



Incremental Learning Meets Transfer Learning: Application to Multi-site Prostate MRI Segmentation

Chenyu You¹(✉), Jinlin Xiang², Kun Su², Xiaoran Zhang³, Siyuan Dong¹, John Onofrey⁴, Lawrence Staib^{1,3,4}, and James S. Duncan^{1,3,4}

¹ Electrical Engineering, Yale University, New Haven, CT, USA
chenyu.you@yale.edu

² Electrical and Computer Engineering, The University of Washington, WA, USA

³ Biomedical Engineering, Yale University, New Haven, CT, USA

⁴ Radiology and Biomedical Imaging, Yale School of Medicine, New Haven, CT, USA

Abstract. Many medical datasets have recently been created for medical image segmentation tasks, and it is natural to question whether we can use them to sequentially train a single model that (1) performs better on all these datasets, and (2) generalizes well and transfers better to the unknown target site domain. Prior works have achieved this goal by jointly training one model on multi-site datasets, which achieve competitive performance on average but such methods rely on the assumption about the availability of all training data, thus limiting its effectiveness in practical deployment. In this paper, we propose a novel multi-site segmentation framework called **incremental-transfer learning (ITL)**, which learns a model from multi-site datasets in an end-to-end sequential fashion. Specifically, “incremental” refers to training sequentially constructed datasets, and “transfer” is achieved by leveraging useful information from the linear combination of embedding features on each dataset. In addition, we introduce our ITL framework, where we train the network including a site-agnostic encoder with pretrained weights and at most two segmentation decoder heads. We also design a novel site-level incremental loss in order to generalize well on the target domain. Second, we show for the first time that leveraging our ITL training scheme is able to alleviate challenging catastrophic forgetting problems in incremental learning. We conduct experiments using five challenging benchmark datasets to validate the effectiveness of our incremental-transfer learning approach. Our approach makes minimal assumptions on computation resources and domain-specific expertise, and hence constitutes a strong starting point in multi-site medical image segmentation.

Keywords: Incremental learning · Transfer learning · Medical image segmentation

C. You and J. Xiang—Equal contribution.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022
S. Albarqouni et al. (Eds.): DeCaF 2022/FAIR 2022, LNCS 13573, pp. 3–16, 2022.
https://doi.org/10.1007/978-3-031-18523-6_1

1 Introduction

Many medical image datasets have been created over the year, and recent breakthrough achieved by supervised training accelerates the pace in medical image segmentation. Despite great promise, many prior works have limited clinical value, since they are separately trained on small datasets in terms of scale, diversity, and heterogeneity of annotations. As a result, such single-site methods [10, 14, 21, 22, 29, 31, 32, 35–41] are vulnerable to unknown target domains, and linearly expand parameters since they assume to train a new model in isolation when adding new datasets. This jeopardizes their trustworthiness and practical deployment in real-world clinical environments.

In this paper, we carry out the **first-of-its-kind** comprehensive exploration of how to build a multi-site model to achieve strong performance on the training domains and can also serve as a strong starting point for better generalization on new domains in the clinical scenarios. Multi-site training [1, 3, 7, 8, 11, 24, 25] has been proposed to consolidate the generalization on multi-site datasets, but it has the following limitations: (1) it still exhibits certain vulnerability to different domains (*i.e.*, different imaging protocols), which yields sub-optimal performance [1, 13, 34]; (2) due to various constraints (*i.e.*, imaging time, privacy, and copyright status), it could become challenging or even infeasible for the requirement on the availability of all training data in a certain time phase. For example, when a new site’s data will be available after training, the model requires retraining, which largely prohibits the practical deployments; and (3) consider the relatively small size of the single medical imaging dataset, simply training a dense network from scratch usually leads to sub-optimal segmentation quality because the model might over-fit to those datasets.

Our **key idea** is to combine the benefits of incremental-learning (IL) and transfer-learning by sequentially training a multi-dataset expert: we continually train a model with corresponding pretrained weights as new site data are incrementally added, which we call **Incremental-Transfer Learning (ITL)**. This setting is appealing as: (1) the common IL setting [4, 5, 15, 17, 23, 27, 28, 42] is to train the base-learner when different site datasets gradually come; thus the effectiveness of this approach heavily depends on the optimality of the base-learner. Consider each single medical image dataset is usually of relatively small size, it is undesirable to build a strong base-learner from scratch; (2) transfer-learning [26, 30, 33, 43, 44] typically leads to better performance and faster convergence in medical image analysis. Inspired by these findings above, we develop a novel training strategy for expanding its high-quality learning abilities to our multi-site incremental setting, considering both *model-level* and *site-level*. Specifically, our system is built upon a site-agnostic encoder with pretrained weights from natural image datasets such as IMAGENET, and at most two segmentation decoder heads wherein only one head is trainable, and the other is fixed associated with specific sites - a parameter-efficient design. Our intuition is that the shared site-agnostic encoder network with pretrained weights encodes regularities across different medical image datasets, while the target and source segmentation decoder heads model the sub-distribution by our proposed site-level incremental loss, resulting




