



THE UNIVERSITY *of* EDINBURGH

## Edinburgh Research Explorer

### **Semi-supervised Meta-learning with Disentanglement for Domain-generalised Medical Image Segmentation**

**Citation for published version:**

Liu, X, Thermos, S, O'Neil, A & Tsaftaris, SA 2021, Semi-supervised Meta-learning with Disentanglement for Domain-generalised Medical Image Segmentation. in Medical Image Computing and Computer Assisted Intervention – MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part II. vol. 12902, Lecture Notes in Computer Science, vol. 12902, Springer, pp. 307-317, 24th International Conference on Medical Image Computing and Computer Assisted Intervention, 27/09/21. [https://doi.org/10.1007/978-3-030-87196-3\\_29](https://doi.org/10.1007/978-3-030-87196-3_29)

**Digital Object Identifier (DOI):**

[10.1007/978-3-030-87196-3\\_29](https://doi.org/10.1007/978-3-030-87196-3_29)

**Link:**

[Link to publication record in Edinburgh Research Explorer](#)

**Document Version:**

Peer reviewed version

**Published In:**

Medical Image Computing and Computer Assisted Intervention – MICCAI 2021

**General rights**

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [openaccess@ed.ac.uk](mailto:openaccess@ed.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# Semi-supervised Meta-learning with Disentanglement for Domain-generalised Medical Image Segmentation

Xiao Liu<sup>1</sup>, Spyridon Thermos<sup>1</sup>, Alison O’Neil<sup>1,3</sup> and Sotirios A. Tsaftaris<sup>1,2</sup>

<sup>1</sup> School of Engineering, University of Edinburgh, Edinburgh EH9 3FB, UK

<sup>2</sup> The Alan Turing Institute, London, UK

<sup>3</sup> Canon Medical Research Europe Ltd., Edinburgh, UK

Xiao.Liu@ed.ac.uk, SThermos@ed.ac.uk, Alison.ONeil@eu.medical.canon,  
S.Tsaftaris@ed.ac.uk

**Abstract.** Generalising deep models to new data from new centres (termed here domains) remains a challenge. This is largely attributed to shifts in data statistics (domain shifts) across source and unseen domains. Recently, gradient-based meta-learning approaches where the training data are split into meta-train and meta-test sets to simulate and handle the domain shifts during training have shown improved generalisation performance. However, the current fully supervised meta-learning approaches are not scalable for medical image segmentation, where large effort is required to create pixel-wise annotations. Meanwhile, in a low data regime, the simulated domain shifts may not approximate the true domain shifts well across source and unseen domains. To address this problem, we propose a novel semi-supervised meta-learning framework with disentanglement. We explicitly model the representations related to domain shifts. Disentangling the representations and combining them to reconstruct the input image allows unlabeled data to be used to better approximate the true domain shifts for meta-learning. Hence, the model can achieve better generalisation performance, especially when there is a limited amount of labeled data. Experiments show that the proposed method is robust on different segmentation tasks and achieves state-of-the-art generalisation performance on two public benchmarks. Code is publicly available at: <https://github.com/vios-s>

**Keywords:** Domain generalisation · Disentanglement · Medical image segmentation.

## 1 Introduction

Despite recent progress in medical image segmentation [3, 8, 17], inference performance on unseen datasets, acquired from distinct scanners or clinical centres, is known to decrease [5, 36]. Such reduction is mainly caused by shifts in data statistics between different clinical centres i.e. *domain shifts* [42], due to variation in patient populations, scanners, and scan acquisition settings [39]. The

variation in population impacts the underlying anatomy and pathology due to factors such as gender, age, ethnicity, which may differ for patients in different locations [27,40]. The variation in scanners and scan acquisition settings impacts the characteristics of the acquired image, such as brightness and contrast [42].

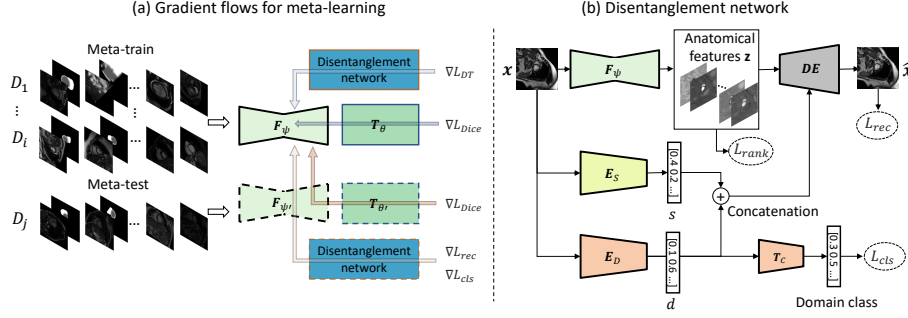
The naive approach to handling domain shift is to acquire and label as many and diverse data as possible, the cost implications and difficulties of which are known to this community. Alternatively one can train a model on source domains to generalise for a target domain with some information on the target domain available i.e. domain adaptation [4] such as cross-site MRI harmonization to enforce the source and target domains to share similar image-specific characteristics [9]. A more strict alternative is to *not use any* information for the target domain, known as *domain generalisation* [21]. Herein, we focus on this more challenging and more widely applicable approach.

