

## ARTICLE OPEN



# Addressing materials' microstructure diversity using transfer learning

Aurèle Goetz<sup>1,2,7</sup>, Ali Riza Durmaz<sup>1,3,7</sup>, Martin Müller<sup>4,5</sup>, Akhil Thomas<sup>1,3</sup>, Dominik Britz<sup>4,5</sup>, Pierre Kerfriden<sup>2,6</sup> and Chris Eberl<sup>1,3</sup>

Materials' microstructures are signatures of their alloying composition and processing history. Automated, quantitative analyses of microstructural constituents were lately accomplished through deep learning approaches. However, their shortcomings are poor data efficiency and domain generalizability across data sets, inherently conflicting the expenses associated with annotating data through experts, and extensive materials diversity. To tackle both, we propose to apply a sub-class of transfer learning methods called unsupervised domain adaptation (UDA). UDA addresses the task of finding domain-invariant features when supplied with annotated source data and unannotated target data, such that performance on the latter is optimized. Exemplarily, this study is conducted on a lath-shaped bainite segmentation task in complex phase steel micrographs. Domains to bridge are selected to be different metallographic specimen preparations and distinct imaging modalities. We show that a state-of-the-art UDA approach substantially fosters the transfer between the investigated domains, underlining this technique's potential to cope with materials variance.

*npj Computational Materials* (2022)8:27; <https://doi.org/10.1038/s41524-022-00703-z>

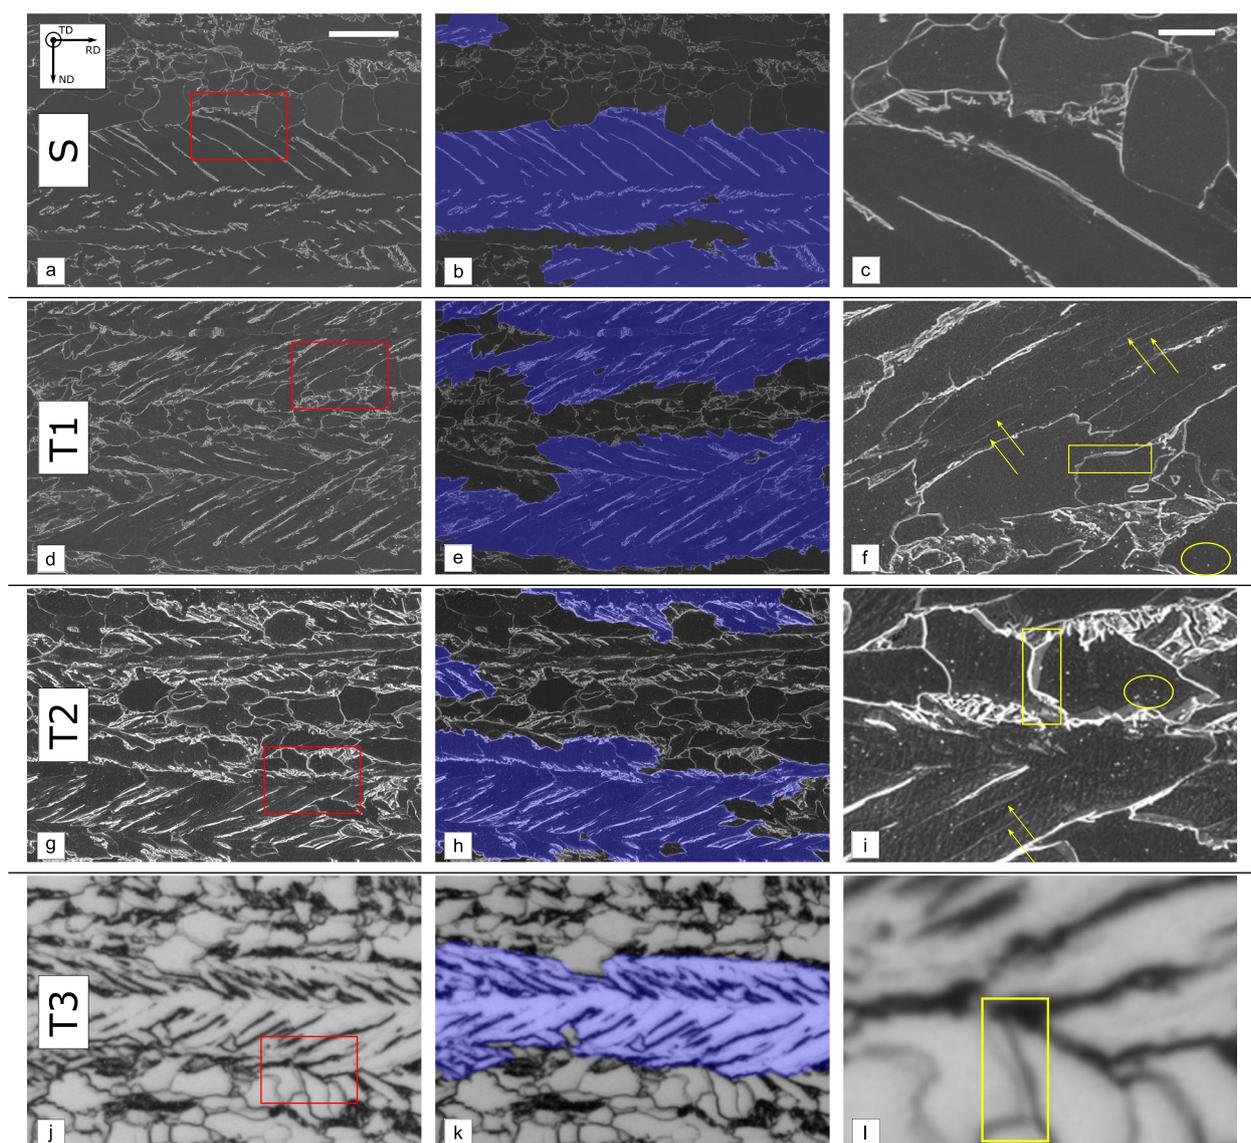
## INTRODUCTION

The inner structure of a material, the so-called microstructure, determines most properties and shows substantial variance depending on its composition and process history. Hence, the quantification and digitization of microstructures play an important role in virtual material design as well as objective and automated quality control. Currently, the quantification of microstructural constituents is performed on material sections, which undergo mechanical polishing and chemical etching routines to expose their inner features in, e.g. light optical microscopy (LOM) or scanning electron microscopy (SEM). Aside from the material itself, this process introduces substantial scatter in the micrographs' appearance. With evolving materials complexity, not only the quantitative micrograph assessment by metallography experts becomes increasingly subjective and costly, but also conventional computer vision algorithms have reached their limits. Therefore, the materials science field recently adopted deep learning (DL) models for this task but is hampered by the amount of available and specifically labeled data. While annotated data is scarce, input micrographs are acquired of various materials, with different settings and optics, by different experts, and with varying processing. This motivates the need for models that can generalize across such data domains without additional labeled data in the target domains. However, deep learning approaches have been shown to exhibit poor generalization capability across miscellaneous data sets<sup>1–3</sup> and materials science data sets<sup>4,5</sup>. Furthermore, in terms of annotated data, materials parameter spaces are often sparsely (i.e., disjointly) populated, contributing to poor model generalization. Advanced generalization techniques could enable model sharing between institutes and their applicability on diverse data sets. Microstructure characterization is one of many tasks, which could benefit from such models.

This work addresses DL-based microstructure quantification and generalization thereof while focusing on a binary segmentation (pixel-wise classification) task introduced in our last paper<sup>5</sup>. It encompasses the binary segmentation of the lath-shaped bainite phase on topography-contrast SEM images of an etched complex phase steel surface. Specimens originate from thermomechanically rolled heavy industrial plates. Resulting microstructure is composed of lath-shaped bainite (later referred to as foreground) as well as polygonal and irregular ferrite with dispersed granular carbon-rich 2nd phase and martensite-austenite (MA) islands (collectively referred to as background). A quick insight into the binary segmentation task is given in Fig. 1b. In this particular case, the annotation process is not only expensive but also complex, to the extent that there is a frequent disagreement between different experts on the exact nature of some phases<sup>6</sup>. This impairs the development of DL models. In our previous work<sup>5</sup>, an extensive effort has been made to combine the SEM micrographs with orientation-contrast from electron back-scatter diffraction (EBSD) images, facilitating a more repeatable and precise labeling process. However, this time-consuming procedure emphasizes the demand for frugal models in terms of labeled training data. Moreover, this work draws the focus on model transferability to different target domains in the context of low source data availability and unavailability of target domain annotations. These are typical boundary conditions in materials science. As a first step, having trained a model on a source dataset, we want to give some insight into the possibility of adopting it to a target domain (domain generalization experiments). In this case, we investigate whether pre-training with domain-extrinsic datasets and subsequent fine-tuning to the source domain can improve generalization to target domains in the low data regime. To this end, we compare different pre-training datasets with models trained from

<sup>1</sup>Fraunhofer Institute for Mechanics of Materials IWM, 79108 Freiburg, Germany. <sup>2</sup>Mines ParisTech PSL University, Centre des Matériaux, 91000 Evry, France. <sup>3</sup>University of Freiburg, 79110 Freiburg, Germany. <sup>4</sup>Department of Materials Science, Saarland University, Saarbrücken 66123, Germany. <sup>5</sup>Material Engineering Center Saarland, 66123 Saarbrücken, Germany. <sup>6</sup>Cardiff University, School of Engineering, Cardiff, UK. <sup>7</sup>These authors contributed equally. Aurèle Goetz, Ali Riza Durmaz.

✉email: aurele.goetz@mines-paristech.fr; ali.riza.durmaz@iwm.fraunhofer.de



**Fig. 1 Visualization of the investigated micrograph sets.** Input images, same superimposed with the annotation, and a detail view are each shown for the source (a–c), target 1 (d–f), target 2 (g–i), and target 3 (j–l) domain. Please refer to Table 6 for the technical details about the different sets. The red frames in the first column highlight the region of interest shown in the third column. Yellow annotations are used in the section “Dataset processing” for discussing major differences between the domains. The scale bar in (a) applies to the first and second column and its width corresponds to 10  $\mu\text{m}$ , while the third column scale bar width amounts to 2  $\mu\text{m}$ .

random initialization. Subsequently, we apply a state-of-the-art unsupervised domain adaptation (UDA) model introduced by Tsai et al.<sup>7</sup> to our phase segmentation task across different domains. With the aid of additional unlabeled target data, this technique attempts to learn domain-invariant features to facilitate the domain transfer. Provided few data, and especially in the materials science field, it is unresolved whether such advanced deep learning techniques can facilitate the transfer across different processing routes or even materials. As example studies, distinct metallographic surface etchings and different imaging modalities were investigated as the domains to bridge. The source domain (S) and the three target domains (T1–T3) are introduced in Fig. 1 and their similarities as well as differences are described in the results.