Fig. 1. Overview of (a) our proposed Incremental Transfer Learning framework, and (b) the multi-site expert model. Note that in this study, we only use one multi-site expert model and one source decoder network, which will not introduce additional parameter.

in an accurate and robust model that transfers better to new domains without sacrificing performance. We conduct a comprehensive evaluation of ITL on five prostate MRI datasets. Our approach can consistently achieve competitive performance and faster convergence compared to the upper-bound baselines (*i.e.*, isolated-site and mixed-site training), and has a clear advantage on overall segmentation performance compared to the lower-bound baselines (*i.e.*, multi-site training). We also find that our simple approach can effectively address the forgetting issues. Our experiments demonstrate the benefits of modeling both multi-site regularities and site-specific attributes, and thereby serve as a strong starting point on this important practical setting.

2 Method

2.1 Problem Setup

In ITL, a model incrementally learns from a sequential site stream wherein new datasets (namely, medical image segmentation tasks with new sites) are gradually added during the training, as illustrated in Fig. 1. More formally, we denote the sequence of multi-site datasets to be trained as a multi-domain data sequence $\mathcal{D} = \{D_1, D_2, \dots, D_N\}$ of N sites, and i -th site D_i contains the training images $X = \{x_j\}_{j=1}^M$ and segmentation labels $Y = \{y_j\}_{j=1}^M$, where $x_j \in \mathbb{R}^{H \times W \times 3}$ is the augmented image input, and $y_j \in \{0, 1\}^{H \times W}$ is the ground-truth label. Here the augmented input setting is appealing: the axial context naturally provided by a 3D volume can uniquely yield more robust semantic representations to the downstream tasks. We assume access to a multi-site expert model $F_i = \{E_i, G_i\}$ for i -th (site) phase, including a pretrained model as a site-agnostic encoder network E_i with the weight θ_i , a target decoder network G_i^t with the weight θ_i^t . During training, we additionally attach a source decoder network G_i^s (*i.e.*, using G_{i-1}^s from previous phrase) with the weight θ_i^s . In the i -th incremental (site) phase, the multi-site expert model has access to two types of domain knowledge: the site-specific knowledge from the current dataset D_i and old exemplars P_i . The latter refers to a set of old exemplars from all previous training datasets $D_{1:i-1}$ in

Table 1. Information about five different sites from three benchmark datasets.

Dataset	Modality	# of cases	Field strength (T)	Resolution (in/through plane) (mm)	Coil	Source
Site0	MRI	30	3	0.6–0.625/3.6–4	Surface	NCI-ISBI13 [2]
Site1	MRI	30	1.5	0.4/3.0	Endorectal	NCI-ISBI13 [2]
Site2	MRI	19	3	0.67–0.79/1.25	-	I2CVB [12]
Site3	MRI	12	1.5	0.625/3.6	Endorectal	PROMISE12 [16]
Site4	MRI	13	1.5 and 3	0.325–0.625/3–3.6	-	PROMISE12 [12]

the memory protocol \mathcal{M} . This is highly nontrivial to preventing the challenging “catastrophic forgetting” problem [20] of the current dataset i against previous sites in clinical practice. Note that, in this study, we only use one multi-site expert model and one source decoder network, which will not introduce additional parameters. Based on the setting above, we define the ITL problem below.

Problem of ITL. *In the current site i , our goal is to continuously learn a multi-site expert model based on the knowledge from both (D_i, P_i) and the pretrained weight, making the model (1) generalizes well on the unseen data at site i , and (2) achieves competitive performance on the previous sites.*

2.2 Preliminary

Our goal is to build a strong multi-site model by learning a site-agnostic encoder with pretrained weights as well as a segmentation decoder over multi-site datasets. This naturally raises several interesting questions: *How well will ITL-based methods perform in multi-site medical image datasets? Will transfer learning make the base learner stronger on the unseen site? If yes, can they perform stably well?* To answer the above questions, a prerequisite is to define the upper bound and lower bound. Here we introduce three common paradigms for multi-site medical image segmentation: (1) isolated-site training, (2) mixed-site training, and (3) multi-site training. It is well-known that the isolated-site and mixed-site training approaches can achieve state-of-the-art performance when evaluating the same dataset, while the performance catastrophically drops when evaluating new datasets. On the other hand, the multi-site training approach often yields inconsistent performance across multiple sites. For all training paradigms, we minimize Dice loss between the predicted outputs and the ground truth label.