In domain generalisation, the overarching goal is to identify suitable representations that encode information about the task at hand whilst being insensitive to domain specific information. There are several active research directions aiming to address this goal, including: direct augmentation of the source domain data [42], feature space regularisation [6,13,26,34,43], alignment of the source domain features or output distributions [24], and learning domain-invariant features with gradient-based meta-learning [11,23,28]. Of the above, gradient-based meta-learning methods have the advantage of not overfitting to dominant source domains which account for the more populous data in the training dataset [11]. Gradient-based meta-learning [21,25] exploits an episodic training paradigm [22] by splitting the source domains into meta-train and meta-test sets at each iteration. The model is trained to handle domain shift by simulating it during training. By using constraints to implicitly eliminate the information related to the simulated domain shifts, the model can learn to extract domain-invariant features. Previous work introduced different constraints in a fully supervised setting e.g. global class alignment and local sample clustering [11], shape-aware constraints [28] or simply the task objective [19,21], where [19] extends [21] to medical image segmentation.<sup>1</sup> These approaches do not scale in medical image segmentation as pixel-wise annotation is time-consuming, laborious and requires expert knowledge. Meanwhile, in a low data regime where centres only provide few labeled data samples, these methods may only learn to extract domain-invariant features from an under-represented data distribution [19,42]. In other words, the simulated domain shifts may not approximate well the true domain shifts between source and unseen domains.

To address this problem, we propose to explicitly disentangle the representations related to domain shifts for meta-learning as we illustrate in Fig. 1. Learning these complete and sufficient representations [1] via reconstruction brings the benefit of unsupervised learning, thus we can better simulate the domain shifts by also using unlabeled data from any of the source domains. We consider two sources of shifts: one due to scanner and scan acquisition setting variation and

---

<sup>1</sup>With the exception of [38] which clusters unlabeled data to generate pseudo labels, but unfortunately is not applicable in segmentation.



**Fig. 1.** (a) A feature network  $F_\psi$  extracts features  $z$  for a task network  $T_\theta$  to predict segmentation masks. The model is trained in a semi-supervised setting, where  $\mathcal{L}_{DT}$ ,  $\mathcal{L}_{rec}$  and  $\mathcal{L}_{cls}$  do not require pixel-wise annotation. At each iteration, the training dataset is split into meta-train and meta-test sets. We propose the constraints ( $\mathcal{L}_{DT}$  and  $\mathcal{L}_{rec}$  and  $\mathcal{L}_{cls}$ ) to encourage disentanglement between the anatomical features  $z$  and common  $s$  or domain specific  $d$  representations for meta-train and meta-test sets. Gradients are computed to update  $F_\psi$  and  $T_\theta$ . (b) A disentanglement network decomposes image  $x$  to common  $s$  and specific to the domain  $d$  representations to be disentangled with  $z$ . See Section 2 for loss definitions.

one due to population variation. Because our task is segmentation, we want to be sensitive to changes in anatomy but insensitive to changes in imaging characteristics be it some common across domains or domain specific. We use spatial features as a representation of anatomy ( $z$ ) and two vectors ( $s, d$ ) to encode common or domain specific imaging characteristics. We apply specific design and learning biases to disentangle the above. For example, a spatial  $z$  is equivariant to segmentation and this has been shown to improve performance [7, 15]. We encourage  $z$  to be disentangled to  $s$  and  $d$  by exploiting also low-rank regularization [24]. Gradient-based meta-learning also encourages  $z$ ,  $s$ , and  $d$  to generalise well to unseen domains whilst at the same time improve (implicitly) their disentanglement. Our **contributions** are summarized as follows:

1. We propose the first, to the best of our knowledge, semi-supervised domain-generalisation framework combining meta-learning and disentanglement.
2. Use of low-rank regularization as a learning bias to encourage better disentanglement and hence improved generalisation performance.
3. Extensive experiments on cardiac and gray matter datasets show improved performance over several baselines especially for the limited annotated data case.

## 2 Proposed method

Consider a multi-domain training dataset  $\mathcal{D} = \{\mathbf{x}_i^k, \mathbf{y}_i^k\}_{i=1}^{N_k}, k \in \{1, 2, \dots, K\}$  that is defined on a joint space  $\mathcal{X} \times \mathcal{Y}$ , where  $\mathbf{x}_i^k$  is the  $i^{th}$  training datum from the  $k^{th}$  source domain with corresponding ground-truth segmentation mask  $\mathbf{y}_i^k$ ,

and  $N_k$  denotes the number of training samples in the  $k^{th}$  source domain. We aim to learn a model containing a feature network  $\mathbf{F}_\psi : \mathcal{X} \rightarrow \mathcal{Z}$  to extract the anatomical features  $\mathbf{z}$  and a task network  $\mathbf{T}_\theta : \mathcal{Z} \rightarrow \mathcal{Y}$  to predict the segmentation masks, where  $\psi$  and  $\theta$  denote the network parameters.

