shift across meta-train \mathcal{D}_{tr} and meta-test \mathcal{D}_{te} . In other words, the model is optimized such that the parameter updates learned on virtual source domains \mathcal{D}_{tr} also improve the performance on the virtual target domains \mathcal{D}_{te} , regarding certain aspects in \mathcal{L}_{meta} .

In segmentation problems, we expect the model to well preserve the complete shape (compactness) and smooth boundary (smoothness) of the segmentations in unseen target domains. To achieve this, apart from the traditional segmentation loss \mathcal{L}_{seg} , we further introduce two complementary loss terms into our meta-objective, $\mathcal{L}_{meta} = \mathcal{L}_{seg} + \lambda_1 \mathcal{L}_{compact} + \lambda_2 \mathcal{L}_{smooth}$ (λ_1 and λ_2 are the weighting trade-offs), to explicitly impose the shape compactness and shape smoothness of the segmentation maps under domain shift for improving generalization performance.

2.2 Meta Shape Compactness Constraint

Traditional segmentation loss functions, *e.g.*, Dice loss and cross entropy loss, typically evaluate the pixel-wise accuracy, without a global constraint to the segmentation shape. Trained in that way, the model often fails to produce complete segmentations under distribution shift. Previous study have demonstrated that for the compact objects, constraining the shape compactness [9] is helpful to promote segmentations for complete shape, as an incomplete segmentation with irregular shape often corresponds to a worse compactness property.

Based on the observation that the prostate region generally presents a compact shape, and such shape prior is independent of observed domains, we propose to explicitly incorporate the compact shape constraint in the meta-objective \mathcal{L}_{meta} , for encouraging the segmentations to well preserve the shape completeness under domain shift. Specifically, we adopt the well-established Iso-Perimetric Quotient [19] measurement to quantify the shape compactness, whose definition is $C_{IPQ} = 4\pi A/P^2$, where P and A are the perimeter and area of the shape,

respectively. In our case, we define the shape compactness loss as the reciprocal form of this C_{IPQ} metric, and expend it in a pixel-wise manner as follows:

$$\mathcal{L}_{compact} = \frac{P^2}{4\pi A} = \frac{\sum_{i \in \Omega} \sqrt{(\nabla p_{u_i})^2 + (\nabla p_{v_i})^2 + \epsilon}}{4\pi(\sum_{i \in \Omega} |p_i| + \epsilon)}, \quad (2)$$

where p is the prediction probability map, Ω is the set of all pixels in the map; ∇p_{u_i} and ∇p_{v_i} are the probability gradients for each pixel i in direction of horizontal and vertical; ϵ ($1e^{-6}$ in our model) is a hyperparameter for computation stability. Overall, the perimeter length P is the sum of gradient magnitude over all pixels $i \in \Omega$; the area A is calculated as the sum of absolute value of map p .

Intuitively, minimizing this objective function encourages segmentation maps with complete shape, because an incomplete segmentation with irregular shape often presents a relatively smaller area A and relatively larger length P , leading to a higher loss value of $\mathcal{L}_{compact}$. Also note that we only impose $\mathcal{L}_{compact}$ in meta-test \mathcal{D}_{te} , as we expect the model to preserve the complete shape on unseen target images, rather than overfit the source data.

2.3 Meta Shape Smoothness Enhancement

In addition to promoting the complete segmentation shape, we further encourage smooth boundary delineation in unseen domains, by regularizing the model to capture domain-invariant contour-relevant and background-relevant embeddings that cluster regardless of domains. This is crucial, given the observation that performance drop at the cross-domain deployment mainly comes from the ambiguous boundary regions. In this regard, we propose a novel objective \mathcal{L}_{smooth} to enhance the boundary delineation, by explicitly promoting the intra-class cohesion and inter-class separation between the contour-relevant and background-relevant embeddings drawn from each sample across all domains \mathcal{D} .