Upper Bound. We consider two training paradigms (*i.e.*, isolated-site and mixed-site training) as our upper bound baseline. For isolated-site training, given each site D_i , we train isolated-site models separately. The architecture of the isolated-site model consists of a pretrained encoder E_i and a segmentation decoder network, same architecture as G_i . Then, we apply different isolated-site

models to predict results based on the site-specific data at inference. However, this approach dramatically increases memory and computational overhead, making it practically challenging at scale. For mixed-site training, we train one full model on the full mixed-site data D , and then use the well-trained model for inference. However, it requires the simultaneous presence of all data in training and inference.

Lower Bound. For multi-site training, we sequentially train only one model coupled with the pretrained weights on all sites. This can get rid of large parameter counts, making it appealing in practice. However, due to the forgetting quandary, it inevitably suffers from severe performance degradation. This naturally questions: *can we improve performance on multi-site medical image segmentation with minimal additional memory footprint?* In the following, we give an affirmative answer.

2.3 Proposed Incremental Transfer Learning Multi-site Method

To address the aforementioned problems, we develop the incremental transfer learning framework to perform well on the training distribution and generalize well on the new site dataset with minimal additional memory. To our best knowledge, we are **the first work** to apply incremental transfer learning to the limited clinical data regimes. To control the parameter efficiency, we decompose the model into a share site-agnostic encoder E_i and two segmentation decoder heads (*i.e.*, source decoder G_i^s and target decoder G_i^t). In this way, we can keep the network parameters the same when adding a new site. Specifically, G_i^s is designed to transfer the knowledge of a previously learned site, and G_i^t is designed to comprehensively train on a new site and previous datasets. During training, we only update G_i^t while G_i^s is frozen. It is worth mentioning that our proposed framework is independent of the encoder architecture, and can be easily plugged in other pretrained vision models.

The full ITL algorithm is summarized in Algorithm 1. We describe our ITL algorithm as follows. We first randomly initialize G_i^t , G_i^s , and then iteratively train our full model (*i.e.*, a pretrained encoder E_i and two decoders G_i^t , G_i^s) with N -site training samples. Bounded by the computational requirements, it is challenging or even infeasible to retain all data for training. Inspired by recent work [23], to maintain the knowledge of previous sites, we “store” all the old site data exemplars in the memory protocol \mathcal{M}_i . In the i -th incremental (site) phase, we first load P_i , and then use both P_i and D_i to train F_i initialized by θ_i^s . This setting is appealing as (1) it can substantially alleviate the imbalance between the old and new site knowledge, and (2) it is efficient to train on them. Of note, we do not use the source decoder when training on the first-site dataset. We formulate ITL as *model-level* and *site-level* optimization.

Model-Level Optimization. To perform better on all these training distributions, we propose improving generic representations by distilling knowledge from previous data. In each incremental phase, we jointly optimize two groups of learnable parameters in our ITL learning by minimizing the *model-level* incremental

Algorithm 1. Incremental-Transfer Learning(ITL) Algorithm

Require: Dataset: \mathcal{D} ; Hyper-parameters: α, δ, γ

- 1: Initialize the \mathcal{M} (*Memory*) : \mathcal{M}
- 2: Initialize the Model F_0 : *Pertrained Encoder* $\rightarrow E_0, G_0$
- 3: **for** $i = 1, 2, 3, \dots, N$ **do**
- 4: **for** All training Sample in D_i and \mathcal{M}_{i-1} **do**
- 5: $\mathcal{L}_{\text{target}} = \sum_{j=0}^{N-1} \alpha_j \mathcal{L}_{\text{Dice}}^{E_i, G_i^t}(\mathcal{M}_j, Y_j)$ or 0 When $N = 1$
- 6: $\mathcal{L}_{\text{source}} = \sum_{j=0}^{N-1} \delta_j \mathcal{L}_{\text{Dice}}^{E_i, G_i^s}(\mathcal{M}_j, Y_j)$ or 0 When $N = 1$
- 7: $\mathcal{L}_{\text{model}} = \mathcal{L}_{\text{source}} + \mathcal{L}_{\text{target}}$
- 8: $\mathcal{L}_{\text{site}} = \mathcal{L}_{\text{Dice}}^{E_i, G_i^t}(D_i, Y_i)$
- 9: $\mathcal{L}_{\text{all}} = \mathcal{L}_{\text{site}} + \mathcal{L}_{\text{model}}$
- 10: $F_i = (E_i, G_i^t)$ by minimizing the \mathcal{L}_{all}
- 11: **end for**
- 12: Update Memory: $\mathcal{M} + \gamma\%D_N \rightarrow \mathcal{M}$
- 13: Save Teacher Model: G_N
- 14: **end for**