## 2.1 Gradient-based meta-learning for domain generalisation

In gradient-based meta-learning for domain generalisation, the domain shift is simulated by training the model on a sequence of episodes [22, 23]. Specifically, the meta-train set  $\mathcal{D}_{tr}$  and the meta-test set  $\mathcal{D}_{te}$  are constructed by randomly splitting the source domains  $\mathcal{D}$  for each iteration of training. Each iteration comprises a meta-train step followed by a meta-test step. For the meta-train step, the parameters  $\psi$  and  $\theta$  of  $\mathbf{F}_\psi$  and  $\mathbf{T}_\theta$  are updated by optimising the meta-train loss  $\mathcal{L}_{meta-train}$  with data from  $\mathcal{D}_{tr}$ , as defined by:

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

where  $\alpha$  is the learning rate for the meta-train update step. Typically,  $\mathcal{L}_{meta-train}$  is the task objective, e.g. the Dice loss [10] for a segmentation task. This step rewards accurate predictions on the meta-train source domains. For the meta-test step, the meta-test source domains  $\mathcal{D}_{te}$  are processed by the updated parameters  $(\psi', \theta')$  and the model is expected to contain certain properties quantified by the  $\mathcal{L}_{meta-test}$  loss.  $\mathcal{L}_{meta-test}$  is computed using the updated parameters  $(\psi', \theta')$ , whilst the gradients are computed towards the original parameters  $(\psi, \theta)$ . The final objective is defined as:

$$\underset{\psi, \theta}{\operatorname{argmin}} \mathcal{L}_{meta-train}(\mathcal{D}_{tr}; \psi, \theta) + \mathcal{L}_{meta-test}(\mathcal{D}_{te}; \psi', \theta'). \quad (2)$$

The intuition behind this scheme is that the model should not only perform well on the source domains, but its future updates should also generalise well to unseen domains. Below, we will describe our meta-train and meta-test objectives but first we present how we disentangle representations related to domain shifts.

## 2.2 Learning disentangled representations

To model appearance in a single-domain setting, typically a single vector-based variational representation is used [7]. Here, due to our multi-domain setting, inspired by [16, 41], we separately encode domain specific imaging characteristics as an additional vector-based variational representation. Hence, we aim to learn two independent vector representations, where one ( $s$ ) captures common imaging characteristics across domains and the other one ( $d$ ) captures specific imaging characteristics for each domain. In addition, we encode spatial anatomy information in a separate representation  $\mathbf{z}$ , which we encourage to be disentangled from  $s$  and  $d$ .

In particular, the input image  $\mathbf{x}$  is first encoded in a common (appearance) representation  $s = \mathbf{E}_S(\mathbf{x})$ , and a domain representation  $d = \mathbf{E}_D(\mathbf{x})$  that is

followed by a shallow domain classifier  $\mathbf{T}_C(d)$  which predicts the source domain ( $\hat{c}$ ) label of  $\mathbf{x}$ . Then, a decoder **DE** combines the extracted features  $\mathbf{z} = \mathbf{F}_\psi(\mathbf{x})$  and the representations  $s$  and  $d$  to reconstruct the input image, i.e.  $\hat{x} = \mathbf{DE}(\mathbf{z}, s, d)$ . Note that **DE** combines  $\mathbf{z}$  and  $s, d$  using adaptive instance normalization (AdaIN) layers [14]. As shown in [15], AdaIN improves disentanglement and encourages  $\mathbf{z}$  to encode spatially equivariant information, i.e. anatomical information useful for segmentation, and  $s, d$  to only encode common or domain specific appearance.

To achieve such “triple” disentanglement we consider several losses: **1)** KL divergences  $\mathcal{L}_{KL}(s, N(0, 1)), \mathcal{L}_{KL}(d, N(0, 1))$  to induce a Gaussian  $N(0, 1)$  prior in  $s$  and  $d$ , encouraging the representations to be robust on unseen domains [12]; **2)** Hilbert-Schmidt Independence Criterion (HSIC) loss  $\mathcal{L}_{HSIC}(s, d)$ , to force  $s$  and  $d$  to be independent from each other [33]; **3)** a classification loss  $\mathcal{L}_{cls}(c, \hat{c})$  such that the domain representation  $d$  is highly correlated with the domain specific information [16]; and **4)** a reconstruction loss  $\mathcal{L}_{rec}(\mathbf{x}, \hat{\mathbf{x}})$ , defined as the  $\ell_1$  distance between  $\mathbf{x}$  and  $\hat{\mathbf{x}}$ , to learn representations without supervision [7, 15].