Specifically, given an image $x_m \in \mathbb{R}^{H \times W \times 3}$ and its one-hot label y_m , we denote its activation map from layer l as $M_m^l \in \mathbb{R}^{H_l \times W_l \times C_l}$, and we interpolate M_m^l into $T_m^l \in \mathbb{R}^{H \times W \times C_l}$ using bilinear interpolation to keep consistency with the dimensions of y_m . To extract the contour-relevant embeddings $E_m^{con} \in \mathbb{R}^{C_l}$ and background-relevant embeddings $E_m^{bg} \in \mathbb{R}^{C_l}$, we first obtain the binary contour mask $c_m \in \mathbb{R}^{H \times W \times 1}$ and binary background mask $b_m \in \mathbb{R}^{H \times W \times 1}$ from y_m using morphological operation. Note that the mask b_m only samples background pixels around the boundary, since we expect to enhance the discriminativeness for pixels around boundary region. Then, the embeddings E_m^{con} and E_m^{bg} can be extracted from T_m^l by conducting weighted average operation over c_m and b_m :

$$E_m^{con} = \frac{\sum_{i \in \Omega} (T_m^l)_i \cdot (c_m)_i}{\sum_{i \in \Omega} (c_m)_i}, \quad E_m^{bg} = \frac{\sum_{i \in \Omega} (T_m^l)_i \cdot (b_m)_i}{\sum_{i \in \Omega} (b_m)_i}, \quad (3)$$

where Ω denotes the set of all pixels in T_m^l , the E_m^{con} and E_m^{bg} are single vectors, representing the contour and background-relevant representations extracted from the whole image x_m . In our implementation, activations from the last two deconvolutional layers are interpolated and concatenated to obtain the embeddings.

Next, we enhance the domain-invariance of E^{con} and E^{bg} in latent space, by encouraging embeddings’ intra-class cohesion and inter-class separation among samples from all source domains \mathcal{D} . Considering that imposing such regularization directly onto the network embeddings might be too strict to impede the convergence of \mathcal{L}_{seg} and $\mathcal{L}_{compact}$, we adopt the contrastive learning [6] to achieve this constraint. Specifically, an embedding network H_ϕ is introduced to project the features $E \in [E^{con}, E^{bg}]$ to a lower-dimensional space, then the distance is computed on the obtained feature vectors from network H_ϕ as $d_\phi(E_m, E_n) = \|H_\phi(E_m) - H_\phi(E_n)\|_2$, where the sample pair (m, n) are randomly drawn from all domains \mathcal{D} , as we expect to harmonize the embeddings space of \mathcal{D}_{te} and \mathcal{D}_{tr} to capture domain-invariant representations around the boundary region. Therefore in our model, the contrastive loss is defined as follows:

$$\ell_{contrastive}(m, n) = \begin{cases} d_\phi(E_m, E_n), & \text{if } \tau(E_m) = \tau(E_n) \\ (\max\{0, \zeta - d_\phi(E_m, E_n)\})^2, & \text{if } \tau(E_m) \neq \tau(E_n) \end{cases}. \quad (4)$$

where the function $\tau(E)$ indicates the class (1 for E being E^{con} , and 0 for E^{bg}) ζ is a pre-defined distance margin following the practice of metric learning (set as 10 in our model). The final objective \mathcal{L}_{smooth} is computed within mini-batch of q samples. We randomly employ either E^{con} or E^{bg} for each sample, and the \mathcal{L}_{smooth} is the average of $\ell_{contrastive}$ over all pairs of (m, n) embeddings:

$$\mathcal{L}_{smooth} = \sum_{m=1}^q \sum_{n=m+1}^q \ell_{contrastive}(m, n) / C(q, 2). \quad (5)$$

where $C(q, 2)$ is the number of combinations. Overall, all training objectives including $\mathcal{L}_{seg}(\mathcal{D}_{tr}; \theta)$ and $\mathcal{L}_{meta}(\mathcal{D}_{tr}, \mathcal{D}_{te}; \theta')$, are optimized together with respect to the original parameters θ . The \mathcal{L}_{smooth} is also optimized with respect to H_ϕ .