Instances where deep learning was applied to solve materials science tasks, are fairly scarce. For instance, Azimi et al.<sup>8</sup> perform a microstructure segmentation task on SEM images of dual-phase steels using a fully convolutional deep neural network coupled with a sort of super-pixel voting approach. Holm et al. investigate

various tasks, amongst others microstructure segmentation, on their ultra-high carbon steel (UHCS) dataset using an end-to-end deep learning approach<sup>9</sup>. Recently, Thomas et al.<sup>4</sup> published a study implementing a U-Net architecture for damage segmentation, detecting fatigue-induced extrusions and cracks on steel micrographs. When training with little data, the first prevalent practice in DL is fine-tuning a pre-trained model<sup>10</sup>. This procedure is often used synonymously with transfer learning even though latter now encompasses many further methods. Pre-training a neural network is based on the assumption that its first layers learn similar features regardless of the trained task. Indeed, these layers act as feature extractors, detecting edges, corners, colors, or blobs. Thus, carrying over the weights of a model pre-trained on a large-scale dataset and fine-tuning them on an another task reduces the demand for training data in the latter task while accelerating convergence. Different strategies exist, ranging from full model fine-tuning with a small learning rate to freezing the initial layers of the model and fine-tuning only the last ones.

Utilizing readily available model weights from pre-training with ImageNet, owing to its apparent size and richness, has become the status quo for weight initialization in deep learning. Recently, He et al.<sup>11</sup> questioned the undifferentiated ImageNet weight usage, showing that it only helps in faster convergence but does not enable performance improvement. However, the low-data regime, i.e., when little data is available for the final task, was shown to be exempt from this finding. Specifically, pre-training on ImageNet culminated in improved COCO object detection<sup>12</sup> only when <10k images were used for fine-tuning. However, in material science, this low-data regime is a lasting attendant circumstance of both materials diversity and annotation complexity, making pre-training a suitable strategy to infer robustness and to raise the performance of the trained models<sup>13</sup>. Nevertheless, pre-training dataset selection remains an unsettled issue. In other domains, this has been the subject of recent research<sup>14,15</sup>. The overarching trends seem logical—using a large pre-training dataset as close as possible to the target task dataset proves beneficial. However, the trade-off between data quantity and domain gap is undefined. Cheplygina et al.<sup>14</sup> condensed 12 studies from the medical field in a detailed review, comparing ImageNet with smaller in-field datasets. The conclusions are not consistent, and no clear trend could be identified. Romero et al.<sup>15</sup> studied the effect of different pre-training for chest X-ray radiographs classification, varying the number of target training images from 50 to 2000. Pre-training was conducted using ImageNet or one of two X-ray datasets (220k chest images and 40k images of various body parts). They show that pre-training is always helping in their low-data regime and emphasize that the chest dataset pre-training yields the best results. However, their miscellaneous body parts radiography dataset only performed comparably to ImageNet, showing the trade-off between data amount and domain gap for choosing a pre-training dataset. Gonthier et al. studied similar aspects for artwork classification, where pre-training on ImageNet outperformed another artwork pre-training dataset containing around 80k images, presumably ascribed to the low quantity of the artwork pre-training dataset<sup>16</sup>. Interestingly, a gradual two-step pre-training, using the first ImageNet and then the intermediate artwork dataset, led to further improvement. This suggests that successive pre-trainings could help to bridge domain gaps continuously.

Unsupervised domain adaptation (UDA) describes training a model with labeled data from a source domain and unlabeled data from a target domain to perform well on the latter. This is of major interest when the target domain labeling process is costly or when source domain annotations are readily available. One very good example is the well-known GTA5<sup>17</sup> or SYNTHIA<sup>18</sup> (source) to Cityscapes<sup>19</sup> (target) task. In both cases, source datasets provide synthetic and inherently annotated urban landscapes. A wide branch of the UDA methods is now improving an adversarial learning framework, which was initially proposed by Ganin et al.<sup>20</sup>. The underlying idea is to force the model to adopt a domain-independent feature representation. This is achieved by passing images of both domains into the main model and feeding intermediate layer feature representations into a discriminator. This discriminator then guesses whether the initial data comes from the source or the target domain. Thereby, the discriminator penalizes internal feature representations when they differ substantially for both domains. At the same time, the source data is used in a supervised fashion to train the main model for the task. This adversarial approach has been adapted to semantic segmentation by Tsai et al. where matching of internal features and segmentation masks was performed<sup>7</sup>. Details about this model are given in the section “Unsupervised domain adaptation—an adversarial framework”. To the largest extent, recent articles about UDA report results on the previously mentioned GTA5 to Cityscape task, and only a few applications can be found in other fields such as the adversarial training approach in the medical domain<sup>21,22</sup>. This techniques’ potential for the materials science community has never been showcased to the best of our knowledge. By doing so,

we hope that this work will spark interest in UDA techniques in this field.

The contributions of this work are the following:

- We study the impact of different data augmentation policies, models, and pre-trainings on the segmentation performance reached on the source domain through supervised learning.
- We show the impact of these different training strategies on the domain generalization (i.e., applicability) of a model to an alternate target domain. Specifically, the domain generalization across SEM micrographs of differently contrasted complex phase steel microstructures and across different imaging modalities is addressed.
- We implement a state-of-the-art UDA approach<sup>7</sup> and show its merit in materials science despite data limitations by bridging the aforementioned domain gaps. The UDA frameworks’ suitability for applications beyond the ones presented here is discussed.

## RESULTS

### Supervised learning on the source domain

In this section, models were trained on the source dataset in a supervised fashion and evaluated on the same domain. The subsequent section addresses the results of the same models tested on the target domains (T1–T3 introduced in Fig. 1 and Table 6). These so-called domain generalizations (DG) results act as baselines for the unsupervised domain adaptation (UDA) results addressed later.

In the source domain, the electrolytical Struers A2 etching conducted on the complex phase steel does not emphasize the sub-grain boundaries of the lath-bainite regions. The etching culminates in comparatively slender carbide film appearance overall in the SEM images (see Fig. 1a).

All models were trained on cropped images (tiles) but tested on full-frame images. Details of the data processing and of the investigated segmentation architectures are outlined in the sections “Dataset processing” and “Segmentation architectures”, respectively. The mean intersection over union (mIoU) averaged over foreground (lath-shaped bainite) and background (other phases) classes are used as a metric for segmentation quality. Five-fold cross-validation was implemented. A VGG16 U-Net architecture has been used to test pre-training and data augmentation settings mentioned in the sections “Pre-training and fine-tuning procedure” and “Data augmentation”, respectively. The objective was to maximize the phase segmentation task performance. A batch size of 12 with a constant learning rate of  $\lambda = 5E-3$  was utilized. The models were trained for 200 epochs, except for experiment S.3, which required more iterations due to the random initialization coupled with the extended augmentation pipeline (increased data variance). This experiment was thus extended to 400 epochs. Additionally, a ResNet-101 Deeplabv2 model was trained in a supervised manner. This segmentation architecture corresponds to the one used in the UDA framework, rendering the results comparable. For this architecture, only the pre-training on ImageNet was implemented. The models were trained for 550 epochs with an initial learning rate of  $\lambda = 1E-3$  and a polynomial decay with a decay factor of 0.9. Note that a larger epoch number was required for this architecture because of ResNet-101’s higher number of model parameters. The mentioned training epoch numbers were chosen to ensure proper convergence of the validation mIoU. In practice, verifying convergence was performed manually and in a cautious manner. The results are given in Table 1. Please note that given ranges correspond to the five folds’ standard deviations. Additionally, a repeatability study was conducted where the mean mIoU’s standard deviations for the

**Table 1.** Performance of the models trained on the source dataset.

Model	Pre-training	Augmentation	Exp.#	mIoU [%]
VGG16 U-Net	Random init.	–	S.1	75.8 ± 2.4
		Basic	S.2	77.8 ± 2.0
		Extended	S.3	78.7 ± 2.6
	ImageNet	–	S.4	76.1 ± 3.0
		Basic	S.5	78.9 ± 2.2
		Extended	S.6	79.2 ± 2.2
	NanoSEM	–	S.7	76.9 ± 2.6
		Basic	S.8	79.2 ± 2.2
		Extended	S.9	80.2 ± 2.4
	ResNet-101 DeepLabv2	ImageNet	Extended	S.10

The given mIoU is always averaged over the five folds used for cross-validation. NanoSEM refers to the two-step pre-training introduced in the section “Pre-training and fine-tuning procedure.” Details about the augmentation pipelines are given in the section “Data augmentation” and in Supplementary Table 2.

experiments S.6 and S.9 reached 0.4% and 0.2%, respectively, indicating statistical separability of the results.

The results in Table 1 demonstrate that pre-training helps the model in this low-data regime. Indeed, we observe the systematic trend that the two-stage NanoSEM pre-training outperforms the ImageNet one, which in turn surpasses random initialization, regardless of the used augmentation pipeline. Similarly, the extended augmentation pipeline consistently grants better results than the basic one, which in turn compares favorably to the unaugmented result. Thus, the best performance is observed with the NanoSEM pre-training and the extended data augmentation pipeline, reaching 80.2% mIoU as an average over the five folds (S.9). Experiment S.10 consists in the supervised training of the DeepLabv2 architecture on the source domain. Concerning the ImageNet pre-training case, DeepLabv2 slightly outperforms the VGG16 architecture on this task (cf. S.10 to S.6).

In the following section, all mIoU values in Table 1 are utilized as reference values to compute equivalent DG experiments’ relative domain transferability (RDT) metrics. The RDT metric measures how the model performance on the target domain compares to the one on the source relatively and is introduced in the section “Evaluation metrics”. Analogously, model S.10 serves as the reference to compute the RDT for all the UDA-based models.

### Model generalization and adaptation to target domains

In the following, the main objective is to achieve good segmentation results on the target datasets introduced in Fig. 1 despite the unavailability of labeled training data in these domains. The microscopic differences between the source and three target datasets become evident in the detail views in Fig 1c, f, i, l. Subsequent passage addresses the three target datasets with continuously increasing domain shifts consecutively.