loss (*i.e.*, $\mathcal{L}_{\text{model}} = \mathcal{L}_{\text{target}} + \mathcal{L}_{\text{source}}$) on all training samples (*i.e.*, $D_i \cup D_{0:i-1}$): (1) a share site-agnostic encoder E_i and a target decoder G_i^t ; (2) a share site-agnostic encoder E_i and a source decoder G_i^s . This helps ITL avoid catastrophic forgetting of prior site-specific knowledge.

Site-Level Optimization. The above *model-level* optimization is used to maintain previously learned knowledge. In contrast, this step is design to train the multi-site model to learn site-specific knowledge on the newly added site. Specifically, we minimize the site-level incremental loss $\mathcal{L}_{\text{site}}$ between the probability distribution from F_i and the ground truth. This essentially learns the site-specific knowledge for the downstream medical image segmentation tasks. Of note, $\mathcal{L}_{\text{source}}$, $\mathcal{L}_{\text{target}}$, and $\mathcal{L}_{\text{site}}$ use the Dice loss. The overall loss combines the *model-level* loss and the *site-level* loss as follows:

$$\mathcal{L}_{\text{all}} = \mathcal{L}_{\text{model}} + \mathcal{L}_{\text{site}}. \quad (1)$$

3 Experiments

Datasets and Settings. We evaluate our proposed incremental transfer learning method on three prostate T2-weighted MRI datasets with different sub-distributions: NCI-ISBI13 [2], I2CVB [12], and PROMISE12 [16]. Due to the diverse data source distributions, they can be split into five multi-site datasets, which is similar to [19]. Table 1 provides some dataset statistics. For pre-processing, we follow the setting in [18] to normalize the intensity, and resample all 2D slices and the corresponding segmentation maps to 384×384 in the axial plane. For all five site datasets, we randomly split each original site dataset into training and testing with a ratio of 4:1. For each site training, we divide the data from the previous site into a small subset with a certain portion (*i.e.*, 1%, 3%, 5%), and combine it with the current site data for training.

Table 2. Comparison of segmentation performance (DSC[%]/95HD[mm]) across datasets. Note that a larger DSC (\uparrow) and a smaller 95HD (\downarrow) indicate better performing ITL models. We use four models pretrained on IMAGENET: ResNet-18, ResNet-34, ResNet-50, and ViT under different portions (*i.e.*, 1%, 3%, 5%) of exemplars from previous data for every incremental phase. We consider multi-site training as the lower bound, isolated-site, and mixed-site training as the upper bound.

Backbone	Scheme	HK		UCL		ISBI		ISBI1.5		I2CVB	
		DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]
RES-18	Multi	59.38	64.17	66.26	54.19	54.38	73.40	66.89	44.49	84.54	11.70
	1%	67.82	56.08	67.12	58.05	59.47	70.46	77.34	34.77	82.94	6.06
	3%	71.60	18.41	82.18	23.92	72.26	20.91	81.53	19.21	84.08	13.75
	5%	81.81	5.50	84.45	13.95	84.52	15.65	89.32	10.11	86.72	11.70
	Isolated	93.46	2.06	88.29	6.20	93.35	2.04	90.89	7.53	88.74	13.93
	Mixed	92.17	7.60	83.38	12.22	91.70	2.46	90.08	9.20	89.12	13.86
RES-34	Multi	57.75	55.13	64.87	52.50	57.47	65.38	65.61	56.83	91.46	8.83
	1%	67.40	24.18	79.55	30.43	69.61	44.69	84.68	18.71	89.38	15.24
	3%	80.90	28.41	82.57	22.18	75.89	26.26	84.68	10.57	90.35	13.15
	5%	80.46	22.92	87.79	17.32	88.14	14.64	90.29	8.57	91.30	8.52
	Isolated	93.87	1.89	89.03	4.05	92.08	2.19	92.57	7.96	91.57	7.98
	Mixed	93.85	1.71	87.81	16.85	91.49	3.35	93.82	5.30	92.58	6.64
RES-50	Multi	63.24	53.98	64.79	56.59	72.95	26.63	69.41	49.89	90.40	8.21
	1%	69.01	60.70	69.85	44.21	75.30	28.74	80.27	20.10	90.08	8.02
	3%	78.72	16.89	83.74	12.81	84.96	8.51	86.95	6.18	92.34	5.24
	5%	92.46	2.92	88.79	10.97	92.16	2.04	92.18	4.87	91.35	2.12
	Isolated	93.73	2.12	89.03	7.23	93.26	4.39	93.48	5.10	93.20	2.40
	Mixed	94.38	1.34	88.28	9.77	92.71	9.43	92.27	5.29	90.45	5.29
ViT	Multi	66.94	53.57	65.85	54.69	92.66	6.37	72.80	51.35	90.56	7.02
	1%	71.99	48.61	85.29	11.35	75.99	17.87	84.73	12.32	90.11	7.23
	3%	79.33	20.84	88.16	7.08	85.48	7.97	87.64	9.95	90.07	6.94
	5%	93.25	1.37	87.62	9.23	92.22	4.82	91.62	2.82	91.87	6.59
	Isolated	94.44	1.88	88.80	8.21	93.23	4.76	92.47	6.27	93.23	6.43
	Mixed	93.30	1.38	87.20	9.21	92.86	9.29	86.92	12.28	92.01	6.99