We further encourage the extracted features  $\mathbf{z}$  to be equivariant across the meta-train source domains and improve disentanglement between  $\mathbf{z}$  and  $s, d$  by applying rank regularization [24]. Specifically, consider a batch  $\{\mathbf{x}_{i_1}^1, \mathbf{x}_{i_2}^2, \dots, \mathbf{x}_{i_{K_{tr}}}^{K_{tr}}\}$  from  $K_{tr}$  meta-train source domains, and  $K_{tr}$  features  $\{\mathbf{z}_{i_1}^1, \mathbf{z}_{i_2}^2, \dots, \mathbf{z}_{i_{K_{tr}}}^{K_{tr}}\}$  extracted using the feature network  $\mathbf{F}_\psi$ . By flattening and concatenating these features, we end up with a matrix  $\mathbf{Z}$  with dimensions  $[C, K_{tr} \times H \times W]$ , where  $C, H, W$  denote the number of channels, height, and width of  $\mathbf{z}$ . Then, by forcing the rank of  $\mathbf{Z}$  to be  $m$  (i.e. the number of the segmentation classes), we encourage  $\mathbf{z}$  to encode only globally-shared information across  $K_{tr}$  source domains in order to predict the segmentation mask [24]. We achieve that by minimizing the  $(m+1)^{th}$  singular value  $\sigma_{m+1}$  of  $\mathbf{Z}$ , i.e.  $\mathcal{L}_{rank} = \sigma_{m+1}$ . Overall,  $\mathcal{L}_{DT}$  is defined as:

$$\begin{aligned} \mathcal{L}_{DT} = & \lambda_{rank} \mathcal{L}_{rank}(\mathbf{z}) + \lambda_{KL}(\mathcal{L}_{KL}(s, N(0, 1)) + \mathcal{L}_{KL}(d, N(0, 1))) \\ & + \lambda_{rec} \mathcal{L}_{rec}(\mathbf{x}, \hat{\mathbf{x}}) + \lambda_{HSIC} \mathcal{L}_{HSIC}(s, d) + \lambda_{cls} \mathcal{L}_{cls}(c, \hat{c}), \end{aligned} \quad (3)$$

where  $c$  is the domain label. We adopt hyperparameter values from [7, 24] as  $\lambda_{rank} = 0.1$ ,  $\lambda_{KL} = 0.1$ ,  $\lambda_{rec} = 1$  and  $\lambda_{cls} = 1$ . Note that all the losses do not need ground-truth masks. The domain class label is easy to obtain, as we know the centre where the data belong.

### 2.3 Meta-train and meta-test objectives

Our *meta-train* objective contains two components:

$$\mathcal{L}_{meta-train} = \lambda_{Dice} \mathcal{L}_{Dice}(\mathbf{y}, \hat{\mathbf{y}}) + \mathcal{L}_{DT}, \quad (4)$$

where  $\lambda_{Dice} = 5$  when labeled data are available.

For the *meta-test* step, the model is expected to: **1)** accurately predict segmentation masks (by applying the task objective), and **2)** disentangle  $\mathbf{z}$  and  $s, d$

to the same level as meta-train sets. A naive strategy for the latter is to use  $\mathcal{L}_{DT}$  for meta-test sets. However, as analyzed in [2, 28], the meta-test step is unstable to train: the gradients from the meta-test loss are second-order statistics of  $\psi$  and  $\theta$ . In our experiments, we observed that having the unsupervised losses  $\mathcal{L}_{KL}$  and  $\mathcal{L}_{HSIC}$  make training even more unstable and even can lead to model collapse. In addition, we use one domain for meta-test in experiments, while  $\mathcal{L}_{rank}$  requires multiple domains. According to [30, 32], considering fixed learning and design biases, the level of disentanglement can be proxied by the reconstruction quality (with ground-truth image  $x$ ) and the domain classification accuracy (with ground-truth label  $c$ ). Hence, we adopt as meta-test loss:

$$\mathcal{L}_{meta-test} = \lambda_{Dice} \mathcal{L}_{Dice}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda_{rec} \mathcal{L}_{rec}(\mathbf{x}, \hat{\mathbf{x}}) + \lambda_{cls} \mathcal{L}_{cls}(c, \hat{c}). \quad (5)$$

Note that for unlabeled data,  $\mathcal{L}_{rec}$  and  $\mathcal{L}_{cls}$  do not need ground-truth masks.