3 Experiments

Dataset and Evaluation Metric. We employ prostate T2-weighted MRI from 6 different data sources with distribution shift (cf. Table 1 for summary of their sample numbers and scanning protocols). Among these data, samples of Site A,B are from NCI-ISBI13 dataset [3]; samples of Site C are from I2CVB dataset [16]; samples of Site D,E,F are from PROMISE12 dataset [20]. Note that the NCI-ISBI13 and PROMISE12 actually include multiple data sources, hence we decompose them in our work. For pre-processing, we resized each sample to 384×384

Table 1. Details of our employed six different sites obtained from public datasets.

Dataset	Institution	Case num	Field strength(T)	Resolution(in/through plane)(mm)	Endorectal Coil	Manufacturer
Site A	RUNMC	30	3	0.6-0.625/3.6-4	Surface	Siemens
Site B	BMC	30	1.5	0.4/3	Endorectal	Philips
Site C	HCRUDB	19	3	0.67-0.79/1.25	No	Siemens
Site D	UCL	13	1.5 and 3	0.325-0.625/3-3.6	No	Siemens
Site E	BIDMC	12	3	0.25/2.2-3	Endorectal	GE
Site F	HK	12	1.5	0.625/3.6	Endorectal	Siemens

in axial plane, and normalized it to zero mean and unit variance. We then clip each sample to only preserve slices of prostate region for consistent objective segmentation regions across sites. We adopt Dice score (Dice) and Average Surface Distance (ASD) as the evaluation metric.

Implementation Details. We implement an adapted Mix-residual-UNet [28] as segmentation backbone. Due to the large variance on slice thickness among different sites, we employ the 2D architecture. The domains number of meta-train and meta-test were set as 2 and 1. The weights λ_1 and λ_2 were set as 1.0 and $5e^{-3}$. The embedding network H_ϕ composes of two fully connected layers with output sizes of 48 and 32. The segmentation network F_θ was trained using Adam optimizer and the learning rates for inner-loop update and meta optimization were both set as $1e^{-4}$. The network H_ϕ was also trained using Adam optimizer with learning rate of $1e^{-4}$. We trained 20K iterations with batch size of 5 for each source domain. For batch normalization layer, we use the statistics of testing data for feature normalization during inference for better generalization performance.

Comparison with State-of-the-art Generalization Methods. We implemented several state-of-the-art generalization methods for comparison, including a data-augmentation based method (BigAug) [29], a classifier regularization based method (Epi-FCR) [18], a latent space regularization method (LatReg) [1] and a meta-learning based method (MASF) [7]. In addition, we conducted experiments with ‘DeepAll’ baseline (i.e., aggregating data from all source domains for training a deep model) and ‘Intra-site’ setting (i.e., training and testing on the same domain, with some outlier cases excluded to provide general internal performance on each site). Following previous practice [7] for domain generalization, we adopt the leave-one-domain-out strategy, *i.e.*, training on $K-1$ domains and testing on the one left-out unseen target domain.

As listed in Table 2, DeepAll presents a strong performance, while the Epi-FCR with classifier regularization shows limited advantage over this baseline. The other approaches of LatReg, BigAug and MASF are more significantly better than DeepAll, with the meta-learning based method yielding the best results among them in our experiments. Notably, our approach (cf. the last row) achieves higher performance over all these state-of-the-art methods across all the six sites, and outperforms the DeepAll model by 2.15% on Dice and 0.60mm on ASD, demonstrating the capability of our shape-aware meta-learning scheme to deal with domain generalization problem. Moreover, Fig. 2 shows the generalization segmentation results of different methods on three typical cases from different unseen sites. We observe that our model with shape-relevant meta regularizers can well preserve the complete shape and smooth boundary for the segmentation in unseen domains, whereas other methods sometimes failed to do so.