Training and Evaluation. In this study, we implement all models using Pytorch. We set H, W as 384, α, δ as 0.5, and the batch size as 5. To mitigate the overfitting, we augment the data by random horizontal flipping, random rotation, and random shift. We adopt ResNet family [9] (*i.e.*, ResNet18, ResNet34, ResNet50) and ViT [6] (*i.e.*, R50+ViT-B/16 hybrid model) as our pretrained encoder. We evaluate the model performance by Dice coefficient (DSC) and 95% Hausdorff Distance (95HD). For a fair comparison, we adopt the same decoder architecture design in [18] are shown in Appendix Table 4, and do not use any post-processing techniques. All of our experiments are conducted on two NVIDIA Titan X GPUs. All the models are trained using Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$. For 100 epochs training, a multi-step learning rate schedule is initialized as 0.001 and then decayed with a power of 0.95 at epochs 60 and 80.

Main Results. We conduct extensive experiments on five benchmark datasets. We adopt four models: ResNet-18, ResNet-34, ResNet-50, and ViT. We select three portions (*i.e.*, 1%, 3%, 5%) of exemplars from previous data for every incremental phase. Our results are presented in Table 2 and Appendix Fig. 2. First and foremost, we can see ITL-based methods generalize across all datasets under two exemplar portions (*i.e.*, 3% and 5%), yielding the competitive

segmentation quality comparable to the upper bound baselines (*i.e.*, isolated-site and mixed-site training), which are much higher than the lower bound counterparts. The 1% exemplar portion seems slightly more challenging for ITL, but its superiority over the lower bound counterparts is still solid. A possible explanation for this finding is that using two exemplar portions (*i.e.*, 3% and 5%) maintains enough information of ITL, which mitigates the catastrophic forgetting, while ITL trained in the setting of 1% exemplar portion is not powerful enough to inherit prior knowledge and generalize well on newly added sites. Second, we consistently observe that ITL using larger models (*i.e.*, ResNet-50 and ViT) generalize substantially better than those using small models (*i.e.*, ResNet-18 and ResNet-34), which demonstrate competitive performance across all datasets. These results suggest that our ITL using the large model as our pretrained encoder leads to substantial gains in the setting of very limited data.

4 Analysis and Discussion

We address several research questions pertaining to our ITL approach. We use a ResNet-18 model as our encoder in our experiments. For comparisons, all models are trained for the same number of epochs, and all results are the average of three independent runs of experiments. To study the effectiveness of our proposed ITL framework, we performed experiments with 5% exemplars ratio.

Table 3. Comparison of segmentation performance in different phases.