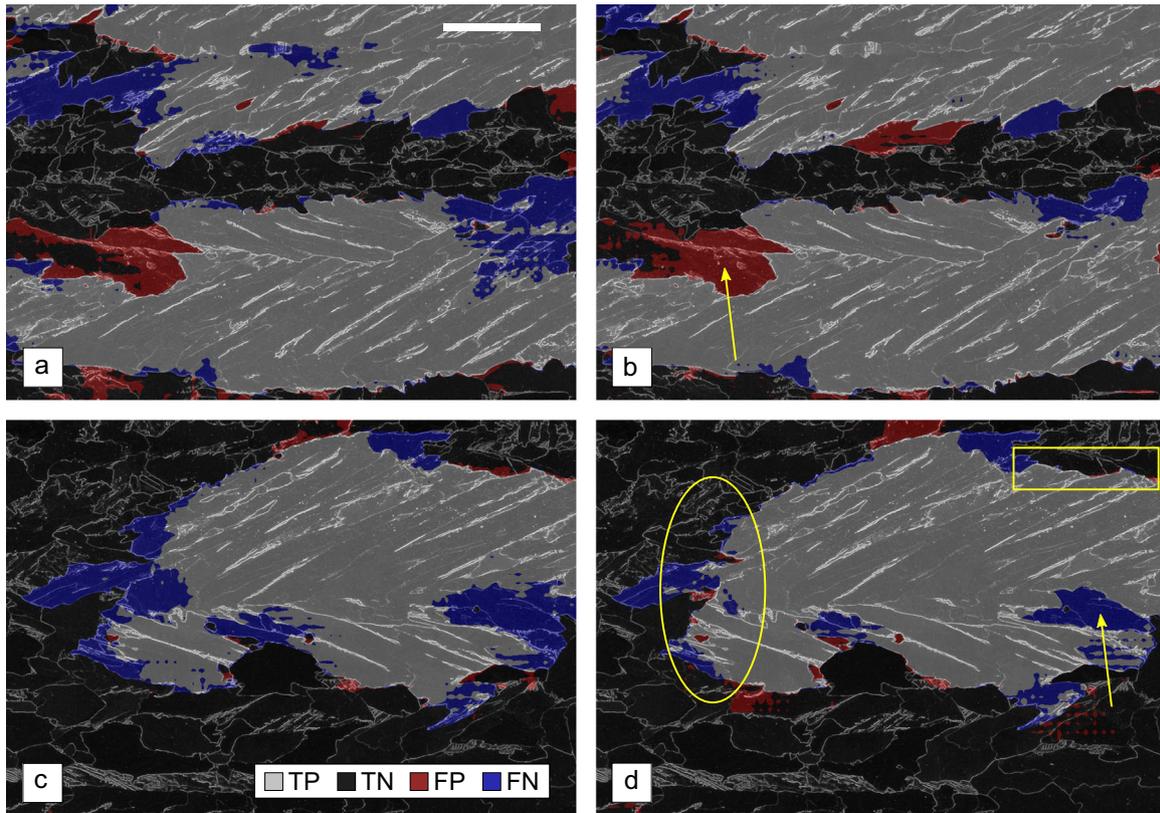
In contrast to the source domain, the Nital etchant applied in T1 reveals the hierarchical structure in the lath-bainite regions (see arrow annotations in Fig. 1f). Moreover, this domain is accompanied by minor etching artifacts as indicated by ellipse annotations there. The contrast of the T1 SEM images is more pronounced as opposed to the source dataset. Overall, the T1 domain represents a small domain shift with respect to the source domain. Models trained in the section “Supervised-learning on the source domain” were tested on this target domain, and the UDA framework of Tsai et al.<sup>7</sup> described in the section “Unsupervised domain adaptation—an adversarial framework” was used to apply

**Table 2.** Performance of the models trained on the source domain evaluated on the target 1 domain along with the UDA reached performance.

Model	Pre-training	Augmentation	Exp.#	mIoU [%]	RDT [%]	
DG-VGG16 U-Net	Random init.	–	T1.1	75.7 ± 2.2	–0.1	
		Basic	T1.2	77.5 ± 1.4	–0.4	
		Extended	T1.3	73.5 ± 3.9	–4.2	
	ImageNet	–	T1.4	77.2 ± 1.0	1.4	
		Basic	T1.5	79.8 ± 1.1	1.2	
		Extended	T1.6	81.6 ± 1.9	3.1	
	NanoSEM	–	T1.7	77.2 ± 0.6	0.4	
		Basic	T1.8	80.6 ± 0.8	1.8	
		Extended	T1.9	81.5 ± 0.7	1.6	
	DG-ResNet-101 DeepLabv2	ImageNet	Extended	T1.10	82.2 ± 1.5	3.6
	UDA-ResNet-101 DeepLabv2	ImageNet	Extended	T1.11	84.7 ± 0.7	6.7

The given mIoU is always averaged over the five folds used for cross-validation. NanoSEM refers to the two-step pre-training introduced in the section “Pre-training and fine-tuning procedure.” Details about the augmentation pipelines are given in the section “Data augmentation” and in Supplementary Table 2.

domain adaptation towards this domain. These UDA models were trained with a polynomial learning rate ( $\lambda = 1E-3$ , decay factor of 0.9) for 3000 epochs and a batch size of 8 both for source and target tiles. During this procedure, no target annotations were supplied to the model. Once again, the number of epochs has been chosen to reach a satisfying convergence of the validation mIoU. The results are given in Table 2 and some visualizations are provided in Fig. 2 to show the advantage of the UDA method over domain generalization. Looking at the results from Table 2, it first appears surprising that most source domain trained models perform excellent on T1, sometimes even exceeding the source performances (cf. Table 1), resulting in positive RDT values. This will be discussed in the section “Discussion”. Secondly, the general tendency that pre-training helps domain generalizability is evident, as random-initialized models yield the lowest RDT values. Despite the aforementioned good model transferability between source and T1, UDA surpasses DG clearly (compare T1.11 to T1.10). Aside from the 2.5% mIoU increase in favor of the UDA approach, the obtained model is more balanced in terms of class mispredictions. Figure 2 illustrates this phenomenon with two examples. Without UDA, the models transferred from scratch exhibit a skewed behavior towards the background class, thus giving substantially more false negatives than false positives. In Fig. 2b, d, the amount of false negatives is reduced while false positives increase slightly, giving an overall better segmentation and an improved phase fraction estimation in the UDA case. Additionally, it can be observed in Fig. 2 that both models (DG and UDA) associate the parallel carbide features with lath-bainite correctly, whereas the DG model particularly struggles to find the proper instance boundaries. The main improvement offered by UDA lies in better boundary localization. For instance, on the left side of Fig. 2d (ellipse annotation), the predicted boundary of the lath-bainite constituent substantially shifts to the left, improving the segmentation to a large extent. The rectangle annotation shows a region example where the phase boundary is localized very accurately, underlining the enhanced model’s confidence.



**Fig. 2 Predictions on the target 1 test images.** The left column (a), (c) gives the predictions of a DeepLabv2 model trained on the source dataset (cf. experiment T1.10). The right column (b), (d) give the predictions of a DeepLabv2 model trained with the UDA framework on the source and target 1 datasets (cf. experiment T1.11). Annotations are used for the description of the results in the text. TP, TN, FP, and FN stand for true positives, true negatives, false positives, and false negatives, respectively, the positive class being the lath-bainite foreground. The scale bar width corresponds to 10  $\mu\text{m}$ .

While some error cases persist (see arrow annotations) they can be explained by misleading peculiarities of relevant regions. For instance, the false positive in Fig. 2b can be traced back to the broader morphology of the contained carbon-rich constituents and their orientation being parallel to the adjacent laths'. On the contrary, the false-negative region in Fig. 2d shows very slender carbide films, incentivizing the model to classify the region as background. Both falsely classified regions remain equally challenging to categorize for metallographers. Furthermore, some checkerboard patterns appear clearly in Fig. 2a and c. These periodic patterns arise in regions of model uncertainty and originate from the single bilinear interpolation in the DeepLabv2 architecture used to restore the input image resolution after the encoding stage<sup>23</sup>. In such uncertain areas, the segmentation could be improved by combining models in a voting scheme (i.e., bagging) for getting better final predictions. Such a bagging strategy will be briefly discussed in the section "Discussion". In contrast, the VGG16 U-Net results (not shown) do not exhibit such artifacts due to its comparatively sophisticated decoder architecture composed of alternating nearest-neighbor interpolation and convolution steps.

Analogous to the T1 domain, the T2 dataset is addressed next, representing a larger domain shift with respect to the source domain. Compared to T1, hierarchical sub-grain features are less visible in T2. Moreover, as highlighted by the rectangle annotations in Fig. 1i, some grain boundaries appear faded due to their grain boundary inclination and the progressed etching state. Along with the pronounced contrast and wider carbide films, this renders it evident that T2 was over-etched. On the other hand, imaging-induced statistical differences are present particularly in this domain, where the lath-bainite phase fraction deviates from the source significantly (see Table 6). Results are given in Table 3

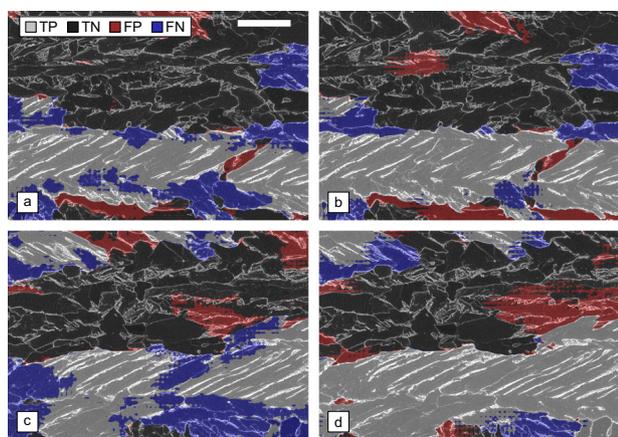
and visualizations in Fig. 3. Regarding domain generalization (DG), two major observations can be made. First, exactly as for T1, it appears that pre-training helps. Second, as opposed to T1, the basic augmentation pipeline improved over the extended one. Moreover, there is a significant RDT drop in T2, indicating this to be a more challenging task for the UDA framework compared to T1. In this case, UDA gives a pronounced advantage, exceeding the DG DeepLabv2 model by 6.3% mIoU. Figure 3 displays how UDA corrects the skew towards the background class, leading to a segmentation that is better and more balanced in terms of misclassifications. While difficult regions at the top of Fig. 3a, b remain challenging for the UDA network, larger lath-shaped regions are segmented more comprehensively (cf. bottom of Fig. 3a and b or c and d). The classification of these difficult regions at the top of Fig. 3a, b is equally complicated for human experts.

Lastly, the T3 dataset is investigated where the objective is to bridge the domain gap between imaging modalities. Similar to T2, the etching state in T3 is advanced as indicated by faded-appearing grain boundaries highlighted by the rectangle annotations in Fig. 1l. In these bright-field LOM images, carbide film morphology cannot be resolved and image features differ substantially due to the modality change. Rather than passing bright-field images to the different networks, pixel intensities were inverted for the reason discussed in the section "Discussion". Along with the UDA performance, the domain generalization DeepLabv2 result is reported as the sole baseline in Table 4. The VGG16 U-Net results on the T3 domain are omitted due to the observed large mIoU scatter between the five-folds, preventing any form of conclusion, and overall poor model performance. While domain generalization seems compromised on this target dataset, a tremendous improvement of 23.6% mIoU is

**Table 3.** Performance of the models trained on the source domain evaluated on the target 2 domain along with the UDA reached performance.