We also report in Table 2 the cross-validation results conducted within each site, *i.e.*, Intra-site. Interestingly, we find that this result for site D/E/F is relatively lower than the other sites, and even worse than the baseline model. The reason would be that the sample numbers of these three sites are fewer than

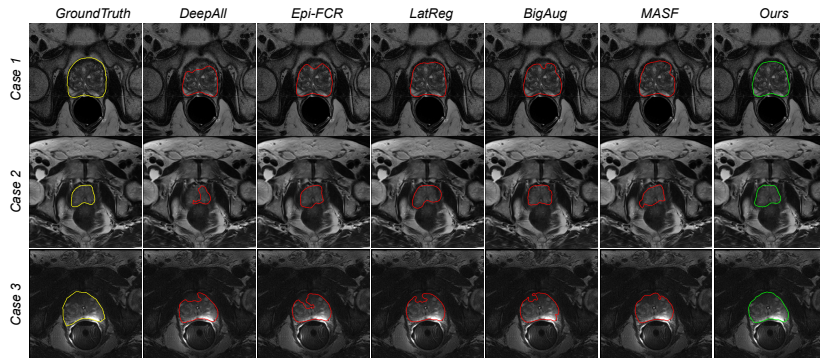


Fig. 2. Qualitative comparison on the generalization results of different methods, with three cases respectively drawn from different unseen domains.

Table 2. Generalization performance of various methods on Dice (%) and ASD (*mm*).

Method	Site A		Site B		Site C		Site D		Site E		Site F		Average	
Intra-site	89.27	1.41	88.17	1.35	88.29	1.56	83.23	3.21	83.67	2.93	85.43	1.91	86.34	2.06
DeepAll (baseline)	87.87	2.05	85.37	1.82	82.94	2.97	86.87	2.25	84.48	2.18	85.58	1.82	85.52	2.18
Epi-FCR [18]	88.35	1.97	85.83	1.73	82.56	2.99	86.97	2.05	85.03	1.89	85.66	1.76	85.74	2.07
LatReg [1]	88.17	1.95	86.65	1.53	83.37	2.91	87.27	2.12	84.68	1.93	86.28	1.65	86.07	2.01
BigAug [29]	88.62	1.70	86.22	1.56	83.76	2.72	87.35	1.98	85.53	1.90	85.83	1.75	86.21	1.93
MASF [7]	88.70	1.69	86.20	1.54	84.16	2.39	87.43	1.91	86.18	1.85	86.57	1.47	86.55	1.81
Plain meta-learning	88.55	1.87	85.92	1.61	83.60	2.52	87.52	1.86	85.39	1.89	86.49	1.63	86.24	1.90
+ $\mathcal{L}_{compact}$	89.08	1.61	87.11	1.49	84.02	2.47	87.96	1.64	86.23	1.80	87.19	1.32	86.93	1.72
+ \mathcal{L}_{smooth}	89.25	1.64	87.14	1.53	84.69	2.17	87.79	1.88	86.00	1.82	87.74	1.24	87.10	1.71
SAML (Ours)	89.66	1.38	87.53	1.46	84.43	2.07	88.67	1.56	87.37	1.77	88.34	1.22	87.67	1.58

the others, consequently intra-site training is ineffective with limited generalization capability. This observation reveals the important fact that, when a certain site suffers from severe data scarcity for model training, aggregating data from other sites (even with distribution shift) can be very helpful to obtain a qualified model. In addition, we also find that our method outperforms the Intra-site model in 4 out of 6 data sites, with superior overall performances on both Dice and ASD, which endorses the potential value of our approach in clinical practice.

Ablation Analysis. We first study the contribution of each key component in our model. As shown in Table 2, the plain meta-learning method only with \mathcal{L}_{seg} can already outperform the DeepAll baseline, leveraging the explicit simulation of domain shift for training. Adding shape compactness constraint into \mathcal{L}_{meta} yields improved Dice and ASD which are higher than MASF. Further incorporating \mathcal{L}_{smooth} (SAML) to encourage domain-invariant embeddings for pixels around the boundary, consistent performance improvements on all six sites are attained. Besides, simply constraining \mathcal{L}_{smooth} on pure meta-learning method (+ \mathcal{L}_{smooth}) also leads to improvements across sites.