HK		UCL		ISBI		ISBI1.5		I2CVB	
DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]
94.06	1.96	-	-	-	-	-	-	-	-
93.68	1.98	88.74	8.72	-	-	-	-	-	-
93.20	1.83	87.38	9.30	92.87	1.82	-	-	-	-
90.37	8.34	86.73	12.75	89.84	13.32	91.57	11.05	-	-
88.88	8.91	85.14	10.97	85.98	14.23	89.74	13.11	88.46	12.15

Does Transfer Learning Lead to Better ITL? We draw two perspectives that may intuitively explain the effectiveness of transfer learning in our proposed ITL framework. As a first test of *whether transfer learning makes the base-learner stronger*, we plot the training loss/validation loss (*i.e.*, \mathcal{L}_{all}) to iteration to demonstrate the convergence improvements in Appendix Fig. 3. We can see that training from pretrained weights can converge faster than training from scratch. Another (perhaps not so surprising) observation we can get from Appendix Fig. 3 is that using pretrained weights usually yields slightly smaller loss compared to training from scratch. We then ask *whether transfer learning produces increased performance on multi-site datasets*. Since each single medical image dataset is usually of relatively small size, training the model from scratch tends to overfit a

particular dataset. To evaluate the impact of transferring learning, we compare w/pretraining to w/o pretraining. As shown in Appendix Table 7, training from scratch does not bring benefits to the ITL framework. Instead of training from scratch, we find that simply incorporating transfer learning significantly boosts the performance of ITL while achieving faster convergence speed, suggesting that transfer learning provides additional regularization against overfitting.

Does ITL Generalizes Well on Multi-site Datasets? We investigate whether the ITL framework generalizes well on multi-site datasets. We report the segmentation results of different phases in Table 3, from which we observe that ITL achieves good performance in different phases. This reveals that our approach is greatly helpful in reducing forgetting issues. We evaluate the proposed ITL methods with two random ordering (*i.e.*, (1) {HK→UCL→ISBI→ISBI1.5→I2CVB}, and (2) {ISBI→ISBI1.5→I2CVB→HK→UCL}). The results are shown in Appendix Table 5. We perform experiments using both ordering strategies and observe comparable performance.

Efficiency of ITL. We report the network size and memory costs in Appendix Table 6. We observe that ITL achieves competitive performance and utilizes less network parameters compared to isolated-site training (upper bound), which requires the new model when adding new site data. We also examine the required memory footprint at each incremental phase. We observe that ITL is significantly more memory-efficient than mixed-site training (upper bound), although the latter remains the same network size when adding a new training phase. These results further demonstrate the efficiency of our proposed ITL framework.

5 Conclusion

In this paper, we present a novel incremental transfer learning framework for incrementally tackling multi-site medical image segmentation tasks. We pose model-level and site-level incremental training strategies for better segmentation, generalization, and transfer performance, especially in limited clinical resource settings. Extensive experimental results on four different baseline architectures demonstrate the effectiveness of our approach, offering a strong starting point to encourage future work in these important practical clinical scenarios.

Appendix




Fig. 2. Visualization of segmentation results on five benchmarks using ResNet-18 as the encoder. Different site results are shown in different colors.

Table 4. Segmentation decoder head architecture

Deocder	
Layer	Feature size
Upsample 1	48×48
Residual block 1	48×48
Upsample 2	96×96
Residual block 2	96×96
Upsample 3	192×192
Residual block 3	192×192
Upsample 4	384×384
Residual block 4	384×384
Output prediction	384×384

Table 5. Comparison of different ordering strategies using ResNet-18. We report mean and standard deviation across three random trials. Note that a larger DSC (\uparrow) and a smaller 95HD (\downarrow) indicate better performing ITL models.

Training sequence	DSC[%]	95HD[mm]
HK \rightarrow UCL \rightarrow ISBI \rightarrow ISBI1.5 \rightarrow I2CVB	85.36 ± 0.33	11.38 ± 0.36
ISBI \rightarrow ISBI1.5 \rightarrow I2CVB \rightarrow HK \rightarrow UCL	86.27 ± 0.27	12.01 ± 0.68

Table 6. Comparison of different training strategies using ResNet-18. We report mean and standard deviation across three random trials.

Scheme	DSC[%]	95HD[mm]	Model size(Mb)	Add new sites?	Memory cost
Isolated	90.95 ± 0.27	6.35 ± 0.68	77.9 × Site Num.	Linearly increase	New data
Mixed	89.29 ± 0.38	9.06 ± 0.84	77.9	Constant	Old data + New data
ITL	85.36 ± 0.33	11.38 ± 0.36	77.9	Constant	5% old data + new data

Table 7. Ablation of each component in the proposed ITL when using ResNet-18 under 5% exemplar portion. We report mean and standard deviation across three random trials. Note that a larger DSC (↑) and a smaller 95HD (↓) indicate better performing ITL models. The best results are in **bold**.