Model	Pre-training	Augmentation	Exp.#	mIoU [%]	RDT [%]
DG-VGG16 U-Net	Random init.	–	T2.1	51.4 ± 4.9	–32.1
		Basic	T2.2	57.6 ± 4.9	–26.0
		Extended	T2.3	52.5 ± 8.5	–31.5
	ImageNet	–	T2.4	58.1 ± 2.4	–23.7
		Basic	T2.5	63.8 ± 1.1	–19.2
		Extended	T2.6	52.8 ± 11.8	–33.2
	NanoSEM	–	T2.7	59.7 ± 3.0	–22.3
		Basic	T2.8	61.2 ± 3.7	–22.7
		Extended	T2.9	60.3 ± 3.3	–24.8
DG-ResNet-101 DeepLabv2	ImageNet	Extended	T2.10	61.0 ± 3.7	–23.2
UDA-ResNet-101 DeepLabv2	ImageNet	Extended	T2.11	67.3 ± 2.0	–15.3

The given mIoU is always averaged over the five folds used for cross-validation. NanoSEM refers to the two-step pre-training introduced in the section “Pre-training and fine-tuning procedure”. Details about the augmentation pipelines are given in the section “Data augmentation” and in Supplementary Table 2.



**Fig. 3 Predictions on the target 2 test images.** The left column (a), (c) gives the predictions of a DeepLabv2 model trained on the source dataset (cf. experiment T2.10). The right column (b), (d) give the predictions of a DeepLabv2 model trained with the UDA framework on the source and target 2 datasets (cf. experiment T2.11). TP, TN, FP, and FN stand for true positives, true negatives, false positives, and false negatives, respectively, the positive class being the lath-bainite foreground. The scale bar width is equivalent to 10  $\mu\text{m}$ .

experienced when using the UDA method. Moreover, the scatter over the five folds is reduced substantially. Considering the few prediction examples in Fig. 4, it is apparent that UDA turns a completely unusable model into a convincing one without requiring any labeled data in the target domain.

## DISCUSSION

Achieving the results above with this source data set (see Table 1) contrasts the common preconception that DL techniques require a large training data quantity. Specifically, satisfactory results on this complex microstructure segmentation task were attained despite using barely more than 100 tiles (27 native SEM images) for training. This is aligned with our prior findings<sup>5</sup> and can be explained by two factors. First, the data has been acquired in a very repeatable manner. Images of each dataset (i.e., domain) were drawn from an individual etched specimen, and reproducible imaging conditions were applied among images of the same domain. It has to be underlined that this results in comparatively

low intra-domain variance, which might not be representative for large-scale datasets acquired by multiple operators or even different institutions. Second, the native micrographs exhibiting a high resolution and rich feature density, 27 such images still represent an appropriate learning foundation for our binary segmentation task and the present microstructure variance.

The results in Table 1 emphasize that pre-training improves the performance of the trained models, giving up to 1.5% mIoU improvement between the best random-initialized model (S3) and the best NanoSEM pre-trained model (S9). Moreover, pre-trained weight initialization led to faster model convergence. Even with further training, it was observed that random initialized models did not catch up to the pre-trained ones. Therefore, the dataset is situated in the low-data regime mentioned by He et al.<sup>11</sup>, where pre-training elevates the performance irrespective of training iterations. In addition, the two-step NanoSEM pre-training shows better performance compared to ImageNet in all cases. In contrast, pre-training solely on NanoSEM resulted in poor model performance (not reported here). These observations are in line with Gonthier et al.<sup>16</sup>, who performed a two-step pre-training process as well to gradually bridge the domain gap between real-world image datasets and artwork datasets. Please note that this gradual pre-training procedure, compared to conventional pre-training, introduces a further learning rate hyperparameter, which is known to affect the final task performance sensitively. Therefore, relatively more learning rate optimization is required for the pre-training and fine-tuning steps. More details on the learning rate variation are given in Supplementary note 1.

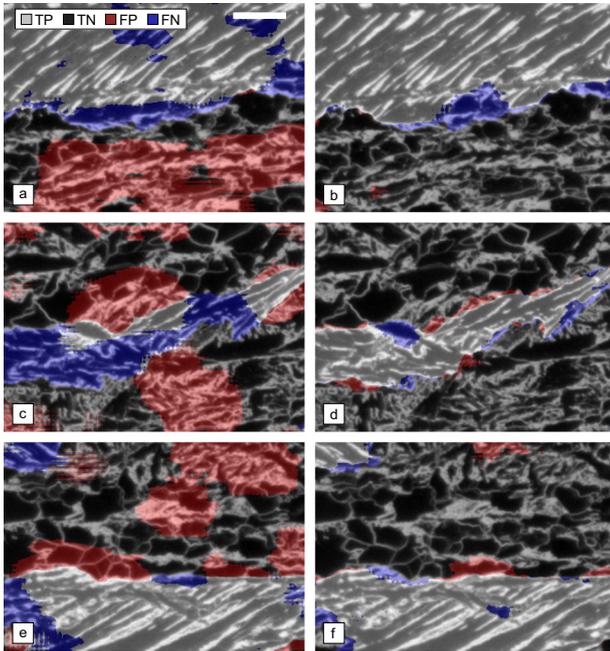
This two-stage pre-training on ImageNet and NanoSEM might be called into question considering the slight performance increase over sole ImageNet pre-training (1% improvement from experiment S6 to S9). However, it has to be emphasized that NanoSEM is far from being the optimal pre-training dataset for our target task. Indeed, it entails the following limitations:

- Structures in certain classes such as MEMS, patterned surfaces, and tips contain shape-related features but barely any apparent microstructural ones, making the learned weights possibly sub-optimal for the final task.
- The image formation is complicated in SEM and depends on a multitude of settings. Most images in the NanoSEM dataset, depending on the class, were either acquired with an Everhart–Thornley (SE2) detector or in-lens detector. Therefore, concerning the detector class, only the latter portion of the pre-training data matches the acquisition of both source and target SEM datasets. Generally, the SE2 detector exhibits a

**Table 4.** Performance of the DeepLabv2 model trained on the source domain evaluated on the target 3 domain along with the UDA reached performance.

Model	Pre-training	Augmentation	Exp.#	mIoU [%]	RDT [%]
DG-ResNet-101 DeepLabv2	ImageNet	Extended	T3.1	49.7 ± 7.6	−37.4
UDA - ResNet-101 DeepLabv2	ImageNet	Extended	T3.2	73.3 ± 0.7	−7.7

The given mIoU is always averaged over the five folds used for cross-validation.



**Fig. 4 Predictions on the intensity-inverted target 3 test images.** The left column (a), (c), (e) give the predictions of a DeepLabv2 model trained on the source dataset (cf. experiment T3.1). The right column (b), (d), (f) give the predictions of a DeepLabv2 model trained with the UDA framework on the source and target 3 datasets (cf. experiment T3.2). TP, TN, FP, and FN stand for true positives, true negatives, false positives, and false negatives, respectively, the positive class being the lath-bainite foreground. The scale bar width corresponds to 10  $\mu\text{m}$ .

more pronounced topography sensitivity due to its location and orientation, while the in-lens detector combines surface topography and, to a lesser extent, material contrast.

- NanoSEM represents a classification task. Hence, only the encoder of our segmentation model could be pre-trained. While it can be assumed to be domain gap dependent, there is no quantitative understanding to which extent and how many layers of a segmentation model would benefit from such a decoder pre-training. Conventional pre-training was reported to primarily help the models' first layers to learn general features<sup>10</sup>.
- For pre-training standards, NanoSEM is comparatively small.

Despite these inadequacies, the underlying rationale of utilizing this NanoSEM dataset for pre-training was that high-level characteristics such as noise levels and typical image textures can be learned. However, we assume that a more extensive dataset involving a micrograph segmentation task of arbitrary alloy would prove beneficial over NanoSEM. For instance, the ultra-high carbon steel micrographs collection subset introduced in ref.<sup>9</sup> would have been appropriate if not for its low quantity. A more promising candidate could be the recently published

Aachen–Heerlen annotated steel microstructure dataset<sup>24</sup> containing annotated martensite-austenite islands. While this datasets' annotations exhibit a systematic offset at instance boundaries potentially causing adverse effects during learning, such tendencies presumably can be unlearned during fine-tuning. The success of pre-training and fine-tuning motivates the demand for more publicly available datasets in the material science field.

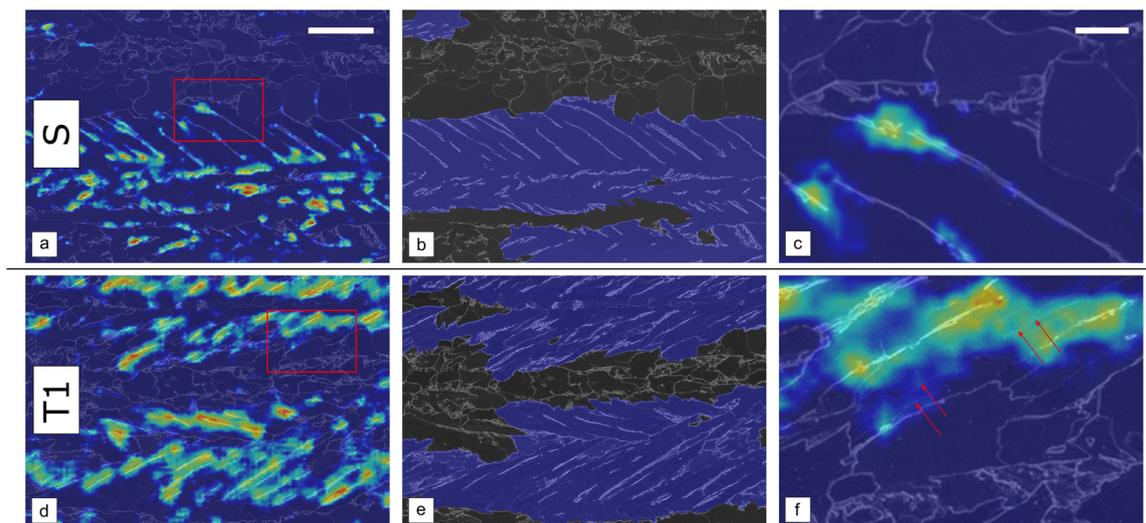
With respect to data augmentation, a systematic increase in performance is observed when applying the two pipelines, which is not surprising considering the low amount of data used for training the models.