### 3 Experiments

#### 3.1 Tasks and datasets

**Multi-centre, multi-vendor & multi-disease cardiac image segmentation (M&Ms) dataset [5]:** The M&Ms challenge dataset contains 320 subjects. Subjects were scanned at 6 clinical centres in 3 different countries using 4 different magnetic resonance scanner vendors (Siemens, Philips, GE, and Canon) i.e. domains A, B, C and D. For each subject, only the end-systole and end-diastole phases are annotated. Voxel resolutions range from  $0.85 \times 0.85 \times 10$  mm to  $1.45 \times 1.45 \times 9.9$  mm. Domain A contains 95 subjects. Domain B contains 125 subjects. Both domains C and D contain 50 subjects.

**Spinal cord gray matter segmentation (SCGM) dataset [36]:** The data from SCGM [36] are collected from 4 different medical centres with different MRI systems (Philips Achieva, Siemens Trio, Siemens Skyra) i.e. domains 1, 2, 3 and 4. The voxel resolutions range from  $0.25 \times 0.25 \times 2.5$  mm to  $0.5 \times 0.5 \times 5$  mm. Each domain contains 10 labeled subjects and 10 unlabelled subjects.

#### 3.2 Baseline models

**nnUNet [18]:** is a self-adapting framework based on 2D and 3D U-Nets [37] which does not specifically target domain generalisation. Given a labelled training dataset, nnUNet automatically adapts its model design and hyperparameters to obtain optimal performance. In the M&Ms challenge, methods based on nnUNet achieved the top performance [5].

**SDNet+Aug. [29]:** disentangles the input image to a spatial anatomy and a non-spatial modality factors. Here we use intensity- and resolution- augmented data in a semi-supervised setting. Compared to our method, “SDNet+Aug.” only poses disentanglement to the latent features without meta-learning.

**LDDG [24]:** is the latest state-of-the-art model for domain-generalised medical image analysis. It also uses a rank loss and when applied in a fully supervised

**Table 1.** Dice (%) results and the standard deviations on M&Ms dataset. For “SDNet+Aug.” and our method, the training data contain all the unlabeled data and 2% or 5% of labeled data from source domains. The other models are trained by 2% or 5% labeled data only. Bold numbers denote the best performance.

	Source	Target	nnUNet	SDNet+Aug.	LDDG	SAML	Ours
2%	B,C,D	A	52.87 <sub>19</sub>	54.48 <sub>18</sub>	59.47 <sub>12</sub>	56.31 <sub>13</sub>	<b>66.01<sub>12</sub></b>
	A,C,D	B	64.63 <sub>17</sub>	67.81 <sub>14</sub>	56.16 <sub>14</sub>	56.32 <sub>15</sub>	<b>72.72<sub>10</sub></b>
	A,B,D	C	72.97 <sub>14</sub>	76.46 <sub>12</sub>	68.21 <sub>11</sub>	75.70 <sub>8.7</sub>	<b>77.54<sub>10</sub></b>
	A,B,C	D	73.27 <sub>11</sub>	74.35 <sub>11</sub>	68.56 <sub>10</sub>	69.94 <sub>9.8</sub>	<b>75.14<sub>8.4</sub></b>
	<b>Average</b>		65.94 <sub>8.3</sub>	68.28 <sub>8.6</sub>	63.16 <sub>5.4</sub>	64.57 <sub>8.5</sub>	<b>72.85<sub>4.3</sub></b>
5%	B,C,D	A	65.30 <sub>17</sub>	71.21 <sub>13</sub>	66.22 <sub>9.1</sub>	67.11 <sub>10</sub>	<b>72.40<sub>12</sub></b>
	A,C,D	B	79.73 <sub>10</sub>	77.31 <sub>10</sub>	69.49 <sub>8.3</sub>	76.35 <sub>7.9</sub>	<b>80.30<sub>9.1</sub></b>
	A,B,D	C	78.06 <sub>11</sub>	81.40 <sub>8.0</sub>	73.40 <sub>9.8</sub>	77.43 <sub>8.3</sub>	<b>82.51<sub>6.6</sub></b>
	A,B,C	D	81.25 <sub>8.3</sub>	79.95 <sub>7.8</sub>	75.66 <sub>8.5</sub>	78.64 <sub>5.8</sub>	<b>83.77<sub>5.1</sub></b>
	<b>Average</b>		76.09 <sub>6.3</sub>	77.47 <sub>3.9</sub>	71.29 <sub>3.6</sub>	74.88 <sub>4.6</sub>	<b>79.75<sub>4.4</sub></b>

setting, LDDG achieved the best generalisation performance on SCGM.

**SAML [28]:** is another gradient-based meta-learning approach. SAML proposed to constraint the compactness and smoothness properties of segmentation masks across meta-train and meta-test sets in a fully supervised setting.