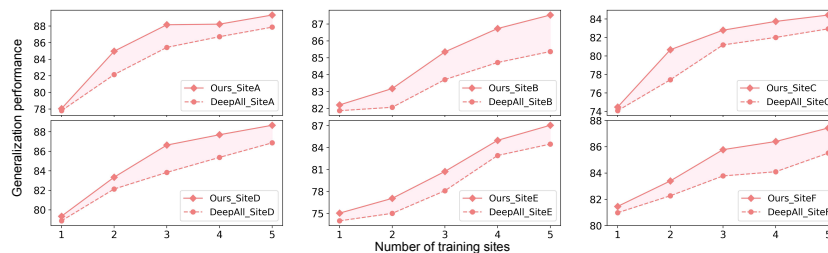


Fig. 3. Curves of generalization performance on unseen domain as the number of training source domain increases, using DeepAll method and our proposed approach.

We further investigate the influence of training domain numbers on the generalization performance of our approach and the DeepAll model. Fig. 3 illustrates how the segmentation performance on each unseen domain would change, as we gradually increase the number of source domains in range $[1, K - 1]$. Obviously, when a model is trained just with a single source domain, directly applying it to target domain receives unsatisfactory results. The generalization performance progresses as the training site number increases, indicating that aggregating wider data sources helps to cover a more comprehensive distribution. Notably, our approach consistently outperforms DeepAll across all numbers of training sites, confirming the stable efficacy of our proposed learning scheme.

4 Conclusion

We present a novel shape-aware meta-learning scheme to improve the model generalization in prostate MRI segmentation. On top of the meta-learning strategy, we introduce two complementary objectives to enhance the segmentation outputs on unseen domain by imposing the shape compactness and smoothness in meta-optimization. Extensive experiments demonstrate the effectiveness. To our best knowledge, this is the first work incorporating shape constraints with meta-learning for domain generalization in medical image segmentation. Our method can be extended to various segmentation scenarios that suffer from domain shift.

Acknowledgement. This work was supported in parts by the following grants: Key-Area Research and Development Program of Guangdong Province, China (2020B010165004), Hong Kong Innovation and Technology Fund (Project No. ITS/426/17FP), Hong Kong RGC TRS Project T42-409/18-R, and National Natural Science Foundation of China with Project No. U1813204.

References

1. Aslani, S., Murino, V., Dayan, M., Tam, R., Sona, D., Hamarneh, G.: Scanner invariant multiple sclerosis lesion segmentation from mri. In: ISBI. pp. 781–785. IEEE (2020)

2. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: *NeurIPS*. pp. 998–1008 (2018)
3. Bloch, N., Madabhushi, A., Huisman, H., Freymann, J., Kirby, J., Grauer, M., Enquobahrie, A., Jaffe, C., Clarke, L., Farahani, K.: Nci-isbi 2013 challenge: automated segmentation of prostate structures. *The Cancer Imaging Archive* **370** (2015)
4. Chen, C., Bai, W., Davies, R.H., Bhuva, A.N., Manisty, C., Moon, J.C., Aung, N., Lee, A.M., Sanghvi, M.M., Fung, K., et al.: Improving the generalizability of convolutional neural network-based segmentation on cmr images. *arXiv preprint arXiv:1907.01268* (2019)
5. Chen, C., Dou, Q., Chen, H., Qin, J., Heng, P.A.: Unsupervised bidirectional cross-modality adaptation via deeply synergistic image and feature alignment for medical image segmentation. *IEEE TMI* (2020)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709* (2020)
7. Dou, Q., de Castro, D.C., Kamnitsas, K., Glocker, B.: Domain generalization via model-agnostic learning of semantic features. In: *NeurIPS*. pp. 6450–6461 (2019)
8. Dou, Q., Liu, Q., Heng, P.A., Glocker, B.: Unpaired multi-modal segmentation via knowledge distillation. *IEEE TMI* (2020)
9. Fan, R., Jin, X., Wang, C.C.: Multiregion segmentation based on compact shape prior. *TASE* **12**(3), 1047–1058 (2014)
10. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. *ICML* (2017)
11. Gibson, E., Hu, Y., Ghavami, N., Ahmed, H.U., Moore, C., Emberton, M., Huisman, H.J., Barratt, D.C.: Inter-site variability in prostate segmentation accuracy using deep learning. In: *MICCAI*. pp. 506–514. Springer (2018)
12. Jia, H., Song, Y., Huang, H., Cai, W., Xia, Y.: Hd-net: Hybrid discriminative network for prostate segmentation in mr images. In: *MICCAI*. pp. 110–118. Springer (2019)
13. Kamnitsas, K., Baumgartner, C., Ledig, C., Newcombe, V., Simpson, J., Kane, A., Menon, D., Nori, A., Criminisi, A., Rueckert, D., et al.: Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In: *IPMI*. pp. 597–609. Springer (2017)
14. Karani, N., Chaitanya, K., Baumgartner, C., Konukoglu, E.: A lifelong learning approach to brain mr segmentation across scanners and protocols. In: *MICCAI*. pp. 476–484. Springer (2018)
15. Kouw, W.M., Ørting, S.N., Petersen, J., Pedersen, K.S., de Bruijne, M.: A cross-center smoothness prior for variational bayesian brain tissue segmentation. In: *IPMI*. pp. 360–371. Springer (2019)
16. Lemaitre, G., Martí, R., Freixenet, J., Vilanova, J.C., Walker, P.M., Meriaudeau, F.: Computer-aided detection and diagnosis for prostate cancer based on mono and multi-parametric mri: a review. *CBM* **60**, 8–31 (2015)
17. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Learning to generalize: Meta-learning for domain generalization. In: *AAAI* (2018)
18. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In: *ICCV*. pp. 1446–1455 (2019)
19. Li, W., Goodchild, M.F., Church, R.: An efficient measure of compactness for two-dimensional shapes and its application in regionalization problems. *IJGIS* **27**(6), 1227–1250 (2013)