Lastly, it appears that the DeepLabv2 architecture achieves better results compared to the U-Net one (compare experiments S.6 and S.10). However, the improvement is relatively small considering the model size difference. A possible explanation is that our segmentation task does not exploit the full representation power of the ResNet-101 DeepLabv2 architecture. Additional results (not reported here) suggested that using larger tiles and thus increasing the context given to the model might enhance the DeepLabv2's performance. Indeed, this architecture is built to learn large receptive fields thanks to its dilated convolutions. Hence, it can benefit from long-distance correlations on tiles when learning the segmentation task.

Pre-training not only helps improve performances on the source domain but also brings generalizability to the trained models. Generally, pre-trained models perform better on target domains (cf. RDT values of pre-trained models in Tables 2 and 3 compared to the corresponding random-initialized experiments). Interestingly, further experiments (not reported here) have shown that models pre-trained with NanoSEM for 100 epochs outperform the 200-epoch ones in terms of domain generalizability towards T2 by around 2% mIoU. It can be assumed that, during prolonged training, weights are tweaked such that very dataset-specific features are progressively replacing general ones.

Contrary to T1, the basic augmentation pipeline consistently outperforms the extended one for T2. This poor domain generalizability (Exp.# T2.3, T2.6, T2.9) suggests that models were rendered invariant to some task-relevant features of T2 when trained with the extended pipeline. It should be emphasized again that this pipeline was optimized for the source domain, which exhibits a substantially wider domain gap with T2 compared to T1. Extended data augmentation causing a drop of generalizability has previously been observed in ref.<sup>4</sup>.

Concerning the UDA framework, the obtained results are very encouraging. For T1, it appeared that due to the minimal domain shift with respect to the source, transferring source-trained models was already performing satisfactorily (cf. experiments T1.6, T1.9). Hence, this problem posed to UDA is not overly challenging. The DG models' achieved mIoUs on this target domain even exceed the source mIoUs. Presumably, this can be attributed to the additional parallel features introduced by the subgrain boundaries (see Fig. 1f, arrow annotations), rendering the prediction easier. This was verified using the GradCAM network visualization technique, which computes the network gradients at a specific layer with respect to a target class and thereby estimates pixel-wise activation. For more information, we refer to ref.<sup>25</sup>.



**Fig. 5** Network visualizations of the ResNet-101 DeepLabv2 model, on the source (S.10, a–c) and target 1 (T1.10, d–f) domain with lath-shaped bainite labels (b) and (e) given as reference. The heat maps indicate regions that were taken into consideration at layer 3\_16 of the ResNet-101<sup>40</sup> to classify the lath-shaped bainite regions. In the detail views (c), (f), which are the same as in Fig. 1, heat maps in the target 1 domain are more extensive and additionally incorporate subgrain boundaries (red arrow annotations). The scale bar in (a) applies to the first and second column and its width corresponds to 10  $\mu\text{m}$ , while the third column scale bar width amounts to 2  $\mu\text{m}$ .

In Fig. 5 this technique has been applied to the source and T1 domain to determine regions that were deemed important by the (same) ResNet-101 model (trained on source). It is clear that the activation is more extensive in the T1 domain and additionally involves the subgrain boundaries inside the laths (red arrow annotations). This supports the theory that these additional features induced by Nital etching are beneficial for the model. In our prior study, we discovered that image downscaling for the source domain culminates in a performance increase since the pixel gap between carbides at lath boundaries is reduced, and information loss is minimal<sup>5</sup>. Therefore, the parallelism of these features can be assessed at earlier network layers. The GradCAM results on layer 3\_16 indicate that the hierarchical microstructure and internal subgrain boundary features revealed in T1 can help to bridge the otherwise feature-sparse bainitic ferrite regions in the source domain to improve learning. Note that the activation in Fig. 5c is high where parallel carbide films are in close vicinity. These parallel carbide films being decisive features indicates that these trained models' performance could be compromised when evaluated on cross-sections in rolling or normal direction due to their distinct microstructural patterns<sup>26</sup>. Moreover, considering the small test sets, it cannot be excluded that the T1 test images potentially being easier to predict on average compared to the source domain test set contributes to the better T1 performance. Overall the UDA framework gave a 2.5% mIoU boost on this target domain (T1.11 compared to T1.10). Furthermore, training the UDA framework with T1 as the target domain gave models that perform better on the source domain with 79.7% mIoU, granting a 0.3% boost compared to experiment S.10. Similarly, such small domain gaps led to the same observation in the context of urban image segmentation<sup>27</sup>.

On the other hand, T2 and T3 have broader gaps with respect to the source domain due to stronger etching and different imaging modalities, respectively. Despite the large domain gap of the T3 dataset with the source one, UDA performed substantially better on this dataset than on T2, culminating in a 23.6% mIoU improvement over the DG experiments, which is mirrored in Fig. 4. As a reference, fully supervised training on T3 presented in our prior work<sup>5</sup> achieved 79% mIoU. Therefore, employing this UDA framework falls short only by 6% mIoU compared to this fully supervised reference, despite not relying on target labels. This result with respect to bridging modalities is promising and in line

with literature where domain adaptation in the medical field was successfully applied to transition between computer tomography and magnetic resonance imaging<sup>21</sup>. Note that UDA models trained with T2 and T3 datasets scored 78.8% and 75.8% mIoU on the source domain, respectively, falling short compared to the fully supervised source domain training (Exp.# S.10: mIoU = 79.4%). Indeed, UDA reaches a compromise between source and target, which is detrimental to the source domain when large domain gaps with the target are involved. This observation confirms that T1–T3 are gradually increasing the domain shift with respect to the source domain. The difference between UDA with respect to T2 and T3 target domain performance could be attributed to the 5 $\times$  larger data amount available for the latter set (see Table 6), where 48 unannotated training tiles for T2 might be insufficient. Another reason could be the phase fraction of the T2 dataset, which is substantially lower than the source dataset. This assumption will be discussed later. Lastly, we consider it unlikely that the SEMs different distortion and noise level characteristics (see the section “Specimen fabrication and image acquisition methodology”) are causing this difference since the UDA framework can cope with different modalities and corresponding data augmentations were applied.

Additionally, to improve over the individually trained UDA models, we also implemented a bagging strategy. This consists in averaging the predictions of multiple models to give an improved segmentation. In our case, we use the models trained with the five different folds and achieve 85.9%, 70.2%, and 75.0% mIoU, which results in 1.2%, 2.9%, and 1.7% increase over the best results presented in the T1–T3 result tables, respectively. The larger improvement for T2 can be attributed to its comparatively weak individual classifiers, making the bagging paradigm relatively more profitable.

Lastly, as opposed to DG models, which are frequently biased towards a class (cf. Figs. 2 and 3), UDA leads to more balanced models. Consequently, this characteristic improves the estimation of phase fractions or other metrics, which do not require full details of the segmentation mask. As an example, predicted phase fractions on the different target datasets with DG and UDA are reported in Table 5. It is evident that UDA improves phase fraction estimation in all cases and additionally reduces the scatter substantially (e.g., by a factor of 10 for T3). The large observed

scatter of DG models for T2 and T3 predictions can be explained by the different training folds each leading to skewed predictions in favor of either background or foreground class. Using UDA systematically reduced the skewed behavior and yielded models that are only slightly biased towards the background class. This suggests that UDA-based models, when applied to the target domains, misinterpret some foreground class features. One potential cause could be incompletely bridged gaps with the source domain. Another possibility for this slight, consistent lath-bainite underestimation could be the labeling process. Unlike the target domains, the source was labeled with supporting EBSD images, which could conceivably result in different annotation patterns. Lastly, one should recall that the fairly small test set size renders the results very sensitive concerning any labeling inconsistencies. For these reasons, the impact of intra-rater reliability is presumably elevated.

In the following, limitations and potentials of the implemented unsupervised domain adaptation framework are discussed. Firstly, it has to be emphasized that adversarial-based frameworks such as AdaptSegNet suffer from training instability, rendering them rather laborious to tune and affecting training repeatability. Facing these training pitfalls while working with low quantity data hampers the deduction of relationships from ablation studies and hyperparameter studies. Furthermore, given the implementations and hardware at hand, typical UDA training runs for 6 h, whereas a DeepLabv2 supervised fine-tuning takes only 45 min.

Another constraint was witnessed when trying to bridge the gap between SEM and LOM data (T3). Indeed, the learning process failed when feeding the model with bright-field LOM tiles, motivating us to invert their pixel values. As the segmentation architecture shares weights between source and target data, it might be difficult to learn filters, which perform well on both modalities while keeping internal features independent of the originating domain.

Moreover, this AdaptSegNet<sup>7</sup> framework is built on the strong prior assumption that the source and target datasets are sharing the same label space distribution. This poses a boundary condition for the segmentation model to give good predictions and fool the discriminator simultaneously. In case of pronounced label space deviations between the source and target datasets, the discriminator should hypothetically quickly learn how to differentiate the

segmentation masks, hampering the transfer learning process. Such a label distribution shift could be due to different phase morphology or phase fractions. For this purpose, the phase fractions are provided in Table 6. Aside from generally low data quantity in the T2 domain, the lath bainite phase was not oversampled during image acquisition as opposed to the other domains. Therefore, selecting a suitable training tile subset to match the source domains' 53% mean lath-bainite phase fraction was unfeasible. Taking the phase fraction histograms of the training tiles (based on expert-reviewed pseudo-labels) into consideration, the discriminator seemingly should learn the tendency that the images from the T2 domain generally show a smaller lath-bainite content. Initially, we considered this to be the primary reason for UDA being more beneficial for T3 compared to T2. However, an additional experiment where we varied the lath-bainite phase fraction of the LOM training dataset provided to the UDA framework invalidated this hypothesis. Specifically, we sampled another target 3<sup>v</sup> train set with a lower mean phase fraction ( $\phi_{\text{train},3^v} = 0.28$  similar to T2) and trained a model with it in the UDA framework. A common test set exempt from both LOM training sets was created for testing purposes consisting of 12 test images with a mean phase fraction of 40%. This LOM test set phase fraction was chosen to be the average between the training phase fractions such that the influence of target train-test shifts could be excluded and domain phase fraction shifts during training could be investigated. The results showed that models trained with T3 reached  $78.0 \pm 1.0$  mIoU, whereas those trained with T3<sup>v</sup> scored  $77.4 \pm 1.8$ . Note that the former value deviates from the result provided in Table 4 due to the distinct test set. This suggests that distinct phase fraction training data of the different domains does not hamper the UDA training. One thing to underline is that the discriminator receives predictions from the segmentation model (see Fig. 7) rather than actual annotation masks, which complicates the distinction based on phase fraction histogram separability. Definitely, the robustness of this adversarial process concerning space label non-conformity is highly auspicious for materials science tasks. For instance, this appears promising for generalizing to different alloys or processing routes, as the phase topology and morphology then can be altered significantly between the source and target sets. Alternatively, if problems due to too different phase distributions were to arise, these could be overcome by feeding the discriminator with tile sub-patches selected based on pseudo-labels to balance both sets' apparent tile phase fractions artificially.