### 3.3 Implementation details

Models are trained using the Adam optimizer [20] with a learning rate of  $2e^{-5}$  for 50K iterations using batch size 4. Images are cropped to  $224 \times 244$  for M&Ms and  $144 \times 144$  for SCGM.  $\mathbf{F}_\psi$  is a 2D UNet [37] to extract  $\mathbf{z}$  features with 8 channels of same height and width as input image. We follow the designs of SDNet [7] for  $\mathbf{E}_S$ ,  $\mathbf{T}_\theta$  and  $\mathbf{DE}$ .  $\mathbf{E}_D$  has the same architecture as  $\mathbf{E}_S$ . Both  $s$  and  $d$  have 8 dimensions.  $\mathbf{T}_C$  is a single fully-connected layer. All models are implemented in PyTorch [35] and are trained using an NVidia 2080 Ti GPU. In the semi-supervised setting, we use specific percentages of the subjects as labeled data and the rest as unlabeled data. We use Dice (%) as the evaluation metric.

### 3.4 Results and discussion

Tables 1 and 2 show that we consistently achieve the best generalisation performance on cardiac and gray matter segmentation. Particularly in the low data regime we improve by  $\approx 5\%$  on M&Ms and  $\approx 3\%$  on SCGM compared to the best performing baseline. For 100% annotations in M&Ms (see Appendix), our model still outperforms the baselines. We also show the qualitative results in Appendix, where the improved performance is visually observed.

**M&Ms:** Compared to “SDNet+Aug.” which can also use (due to disentanglement) unlabeled data, our model consistently outperforms across all cases. The results agree with the conclusion in [31]: without specific designs tuned to the tasks, disentanglement can not provide guaranteed generalisation ability. For



**Table 2.** Dice (%) results and the standard deviations on SCGM dataset. For “SD-Net+Aug.” and our method, the training data contain all the unlabeled data and 20% or 100% of labeled data from source domains. The other models are trained by 20% or 100% of labeled data only. Bold numbers denote the best performance.

Source	Target	nnUNet	SDNet+Aug.	LDDG	SAML	Ours	
20%	2,3,4	1	59.07 <sub>21</sub>	83.07 <sub>16</sub>	77.71 <sub>9.1</sub>	78.71 <sub>25</sub>	<b>87.45</b> <sub>6.3</sub>
	1,3,4	2	69.94 <sub>12</sub>	80.01 <sub>5.2</sub>	44.08 <sub>12</sub>	75.58 <sub>12</sub>	<b>81.05</b> <sub>5.2</sub>
	1,2,4	3	60.25 <sub>7.2</sub>	58.57 <sub>10</sub>	48.04 <sub>5.5</sub>	54.36 <sub>7.6</sub>	<b>61.85</b> <sub>7.3</sub>
	1,2,3	4	70.13 <sub>4.3</sub>	85.27 <sub>2.2</sub>	83.42 <sub>2.7</sub>	85.36 <sub>2.8</sub>	<b>87.96</b> <sub>2.1</sub>
	Average		64.85 <sub>5.2</sub>	76.73 <sub>11</sub>	63.31 <sub>17</sub>	73.50 <sub>12</sub>	<b>79.58</b> <sub>11</sub>
100%	2,3,4	1	75.27 <sub>8.3</sub>	<b>90.25</b> <sub>4.5</sub>	88.21 <sub>4.9</sub>	90.22 <sub>5.6</sub>	90.01 <sub>4.9</sub>
	1,3,4	2	76.32 <sub>2.9</sub>	84.13 <sub>4.2</sub>	83.76 <sub>3.1</sub>	<b>86.65</b> <sub>3.5</sub>	85.48 <sub>2.3</sub>
	1,2,4	3	62.59 <sub>6.9</sub>	62.18 <sub>10</sub>	56.11 <sub>9.3</sub>	58.27 <sub>9.4</sub>	<b>64.23</b> <sub>9.7</sub>
	1,2,3	4	71.87 <sub>2.5</sub>	88.93 <sub>1.9</sub>	89.08 <sub>2.7</sub>	88.66 <sub>2.6</sub>	<b>89.26</b> <sub>2.5</sub>
	Average		71.51 <sub>5.4</sub>	81.37 <sub>11</sub>	79.29 <sub>13</sub>	80.95 <sub>13</sub>	<b>82.25</b> <sub>11</sub>

LDDG and SAML, the generalisation performance significantly drops with small amounts of labeled data. Note that nnUNet adapts the model design per each run/training set. However, adapting model design for different training data limits the scalability of nnUNet. In Appendix, we also show that nnUNet possibly overfits the source domains in some cases.

**SCGM:** We obtain consistent improvements also on SCGM demonstrating application in other organs. Our model also benefits from the additional 10 unlabeled subjects of each domain leading to better performance overall.