Backbone	Component	HK		UCL		ISBI		ISBI1.5		I2CVB		Avg. DSC	Avg. 95HD
		DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]	DSC[%]	95HD[mm]		
RES-18	Pretraining only	59.38	64.17	66.26	54.19	54.38	73.40	66.89	44.49	81.54	28.70	65.69 ± 1.51	52.99 ± 0.72
	$\mathcal{L}_{\text{model}}$ only	74.02	18.86	73.79	31.90	51.23	51.66	80.89	21.96	80.88	58.44	72.16 ± 0.38	36.56 ± 0.85
	Pretraining + $\mathcal{L}_{\text{model}}$	81.81	5.50	84.45	13.95	84.52	15.65	89.32	10.11	86.72	11.70	85.36 ± 0.33	11.38 ± 0.36




Fig. 3. Comparison of training from scratch against using pretraining. We use ResNet-18 on IMAGENET as the encoder. Under 5% exemplar portion, we plot (a) training loss (Scratch), (b) training loss (Pretraining), (c) validation loss (Scratch), (d) validation loss (Pretraining)

References

- Aslani, S., Murino, V., Dayan, M., Tam, R., Sona, D., Hamarneh, G.: Scanner invariant multiple sclerosis lesion segmentation from MRI. In: ISBI. IEEE (2020)
- Bloch, N., et al.: NCI-ISBI 2013 challenge: automated segmentation of prostate structures. *Cancer Imaging Archive* **370** (2015)
- Chang, W.G., You, T., Seo, S., Kwak, S., Han, B.: Domain-specific batch normalization for unsupervised domain adaptation. In: CVPR (2019)
- Chaudhry, A., Dokania, P.K., Ajanthan, T., Torr, P.H.: Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: ECCV (2018)
- Davidson, G., Mozer, M.C.: Sequential mastery of multiple visual tasks: networks naturally learn to learn and forget to forget. In: CVPR (2020)
- Dosovitskiy, A., et al.: An image is worth 16x16 words: transformers for image recognition at scale. In: ICLR (2020)
- Dou, Q., Liu, Q., Heng, P.A., Glocker, B.: Unpaired multi-modal segmentation via knowledge distillation. *IEEE Trans. Med. Imaging* **39**(7), 2415–2425 (2020)
- Gibson, E., et al.: Inter-site variability in prostate segmentation accuracy using deep learning. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11073, pp. 506–514. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_58

9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR, pp. 770–778 (2016)
10. Jia, H., Song, Y., Huang, H., Cai, W., Xia, Y.: HD-Net: hybrid discriminative network for prostate segmentation in MR images. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 110–118. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_13
11. Karani, N., Chaitanya, K., Baumgartner, C., Konukoglu, E.: A lifelong learning approach to brain MR segmentation across scanners and protocols. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11070, pp. 476–484. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00928-1_54
12. Lemaître, 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)
13. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In: CVPR (2019)
14. Li, X., Yu, L., Chen, H., Fu, C.W., Heng, P.A.: Semi-supervised skin lesion segmentation via transformation consistent self-ensembling model. arXiv preprint [arXiv:1808.03887](https://arxiv.org/abs/1808.03887) (2018)
15. Li, Z., Hoiem, D.: Learning without forgetting. IEEE Trans. Pattern Anal. Mach. Intell. **40**(12), 2935–2947 (2017)
16. Litjens, G., et al.: Evaluation of prostate segmentation algorithms for MRI: the PROMISE12 challenge. MIA **18**(2), 359–373 (2014)
17. Liu, P., Xiao, L., Zhou, S.K.: Incremental learning for multi-organ segmentation with partially labeled datasets. In: MICCAI (2021)
18. Liu, Q., Dou, Q., Heng, P.-A.: Shape-aware meta-learning for generalizing prostate MRI segmentation to unseen domains. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12262, pp. 475–485. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59713-9_46
19. Liu, Q., Dou, Q., Yu, L., Heng, P.A.: MS-net: multi-site network for improving prostate segmentation with heterogeneous MRI data. IEEE Trans. Med. Imaging **39**(9), 2713–2724 (2020)
20. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: the sequential learning problem. In: Psychology of Learning and Motivation, vol. 24, pp. 109–165. Elsevier (1989)
21. 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)
22. Nie, D., Gao, Y., Wang, L., Shen, D.: ASDNet: attention based semi-supervised deep networks for medical image segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11073, pp. 370–378. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_43
23. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: iCaRL: incremental classifier and representation learning. In: CVPR (2017)
24. Rundo, L., et al.: Use-net: incorporating squeeze-and-excitation blocks into u-net for prostate zonal segmentation of multi-institutional MRI datasets. Neurocomputing **365**, 31–43 (2019)
25. Rundo, L., et al.: CNN-based prostate zonal segmentation on T2-weighted MR images: a cross-dataset study. In: Esposito, A., Faundez-Zanuy, M., Morabito, F.C., Pasero, E. (eds.) Neural Approaches to Dynamics of Signal Exchanges. SIST, vol.