Despite the aforementioned limitations of AdaptSegNet, this framework has been successfully applied as part of this work, even in a low-data amount scenario. Substantial improvements in performance on the target datasets were observed despite only providing few tens of unlabeled micrographs. Even modality transfers from SEM to LOM could be facilitated successfully with such data. The UDA frameworks' insensitivity with respect to different phase fractions in source and target domains yields hope to enable generalization across different alloys and heat treatments. We consider this framework a good trade-off between

**Table 5.** Estimated phase fractions on the test sets of the different target domains with DG and UDA models.

Target domain	$\phi_{\text{DG}}$ [%]	$\phi_{\text{UDA}}$ [%]	$\phi_{\text{label}}$ [%]
Target 1 (T1)	$39.8 \pm 1.0$	$41.3 \pm 0.9$	43.6
Target 2 (T2)	$48.4 \pm 10.7$	$44.6 \pm 4.9$	46.7
Target 3 (T3)	$51.3 \pm 21.8$	$42.8 \pm 2.6$	44.6

Values are averaged over the five models trained on the different folds and standard deviations are provided.

**Table 6.** Micrograph set descriptions.

Set name	Etching	Imaging setup	# train tiles	# test images	$\phi_{\text{train}}$
Source (S)	Struers A2 electrolytic etching; o	SEM setup 1	112	7 (28 tiles)	0.53
Target 1 (T1)	Nital (2 vol.% HNO <sub>3</sub> ); o	SEM setup 1	80	5 (20 tiles)	0.54
Target 2 (T2)	Nital (2 vol.% HNO <sub>3</sub> ); +	SEM setup 2	48	3 (12 tiles)	0.29
Target 3 (T3)	Nital (2 vol.% HNO <sub>3</sub> ); +	LOM	234	50 (200 tiles)	0.53

Mean lath-bainite phase fractions  $\phi_{\text{train}}$  of the train tiles are given. Note that no annotation was available for the target 1 and target 2 training sets. Hence, their phase fractions were estimated based on pseudo-labels and confirmed by a metallography expert. Symbols "o" and "+" refer to normal and deliberately prolonged etching duration, respectively.

complexity and reached performance, and therefore an excellent introduction to UDA for the material science community. However, in view of the fast-paced ML research, it has been outperformed on the GTA5-to-cityscape task. Several improvements have been published over the past 3 years, most of the time using the work of Tsai et al.<sup>7</sup> as a reference and starting point<sup>28–32</sup>. All these studies rely on the GTA5-to-Cityscapes reference task, and some of them exploit specific characteristics of these datasets<sup>29</sup>. Therefore, the approaches in these works are not directly applicable for our binary segmentation task but potentially relevant for other material science tasks. Nevertheless, some models could potentially improve over AdaptSegNet in our setting. One example is the ADVENT model of Vu et al.<sup>30</sup>, which makes use of entropy maps instead of segmentation maps as input for the discriminator. It encourages entropy minimization in the target domain by matching the source and target entropy distributions. This entropy minimization paradigm is borrowed from semi-supervised learning. Also inspired by semi-supervised techniques, Pan et al.<sup>28</sup> implement the AdaptSegNet framework with an extra pseudo-labeling step. The easiest-to-predict half of the target data is pseudo-labeled in a first training iteration and then utilized as “source” domain data for a second training, using the rest of the target data as the target domain. This is motivated by the intra-domain variance in the target domain. In our case, this approach would probably not be overly beneficial as our intra-domain variance has been reduced to the minimum by repeatable data acquisition. Potentially, such an approach can take intra-domain variance emerging from grain morphology differences due to different imaging locations on the rolled sheet cross-sections into account. Recently, Yu et al. published an improvement of AdaptSegNet<sup>7</sup> including an attention mechanism in order to focus domain adaptation on the parts of the images that are the most difficult to transfer from the source to the target. While the work at hand has focused on adversarial UDA techniques only, promising style transfer GAN approaches are also good candidates for UDA methods and are currently an active research topic<sup>33</sup>.

Whether to employ UDA for training a model depends on three criteria. First, the effort associated with labeling source and target domain data needs to be considered since UDA avoids this cost for the target. For instance, UDA is especially favorable when synthetic source data (e.g., simulation data) with inherent labels and expensive target annotations are concerned. Second, the features contained in the source and target input data determine the attainable annotation accuracy. A setting where significantly more precise labels can be obtained in a source domain compared to the target proves beneficial to UDA. Thus, assuming comparatively poor-quality target labels, the performance gap between a UDA training and the direct supervised training on the target domain diminishes. The transfer from SEM to LOM provides a good example as SEM image features not only render the phase annotation process easier but also the bainite sub-class differentiation possible in the first place. Therefore, considering that a SEM acquires data substantially slower and is not affordable for every research laboratory, training a UDA model with external annotated SEM data to transfer to LOM can increase accessibility to high-quality models and potentially even enable specific tasks. Lastly, the domain gap to bridge has to remain maintainable. Source and target domains need to share enough common image features for the model to learn a domain-independent representation of the data. While this work gives first insights into the UDA scope of application, its precise limits still need to be explored.

The variety of materials and characterization methodologies utilized in materials science and engineering is boundless. This along with the increasing degree of automation and image acquisition rates, will very likely cause a drastic discrepancy between labeled and unlabeled data quantities in the future. Presumably, this will further raise the interest for unsupervised-learning methods such as UDA since these hold the potential to

alleviate the demand for expensive annotations. Given the virtual materials design acceleration we currently undergo through the digital transformation, unsupervised methods will be indispensable to foster the efficient experimental confirmation of computationally optimized microstructures.

## METHODS

### Specimen fabrication and image acquisition methodology

This work is based on SEM images of low-carbon complex-phase steel which was the subject of our last work<sup>5</sup>. Samples were obtained from thermomechanically rolled heavy industrial plates. Their microstructure encompasses lath-shaped bainite surrounded by polygonal and irregular ferrite with dispersed granular carbon-rich 2nd phase and MA islands. Micrographs were taken in the plate’s transversal direction (TD), between quarter- and mid-thickness of the plate. Rolling-induced stress and cooling rate gradients result in a small microstructure variance, wherein some images taken from comparatively surface-near regions, the polygonal ferrite grains are elongated in the rolling direction (RD). During imaging, segregation zones in the plate core were avoided. The specimens were ground using 80–1200 grid SiC papers, and then subjected to polishing with 6, 3, and finally, 1  $\mu\text{m}$  diamond grain sizes. Using different etching and imaging conditions as well as image modalities, four image sets were drawn from these specimens. The configurations are presented in Table 6.

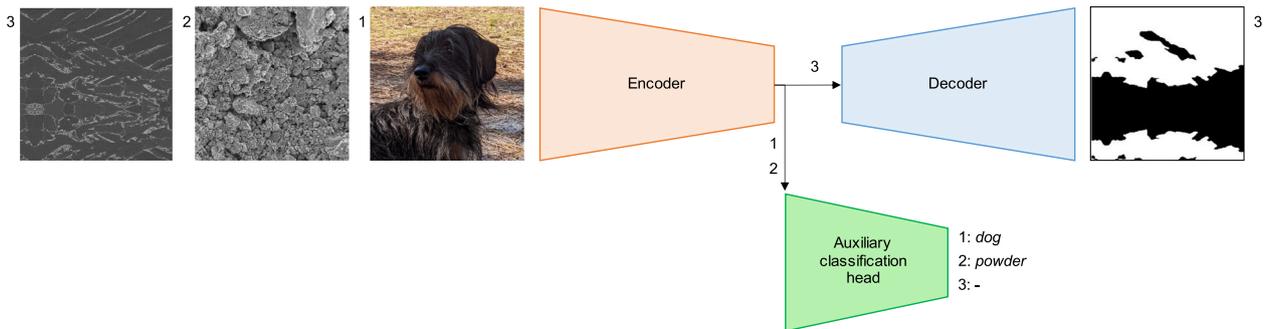
The etching duration was controlled by a metallographer waiting for a macroscopic contrast to be visible to the naked eye. The etching reveals grain boundaries since the reaction kinetics depend on the local chemical composition and crystallographic orientation. Therefore, carbide films and a few MA constituents are exposed. An example input image, the same with the superimposed label, and a detailed view of each data set (i.e., domain) is presented in Fig. 1.

In terms of imaging conditions, the SEM setup 1 was utilizing a Zeiss Merlin FEG-SEM using secondary electron contrast (in-lens) at a magnification of  $\times 2000$  with an image size of  $2048 \times 1433$  (annotation bar cropped), which represents  $56.7 \times 42.5 \mu\text{m}^2$  (pixel size = 27.7 nm). The SEM was set at an acceleration voltage of 5 kV, a probe current of 300 pA, and a working distance of 5 mm. Small acceleration voltages reduce the interaction volume and increase surface sensitivity. In SEM setup 2, the micrographs were recorded in a Zeiss Supra FEG-SEM with another in-lens detector. The acceleration voltage, the probe current and the working distance were the same as in the first setup. The magnification was set at  $\times 1000$  with a higher image resolution, giving virtually identical physical pixel size compared to the first setup. Due to subsequent stitching and cropping, images of SEM setup 2 ultimately had the same size as others. Contrast and brightness settings differed from the first imaging setup, see Fig. 1. Within each domain, micrographs were acquired with the same image contrast and brightness settings in the SEM. Lastly, the LOM images were recorded in an Olympus LEXT OLS 4100. Micrographs were taken at a magnification of  $\times 1000$  with an image resolution of  $1024 \times 1024$  pixels, corresponding to an area of  $129.6 \times 129.6 \mu\text{m}^2$  (pixel size = 126.6 nm). All LOM images were acquired with the same exposure settings.

Segmentation masks were drawn manually by human experts on a digital drawing tablet (Wacom). Correlative EBSD maps on the source domain micrographs rendered the annotation more reproducible and accurate. For more details on the acquisition, multi-modal registration and annotation process we refer to ref. 5. Note that for the target domains, with the exception of T3, annotations were only available for a small portion, i.e. the test images.