**Ablations:** Here we conduct ablations on key losses crucial to disentanglement and the extraction of good anatomical features for good generalisation performance. We omit ablations on the KL losses as [12, 16] showcase that variational encoding helps to learn robust vector representation for better generalisation. To illustrate that  $\mathcal{L}_{rank}$  helps to disentangle  $\mathbf{z}$  to  $(s, d)$ , and improves performance, we use Distance Correlation ( $DC$ ) [30] to measure disentanglement (lower  $DC$  means higher level of disentanglement). For M&Ms 5% cases, without  $\mathcal{L}_{rank}$ , the average  $DC$  on the test dataset between  $\mathbf{z}$  and  $(s, d)$  is 0.22 (an increase compared to 0.19 with  $\mathcal{L}_{rank}$ ), and the average Dice is 78.54% (a decrease compared to 79.75% with  $\mathcal{L}_{rank}$ ). We also ablate  $\mathcal{L}_{cls}$  and  $\mathcal{L}_{HSIC}$ . The proposed model on M&Ms 5% cases had an average Dice 79.75% but without  $\mathcal{L}_{cls}$ , average Dice drops to 77.45% and without  $\mathcal{L}_{HSIC}$ , average Dice drops to 77.86%.

## 4 Conclusion

We presented a novel semi-supervised meta-learning framework for domain generalisation. Using disentanglement our approach models domain shifts, and thanks to our reconstruction approach to disentanglement, our model can be trained also with unlabeled data. By applying the designed constraints including also the low-rank regularization to the gradient-based meta-learning approach, the model extracts robust anatomical features useful for predicting segmentation

masks in a semi-supervised manner. Extensive quantitative results, especially when insufficient annotated data are available, indicate remarkable improvements compared to previous state-of-the-art approaches. Method performance perhaps may improve by using additional unlabeled data from other domains, which we leave as future work.

## 5 Acknowledgement

This work was supported by the University of Edinburgh, the Royal Academy of Engineering and Canon Medical Research Europe by a PhD studentship to Xiao Liu. This work was partially supported by the Alan Turing Institute under the EPSRC grant EP/N510129/1. We thank Nvidia for donating a Titan-X GPU. S.A. Tsafaris acknowledges the support of Canon Medical and the Royal Academy of Engineering and the Research Chairs and Senior Research Fellowships scheme (grant RCSR1819\8\25).

## References

1. Achille, A., Soatto, S.: Emergence of invariance and disentanglement in deep representations. *JMLR* **19**(1), 1947–1980 (2018)
2. Antoniou, A., Edwards, H., Storkey, A.: How to train your maml. In *proc. ICLR* (2019)
3. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE TMI* **37**(11), 2514–2525 (2018)
4. Bian, C., Yuan, C., Wang, J., Li, M., Yang, X., Yu, S., Ma, K., Yuan, J., Zheng, Y.: Uncertainty-aware domain alignment for anatomical structure segmentation. *MedIA* **64**, 101732 (2020)
5. Campello, V.M., Gkontra, P., Izquierdo, C., Martín-Isla, C., Sojoudi, A., Full, P.M., Maier-Hein, K., Zhang, Y., He, Z., Ma, J., et al.: Multi-centre, multi-vendor and multi-disease cardiac segmentation: The m&ms challenge. *IEEE TMI* (2021)
6. Carlucci, F.M., D’Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalisation by solving jigsaw puzzles. In: *Proc. CVPR*. pp. 2229–2238 (2019)
7. Chatsias, A., Joyce, T., Papanastasiou, G., Semple, S., Williams, M., Newby, D.E., Dharmakumar, R., Tsafaris, S.A.: Disentangled representation learning in cardiac image analysis. *MedIA* **58** (2019)
8. Chen, C., Qin, C., Qiu, H., Tarroni, G., Duan, J., Bai, W., Rueckert, D.: Deep learning for cardiac image segmentation: A review. *Frontiers in Cardiovascular Medicine* **7**(25), 1–33 (2020)
9. Dewey, B.E., Zuo, L., Carass, A., He, Y., Liu, Y., Mowry, E.M., Newsome, S., Oh, J., Calabresi, P.A., Prince, J.L.: A disentangled latent space for cross-site mri harmonization. In: *Proc. MICCAI*. pp. 720–729. Springer (2020)
10. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**(3), 297–302 (1945)
11. Dou, Q., Castro, D.C., Kamnitsas, K., Glocker, B.: Domain generalisation via model-agnostic learning of semantic features. *Proc. NeurIPS* (2019)

12. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-vae: Learning basic visual concepts with a constrained variational framework. In: Proc. ICLR. ICLR (2016)
13. Huang, J., Guan, D., Xiao, A., Lu, S.: Fskr: Frequency space domain randomization for domain generalization. In: proc. CVPR (2021)
14. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: In proc. ICCV. pp. 1501–1510 (2017)
15. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: Proc. ECCV. pp. 172–189 (2018)
16. Ilse, M., Tomczak, J.M., Louizos, C., Welling, M.: Diva: Domain invariant variational autoencoders. In: Proc. MIDL. pp. 322–348. PMLR (2020)
17. Isensee, F., Jaeger, P.F., Full, P.M., Wolf, I., Engelhardt, S., Maier-Hein, K.H.: Automatic cardiac disease assessment on cine-mri via time-series segmentation and domain specific features. In: International workshop on STACOM. pp. 120–129. Springer (2017)
18. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021)
19. Khandelwal, P., Yushkevich, P.: Domain generalizer: A few-shot meta learning framework for domain generalization in medical imaging. In: Proc. DART workshop, pp. 73–84. Springer (2020)
20. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: proc. ICLR (2015)
21. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.: Learning to generalise: Meta-learning for domain generalisation. In: Proc. AAAI (2018)
22. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalisation. In: Proc. ICCV. pp. 1446–1455 (2019)
23. Li, H., Pan, S.J., Wang, S., Kot, A.C.: Domain generalisation with adversarial feature learning. In: Proc. CVPR. pp. 5400–5409 (2018)
24. Li, H., Wang, Y., Wan, R., Wang, S., Li, T.Q., Kot, A.C.: Domain generalisation for medical imaging classification with linear-dependency regularization. *Proc. NeurIPS* (2020)
25. Li, X., Yu, L., Jin, Y., Fu, C.W., Xing, L., Heng, P.A.: Difficulty-aware meta-learning for rare disease diagnosis. In: Proc. MICCAI. pp. 357–366. Springer (2020)
26. Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., Tao, D.: Deep domain generalization via conditional invariant adversarial networks. In: In proc. ECCV. pp. 624–639 (2018)
27. Li, Y., Chen, J., Xie, X., Ma, K., Zheng, Y.: Self-loop uncertainty: A novel pseudo-label for semi-supervised medical image segmentation. In: Proc. MICCAI. pp. 614–623. Springer (2020)
28. Liu, Q., Dou, Q., Heng, P.A.: Shape-aware meta-learning for generalising prostate mri segmentation to unseen domains. In: Proc. MICCAI. pp. 475–485. Springer (2020)
29. Liu, X., Thermos, S., Chatsias, A., O’Neil, A., Tsiftaris, S.A.: Disentangled representations for domain-generalised cardiac segmentation. *International workshop on STACOM* (2020)
30. Liu, X., Thermos, S., Valvano, G., Chatsias, A., O’Neil, A., Tsiftaris, S.A.: Metrics for exposing the biases of content-style disentanglement. *arXiv preprint arXiv:2008.12378* (2020)
31. Llera Montero, M., JH Ludwig, C., Ponte Costa, R., Malhotra, G., Bowers, J.: The role of disentanglement in generalisation. In: Proc. ICLR. ICLR (2021)

32. Locatello, F., Bauer, S., Lucic, M., Raetsch, G., Gelly, S., Schölkopf, B., Bachem, O.: Challenging common assumptions in the unsupervised learning of disentangled representations. In: Proc. ICML. pp. 4114–4124. PMLR (2019)
33. Ma, W.D.K., Lewis, J., Kleijn, W.B.: The hsic bottleneck: Deep learning without back-propagation. In: Proc. AAAI. pp. 5085–5092 (2020)
34. Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalisation via invariant feature representation. In: Proc. ICML. pp. 10–18. PMLR (2013)
35. Paszke, A., Gross, S., Massa, F., Lerer, A., et. al: PyTorch: An imperative style, high-performance deep learning library. In: In proc. NeurIPS. pp. 8026–8037 (2019)
36. Prados, F., Ashburner, J., Blaiotta, C., Brosch, T., Carballido-Gamio, J., Cardoso, M.J., Conrad, B.N., Datta, E., Dávid, G., De Leener, B., et al.: Spinal cord grey matter segmentation challenge. *Neuroimage* **152**, 312–329 (2017)
37. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: In proc. MICCAI. pp. 234–241. Springer (2015)
38. Sharifi-Noghabi, H., Asghari, H., Mehra, N., Ester, M.: Domain generalisation via semi-supervised meta learning. arXiv preprint arXiv:2009.12658 (2020)
39. Tao, Q., Yan, W., Wang, Y., Paiman, E.H., Shamonin, et al.: Deep learning-based method for fully automatic quantification of left ventricle function from cine mr images: a multivendor, multicenter study. *Radiology* **290**(1), 81–88 (2019)
40. Wang, J., Zhou, S., Fang, C., Wang, L., Wang, J.: Meta corrupted pixels mining for medical image segmentation. In: Proc. MICCAI. pp. 335–345. Springer (2020)
41. Yu, X., Chen, Y., Li, T., Liu, S., Li, G.: Multi-mapping image-to-image translation via learning disentanglement. *Proc. NeurIPS* (2019)
42. Zhang, L., Wang, X., Yang, D., Sanford, T., et al.: generalising deep learning for medical image segmentation to unseen domains via deep stacked transformation. *IEEE TMI* **39**(7), 2531–2540 (2020)
43. Zhao, S., Gong, M., Liu, T., Fu, H., Tao, D.: Domain generalization via entropy regularization. In proc. NeurIPS **33** (2020)