20. Litjens, G., Toth, R., van de Ven, W., Hoeks, C., Kerkstra, S., van Ginneken, B., Vincent, G., Guillard, G., Birbeck, N., Zhang, J., et al.: Evaluation of prostate segmentation algorithms for mri: the promise12 challenge. *MIA* **18**(2), 359–373 (2014)
21. Liu, Q., Dou, Q., Yu, L., Heng, P.A.: Ms-net: Multi-site network for improving prostate segmentation with heterogeneous mri data. *IEEE TMI* (2020)
22. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *3DV*. pp. 565–571. IEEE (2016)
23. Otálora, S., Atzori, M., Andrearczyk, V., Khan, A., Müller, H.: Staining invariant features for improving generalization of deep convolutional neural networks in computational pathology. *Front. Bioeng. Biotechnol.* **7**, 198 (2019)
24. Paschali, M., Conjeti, S., Navarro, F., Navab, N.: Generalizability vs. robustness: investigating medical imaging networks using adversarial examples. In: *MICCAI*. pp. 493–501. Springer (2018)
25. Yang, X., Dou, H., Li, R., Wang, X., Bian, C., Li, S., Ni, D., Heng, P.A.: Generalizing deep models for ultrasound image segmentation. In: *MICCAI*. pp. 497–505. Springer (2018)
26. Yao, L., Prosky, J., Covington, B., Lyman, K.: A strong baseline for domain adaptation and generalization in medical imaging. *MIDL* (2019)
27. Yoon, C., Hamarneh, G., Garbi, R.: Generalizable feature learning in the presence of data bias and domain class imbalance with application to skin lesion classification. In: *MICCAI*. pp. 365–373. Springer (2019)
28. Yu, L., Yang, X., Chen, H., Qin, J., Heng, P.A.: Volumetric convnets with mixed residual connections for automated prostate segmentation from 3d mr images. In: *AAAI* (2017)
29. Zhang, L., Wang, X., Yang, D., Sanford, T., Harmon, S., Turkbey, B., Wood, B.J., Roth, H., Myronenko, A., Xu, D., et al.: Generalizing deep learning for medical image segmentation to unseen domains via deep stacked transformation. *IEEE TMI* (2020)
30. Zhu, Q., Du, B., Yan, P.: Boundary-weighted domain adaptive neural network for prostate mr image segmentation. *IEEE TMI* **39**(3), 753–763 (2019)