- 151, pp. 269–280. Springer, Singapore (2020). https://doi.org/10.1007/978-981-13-8950-4_25
26. Shi, G., Xiao, L., Chen, Y., Zhou, S.K.: Marginal loss and exclusion loss for partially supervised multi-organ segmentation. *Med. Image Anal.* **70**, 101979 (2021)
 27. Wu, Y., et al.: Large scale incremental learning. In: CVPR (2019)
 28. Xiang, J., Shlizerman, E.: TKIL: tangent kernel approach for class balanced incremental learning. arXiv preprint [arXiv:2206.08492](https://arxiv.org/abs/2206.08492) (2022)
 29. Yang, L., et al.: NuSeT: a deep learning tool for reliably separating and analyzing crowded cells. *PLoS Comput. Biol.* **16**(9), e1008193 (2020)
 30. Yao, Q., Xiao, L., Liu, P., Zhou, S.K.: Label-free segmentation of COVID-19 lesions in lung CT. *IEEE Trans. Med. Imaging* **40**(10), 2808–2819 (2021)
 31. You, C., Dai, W., Staib, L., Duncan, J.S.: Bootstrapping semi-supervised medical image segmentation with anatomical-aware contrastive distillation. arXiv preprint [arXiv:2206.02307](https://arxiv.org/abs/2206.02307) (2022)
 32. You, C., Yang, J., Chapiro, J., Duncan, J.S.: Unsupervised Wasserstein distance guided domain adaptation for 3D multi-domain liver segmentation. In: Cardoso, J., et al. (eds.) IMIMIC/MIL3ID/LABELS -2020. LNCS, vol. 12446, pp. 155–163. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-61166-8_17
 33. You, C., et al.: Class-aware generative adversarial transformers for medical image segmentation. arXiv preprint [arXiv:2201.10737](https://arxiv.org/abs/2201.10737) (2022)
 34. You, C., Zhao, R., Staib, L., Duncan, J.S.: Momentum contrastive voxel-wise representation learning for semi-supervised volumetric medical image segmentation. arXiv preprint [arXiv:2105.07059](https://arxiv.org/abs/2105.07059) (2021)
 35. You, C., Zhou, Y., Zhao, R., Staib, L., Duncan, J.S.: SimCVD: simple contrastive voxel-wise representation distillation for semi-supervised medical image segmentation. *IEEE Trans. Med. Imaging* (2022)
 36. 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)
 37. Zhang, X., et al.: Automatic spinal cord segmentation from axial-view MRI slices using CNN with grayscale regularized active contour propagation. *Comput. Biol. Med.* **132**, 104345 (2021)
 38. Zhang, X., Martin, D.G., Noga, M., Punithakumar, K.: Fully automated left atrial segmentation from MR image sequences using deep convolutional neural network and unscented Kalman filter. In: 2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), pp. 2316–2323. IEEE (2018)
 39. Zhang, X., Noga, M., Martin, D.G., Punithakumar, K.: Fully automated left atrium segmentation from anatomical cine long-axis MRI sequences using deep convolutional neural network with unscented Kalman filter. *Med. Image Anal.* **68**, 101916 (2021)
 40. Zhang, X., Noga, M., Punithakumar, K.: Fully automated deep learning based segmentation of normal, infarcted and edema regions from multiple cardiac MRI sequences. In: Zhuang, X., Li, L. (eds.) MyoPS 2020. LNCS, vol. 12554, pp. 82–91. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-65651-5_8
 41. Zhang, Y., Yang, L., Chen, J., Fredericksen, M., Hughes, D.P., Chen, D.Z.: Deep adversarial networks for biomedical image segmentation utilizing unannotated images. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10435, pp. 408–416. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_47
 42. Zheng, Y., Xiang, J., Su, K., Shlizerman, E.: BI-MAML: balanced incremental approach for meta learning. arXiv preprint [arXiv:2006.07412](https://arxiv.org/abs/2006.07412) (2020)

43. Zhou, S.K., et al.: A review of deep learning in medical imaging: imaging traits, technology trends, case studies with progress highlights, and future promises. *Proc. IEEE* **109**(5), 820–838 (2021)
44. Zhu, J., Li, Y., Hu, Y., Ma, K., Zhou, S.K., Zheng, Y.: Rubik’s Cube+: a self-supervised feature learning framework for 3D medical image analysis. *Med. Image Anal.* **64**, 101746 (2020)