### Dataset processing

The LOM images were assimilated to the native SEM datasets in terms of physical pixel size and field of view by scaling and cropping operations (cf. Fig. 1a and j). Subsequent processing was identical for all datasets. Images were mirror-padded to make them square so four tiles can be extracted per image. Following the results of our last paper<sup>5</sup>, the images were downscaled by a factor  $\times 0.5$  before tiling, effectively increasing the context passed to the models’ receptive field<sup>34</sup>. The obtained training tiles are then of size  $636 \times 636$  ( $512 \times 512$  plus an additional 62-pixel overlap extending into the adjacent tiles at each tile border). This overlap-tile strategy was introduced by Ronneberger et al.<sup>35</sup> to cope with memory restrictions and helps the model to circumvent tile border effects. For segmentation loss computation, the overlap regions extending in neighboring tiles were discarded (i.e.,  $512 \times 512$  center region was used solely). To align the



**Fig. 6 Summary of the pre-training and fine-tuning procedure.** Starting from ImageNet pre-trained weights (1), we optionally use the NanoSEM dataset<sup>37</sup> for pre-training the encoder of the model on this classification task (2). The weights of the encoder are then directly transferred for the final fine-tuning on the source domain (3), while the decoder is random-initialized.

domains' phase fractions as much as possible, few training tiles of the target domains were discarded, ultimately resulting in the tile numbers and phase fractions listed in Table 6. The source dataset was split in five folds for cross-validation. In contrast, the annotated portion in the target domains was too small to perform cross-validation. Therefore, only a single train and test set has been built for the target datasets. While training was performed with tiles, the model evaluation was conducted on full images. Bright-field LOM images (T3) pixel values were inverted to give them a dark background similar to SEM.

### Pre-training datasets

The first pre-training dataset used in this study is ImageNet<sup>36</sup>. Further, the apparent domain gap between ImageNet and our target datasets motivated us to test an additional pre-training dataset with a smaller domain gap. Therefore, we selected a SEM dataset of nanoscientific objects<sup>37</sup>, which comprises ~22k images, non-uniformly distributed in 10 classes: biological, fibers, films, and coated surfaces, microelectromechanical systems (MEMS) and electrodes, nanowires, particles, porous sponges, pattern surfaces, powders, and tips. In the following, we refer to this pre-training dataset with "NanoSEM". Before using the NanoSEM dataset, a pre-processing cleaning step was performed. Images with many burned-in measurement annotations were discarded and SEM annotation bars were cropped to avoid spurious correlations between annotation bars and class predictions, known to occur otherwise. Finally, the pre-training dataset amounted to 18,750 images. More details about the pre-training methodology and results are supplied in the section "Pre-training and fine-tuning procedure".

### Segmentation architectures

As part of this work, two main segmentation architectures are implemented. The well-established U-Net<sup>35</sup> is used in the first place to investigate different pre-training strategies. This fully convolutional architecture is, as its name implies, composed of an encoder-decoder structure with skip connections between the corresponding levels of the encoder and decoder. It gave outstanding results on medical segmentation tasks even with very little training data.

Among many different segmentation models that were proposed after the U-Net, one series of models marked a turning point in this field. Chen et al. published the first version under the name DeepLab<sup>38</sup>. In this work, the second DeepLab version is implemented (DeepLabv2). This architecture uses so-called dilated convolutions (or atrous convolutions), which help the model to enlarge its field of view (receptive field) and take patterns at larger scales into account appropriately. The main idea of DeepLabv2 is to learn and aggregate patterns at different scales with dilated convolutions having different dilation rates. This aggregation of dilated convolutions effectively causes a more uniform distribution in the effective receptive field<sup>34</sup>.

The encoder used for the U-Net is a portion of the VGG16 classification network<sup>39</sup>, while DeepLabv2 was built with a ResNet-101<sup>40</sup> encoder. The exact architecture for both cases are given in Supplementary Fig. 1 and in ref. 7, respectively.

For segmentation training, a binary cross-entropy loss and an Adam optimizer was employed. Learning rates, batch sizes and training times vary along this study. Thus, these parameters are specified in the section "Results" directly for the different experiments. Reported learning rates

were selected by grid-search optimization. The models were trained on a GPU cluster node consisting of four parallel NVIDIA Ampere A100 GPUs.

### Pre-training and fine-tuning procedure

For the ImageNet pre-training, we used pre-trained weights provided by the python package "Segmentation models pytorch"<sup>41</sup>.

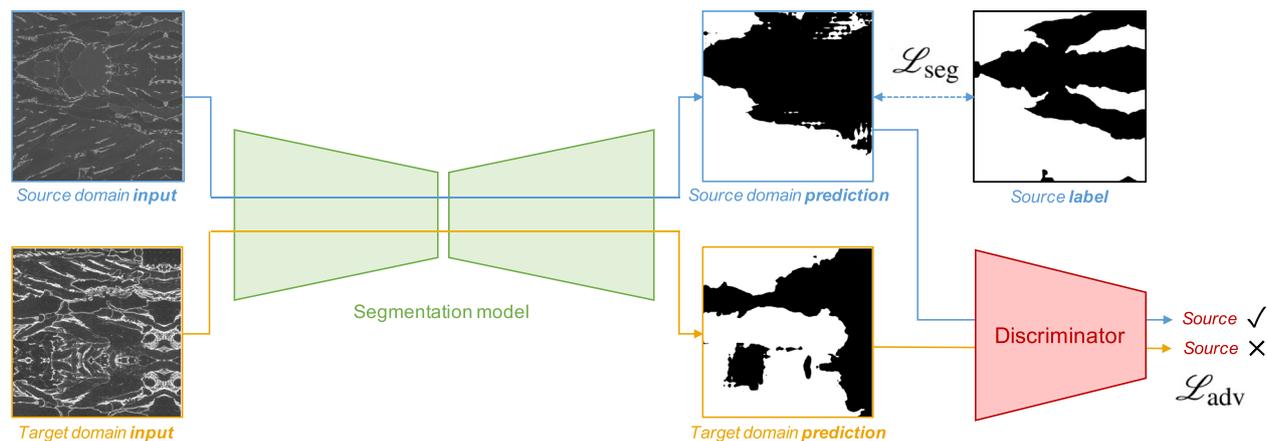
Concerning our self-performed NanoSEM pre-training, we passed the U-Net encoder output to an auxiliary classification head<sup>41</sup>. This head consists of a global average pooling layer, followed by 50% dropout and a linear layer with a sigmoid activation. The auxiliary classification head facilitates encoder training on classification datasets. ImageNet weights were used as an initialization for the trainings, making this process a two-step pre-training (from ImageNet to NanoSEM to the final task). The 18,750 NanoSEM images were split into 15k images for training and 3750 images for testing purposes. Pre-training used an Adam optimizer with a constant and encoder-layer independent learning rate and was run for 100 epochs. The obtained pre-trained models were transferred to the segmentation task by just copying the weights of the model performing best on the pre-training classification task. The full model was then fine-tuned on the source dataset (without frozen layers) with a reduced learning rate for the pre-trained encoder (10x lower than the decoder learning rate). The two-stage pre-training process along with fine-tuning is summarized in Fig. 6. Please note that the individual learning rates applied at the pre-training and fine-tuning stage are of major importance. An optimization of the learning rate used for pre-training on NanoSEM has been carried out (details are given in Supplementary note 1).

In case of the VGG16 U-Net model, aside from random initialization, either the two-stage pre-training or ImageNet pre-trained weights were used as initial conditions before pursuing fine-tuning to the source domain. In contrast, for the DeepLabv2 models solely ImageNet pre-trained weights were used. Testing these models on target datasets (T1–T3) allowed to investigate the impact of pre-training on the domain generalization (DG) capability of the models.

### Unsupervised domain adaptation—an adversarial framework

To additionally take advantage of unlabeled target domain data, which often can be created in abundance effortlessly, the UDA method of Tsai et al.<sup>7</sup> has been implemented. This adversarial framework can be used to train a semantic segmentation unsupervised domain adaptation task. It is based on the original idea proposed by Ganin et al.<sup>20</sup> The code of<sup>7</sup> was adapted to make it compatible with our data. Figure 7 depicts the training process in a simplified fashion.

As proposed in the original framework<sup>7</sup>, we utilize a DeepLabv2 with a ResNet-101 as the segmentation architecture, which facilitates comparability with the corresponding domain generalization experiments described at the end of the section "Pre-training and fine-tuning procedure". In the UDA framework, annotated source and unannotated target domain data are fed into this segmentation model (shared weights). The source domain prediction is used for training the segmentation model in a supervised manner, given that labels are available in this domain. This gives a first part of the loss function ( $\mathcal{L}_{seg}$ ), evaluated as a binary cross-entropy. Furthermore, source and target domain predictions are passed to a discriminator model, which attempts to classify from which domain the prediction comes. The second part of the loss, the so-called adversarial loss ( $\mathcal{L}_{adv}$ ), quantifies the ability of the segmentation model to fool the



**Fig. 7** Simplified schematic representation of AdaptSegNet<sup>7</sup>. Tile examples from the source and target 2 domains are given on the left. The source prediction is used in a supervised manner with cross-entropy loss. Both domains' prediction are fed into the discriminator.

discriminator. It is also computed as a binary cross-entropy for the domain classification. Additional to the segmentation outputs, network-internal feature representations of both domains are extracted from an auxiliary segmentation head, reshaped to segmentation mask size (auxiliary segmentation), and passed to the discriminator ( $\mathcal{L}_{adv}^{aux}$ ). This is not represented in Fig. 7 for the sake of simplicity. Moreover, the source domain auxiliary segmentation is compared to the annotation mask ( $\mathcal{L}_{seg}^{aux}$ ). The different loss parts are weighted so emphasis can be put on either of the segmentation or adversarial losses introducing three further hyper-parameters. Note that the loss portions related to the auxiliary feature output are typically less weighted, making Fig. 7 representation a good first approximation of the model. The four aforementioned loss parts ( $\mathcal{L}_{seg}$ ,  $\mathcal{L}_{adv}$ ,  $\mathcal{L}_{adv}^{aux}$ ,  $\mathcal{L}_{seg}^{aux}$ ) compose the training loss of the segmentation model, whereas the discriminator is optimized based on a domain-classification cross-entropy loss. When back-propagating the combined loss of the segmentation model, the weights of the discriminator are temporarily frozen. Both the segmentation and discriminator models are trained in an end-to-end fashion. Complementary technical details are given in the Supplementary note 2.

### Data augmentation

Data augmentation is a common practice in ML for increasing the labeled data amount without additional annotation cost. It consists in applying transformations to the data before passing it to the model. The main objective is to render the network invariant to specific transformations. As part of this work, a simple flip and 90° rotation pipeline (with probability 0.5) were first tested (marked as basic). Moreover, an extended pipeline making use of, amongst others, elastic transformation and optimized for the source domain in our last publication<sup>5</sup> was implemented (marked as extended). Both pipelines have been built with the Albumentations package<sup>42</sup>. Full details about the pipelines are given in Supplementary Table 2. Data augmentation was applied to train all our models except for the NanoSEM pre-training. In the UDA framework, both source and target datasets were augmented using the optimized pipeline.

### Evaluation metrics

For the segmentation models evaluation, we use an intersection over union (IoU) metric (see Eq. (1)), averaged over background and foreground classes (mIoU). In order to evaluate the trained models' generalizability, we used a relative mIoU deviation between the source domain performance with sole supervised learning and the target domain performance with the concerned generalization method (either domain generalization or UDA). We refer to this metric as relative domain transferability (RDT; see Eq. (2)). Using a relative deviation avoids overrating models that are performing better on the target domains because of their inherent advantage on the source domain. For instance, a hypothetical model with 70% and 65% mIoU on the source and target (RDT = -0.07) generalizes better than one yielding 80% and 70% mIoU (RDT = -0.13), despite latter's better target

performance.

$$\text{IoU} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}} \quad (1)$$

$$\text{RDT} = \frac{\text{mIoU}_T - \text{mIoU}_S}{\text{mIoU}_S} \quad (2)$$

TP, TN, FP, and FN represent true positive, true negative, false positive, and false negative pixels, respectively.  $\text{mIoU}_T$  and  $\text{mIoU}_S$  are the class-averaged model performance on the target domain and reference source domain, respectively.

### DATA AVAILABILITY

The datasets generated during and/or analyzed during the current study are not publicly available because they are part of an ongoing study and are subject to third-party (AG der Dillinger Hüttenwerke) restrictions.

### CODE AVAILABILITY

The codes used in this study are not publicly available because they are part of an ongoing study.

Received: 28 July 2021; Accepted: 5 January 2022;  
Published online: 03 February 2022

### REFERENCES

- Torrabá, A. & Efros, A. A. Unbiased look at dataset bias. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 1521–1528 (IEEE, Colorado Springs, CO, USA, 2011).
- Welinder, P., Welling, M. & Perona, P. A lazy man's approach to benchmarking: Semisupervised classifier evaluation and recalibration. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 3262–3269 (IEEE, Portland, OR, USA, 2013).
- Beery, S., Van Horn, G. & Perona, P. Recognition in terra incognita (eds Ferrari, V., Hebert, M., Sminchisescu, C. & Weiss Y.) In *Proc. European Conference on Computer Vision*, 456–473 (Springer, Cham, Munich, Germany, 2018).
- Thomas, A., Durmaz, A. R., Straub, T. & Eberl, C. Automated quantitative analyses of fatigue-induced surface damage by deep learning. *Materials* **13**, 3298 (2020).
- Durmaz, A. et al. A deep learning approach for complex microstructure inference. *Nat. Commun.* **12**, 6272 (2021).
- Fielding, L. The bainite controversy. *Mater. Sci. Technol.* **29**, 383–399 (2013).
- Tsai, Y.-H. et al. Learning to adapt structured output space for semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 7472–7481 (IEEE, Salt Lake City, USA, 2018).
- Azimi, S. M., Britz, D., Engstler, M., Fritz, M. & Mücklich, F. Advanced steel microstructural classification by deep learning methods. *Sci. Rep.* **8**, 2128 (2018).
- DeCost, B. L., Lei, B., Francis, T. & Holm, E. A. High throughput quantitative metallography for complex microstructures using deep learning: a case study in ultrahigh carbon steel. *Microsc. Microanal.* **25**, 21–29 (2019).

10. Yosinski, J., Clune, J., Bengio, Y. & Lipson, H. How transferable are features in deep neural networks? (eds Z. Ghahramani et al.) In *Proc. 27th International Conference on Neural Information Processing Systems*, Vol. 2. 3320–3328 (MIT Press, Cambridge, MA, USA, 2014).
11. He, K., Girshick, R. B. & Dollár, P. Rethinking imagenet pre-training. In *Proc. IEEE International Conference on Computer Vision*, 4917–4926 (IEEE, Seoul, South Korea, 2019).
12. Lin, T.-Y. et al. Microsoft coco: Common objects in context. (eds Fleet, D., Pajdla, T., Schiele, B. & Tuytelaars, T.) In *Proc. European Conference on Computer Vision*, 740–755 (Springer, Cham, Zurich, Switzerland, 2014).
13. Tajbakhsh, N. et al. Convolutional neural networks for medical image analysis: full training or fine tuning? *IEEE Trans. Med. Imaging* **35**, 1299–1312 (2016).
14. Cheplygina, V. Cats or cat scans: Transfer learning from natural or medical image source data sets? *Curr. Opin. Biomed. Eng.* **9**, 21–27 (2019).
15. Romero, M., Interian, Y., Solberg, T. & Valdes, G. Targeted transfer learning to improve performance in small medical physics datasets. *Med. Phys.* **47**, 6246–6256 (2020).
16. Gonthier, N., Gousseau, Y. & Ladjal, S. An analysis of the transfer learning of convolutional neural networks for artistic images. (eds Del Bimbo A. et al.) In *Pattern Recognition. ICPRI International Workshops and Challenges*, 546–561 (Springer, Cham, 2021).
17. Richter, S. R., Vineet, V., Roth, S. & Koltun, V. Playing for data: ground truth from computer games. (eds Leibe, B., Matas, J., Sebe, N. & Welling, M.) In *Proc. European Conference on Computer Vision*, 102–118 (Springer, Amsterdam, The Netherlands, 2016).
18. Ros, G., Sellart, L., Materzynska, J., Vazquez, D. & Lopez, A. M. The synthia dataset: a large collection of synthetic images for semantic segmentation of urban scenes. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 3234–3243 (IEEE, Las Vegas, NV, USA, 2016).
19. Cordts, M. et al. The cityscapes dataset for semantic urban scene understanding. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 3213–3223 (IEEE, Las Vegas, NV, USA, 2016).
20. Ganin, Y. & Lempitsky, V. Unsupervised domain adaptation by backpropagation. (eds Bach, F. & Blei, D.) In *Proc. 32nd International Conference on Machine Learning*, 1180–1189 (ICML, Lille, France, 2015).
21. Dou, Q., Ouyang, C., Chen, C., Chen, H. & Heng, P. -A. Unsupervised cross-modality domain adaptation of convnets for biomedical image segmentations with adversarial loss. (ed Lang, J.) In *Proc. 27th International Joint Conference on Artificial Intelligence*, 691–697 (AAAI Press, 2018).
22. Zhang, Y. et al. Collaborative unsupervised domain adaptation for medical image diagnosis. In *Proc. IEEE Transactions on Image Processing* (IEEE, 2020).
23. Aitken, A. et al. Checkerboard artifact free sub-pixel convolution: a note on sub-pixel convolution, resize convolution and convolution resize. Preprint at [bioRxiv https://arxiv.org/abs/1707.02937](https://arxiv.org/abs/1707.02937) (2017).
24. Iren, D. et al. Aachen-Heerlen annotated steel microstructure dataset. *Sci. Data* **8**, 1–9 (2021).
25. Selvaraju, R. R. et al. Grad-cam: Visual explanations from deep networks via gradient-based localization. (ed Ikeuchi, K.) In *Proc. IEEE International Conference on Computer Vision*, 618–626 (IEEE, Venice, Italy, 2017).
26. Thornton, K. & Poulsen, H. F. Three-dimensional materials science: an intersection of three-dimensional reconstructions and simulations. *MRS Bull.* **33**, 587–595 (2008).
27. Bolte, J.-A. et al. Unsupervised domain adaptation to improve image segmentation quality both in the source and target domain. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 1404–1413 (IEEE, Long Beach, CA, USA, 2019).
28. Pan, F., Shin, I., Rameau, F., Lee, S. & Kweon, I. S. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 3763–3772 (IEEE, Seattle, WA, USA, 2020).
29. Tranheden, W., Olsson, V., Pinto, J. & Svensson, L. Dacs: Domain adaptation via cross-domain mixed sampling. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 1378–1388 (IEEE, Waikoloa, HI, USA, 2021).
30. Vu, T.-H., Jain, H., Bucher, M., Cord, M. & Pérez, P. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2512–2521 (IEEE, Long Beach, California, USA, 2019).
31. Yang, Y. & Soatto, S. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 4084–4094 (IEEE, Snowmass Village, CO, USA, 2020).
32. Yu, F. et al. Dast: Unsupervised domain adaptation in semantic segmentation based on discriminator attention and selftraining. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10754–10762 (AAAI Press, Palo Alto, CA, USA, 2021).
33. Zhu, J.-Y., Park, T., Isola, P. & Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. (ed Ikeuchi, K.) In *Proc. IEEE International Conference on Computer Vision*, 2242–2251 (IEEE, Venice, Italy, 2017).
34. Luo, W., Li, Y., Urtasun, R. & Zemel, R. Understanding the effective receptive field in deep convolutional neural networks. (eds Lee, D., Sugiyama, M., Luxburg, U., Guyon, I. & Garnett, R.) In *Proc. 30th International Conference on Neural Information Processing Systems*, 4905–4913 (Curran Associates Inc., Red Hook, NY, USA, 2016).
35. Ronneberger, O., Fischer, P. & Brox, T. U-net: convolutional networks for biomedical image segmentation. *Lect. Notes Computer Sci.* **9351**, 234–241 (2015).
36. Deng, J. et al. Imagenet: A large-scale hierarchical image database. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 248–255 (IEEE, Miami, FL, USA, 2009).
37. Aversa, R., Modarres, M. H., Cozzini, S., Ciancio, R. & Chiusole, A. The first annotated set of scanning electron microscopy images for nanoscience. *Sci. Data* **5**, 1–10 (2018).
38. Chen, L., Papandreou, G., Kokkinos, I., Murphy, K. & Yuille, A. L. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**, 834–848 (2018).
39. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. Preprint at <https://arxiv.org/abs/1409.1556> (2015).
40. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 770–778 (IEEE, Las Vegas, NV, USA, 2016).
41. Yakubovskiy, P. Segmentation models pytorch. [https://github.com/qubvel/segmentation\\_models.pytorch](https://github.com/qubvel/segmentation_models.pytorch) (2020).
42. Buslaev, A. et al. Albumentations: fast and flexible image augmentations. *Information* **11**, 125 (2020).

## ACKNOWLEDGEMENTS

We express our appreciation towards project DEAL for providing Open Access funding.

## AUTHOR CONTRIBUTIONS

Conceptualization: A.D.; Data curation: A.D., A.G., M.M.; Formal Analysis: A.G.; Investigation: A.D., A.G. Methodology: A.D., A.G., A.T.; Project administration: A.D., P.K.; Resources: A.D., C.E., D.B.; Software: A.G., A.T.; Supervision: A.D., C.E., P.K.; Visualization: A.D., A.G.; Writing—original draft: A.D., A.G., C.E., M.M., P.K.; Writing—review & editing: all authors.

## FUNDING

Open Access funding enabled and organized by Projekt DEAL.

## COMPETING INTERESTS

The authors declare no competing interests.

## ADDITIONAL INFORMATION

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41524-022-00703-z>.

**Correspondence** and requests for materials should be addressed to Aurèle Goetz or Ali Riza Durmaz.

**Reprints and permission information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022