

This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

## Knowing What to Label for Few Shot Microscopy Image Cell Segmentation

Youssef Dawoud <sup>1\*</sup>

Arij Bouazizi <sup>1,2</sup> Vasileios Belagiannis <sup>5</sup> Gustavo Carneiro<sup>4</sup>

<sup>1</sup> Universität Ulm, Ulm, Germany
 <sup>2</sup> Mercedes-Benz AG, Stuttgart, Germany
 3 Ulm University Medical Center, Ulm, Germany
 4 Centre for Vision, Speech and Signal Processing, University of Surrey, United Kingdom
 5 Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

youssef.dawoud@uni-ulm.de

#### Abstract

In microscopy image cell segmentation, it is common to train a deep neural network on source data, containing different types of microscopy images, and then fine-tune it using a support set comprising a few randomly selected and annotated training target images. In this paper, we argue that the random selection of unlabelled training target images to be annotated and included in the support set may not enable an effective fine-tuning process, so we propose a new approach to optimise this image selection process. Our approach involves a new scoring function to find informative unlabelled target images. In particular, we propose to measure the consistency in the model predictions on target images against specific data augmentations. However, we observe that the model trained with source datasets does not reliably evaluate consistency on target images. To alleviate this problem, we propose novel self-supervised pretext tasks to compute the scores of unlabelled target images. Finally, the top few images with the least consistency scores are added to the support set for oracle (i.e., expert) annotation and later used to fine-tune the model to the target images. In our evaluations that involve the segmentation of five different types of cell images, we demonstrate promising results on several target test sets compared to the random selection approach as well as other selection approaches, such as Shannon's entropy and Monte-Carlo dropout.

## 1. Introduction

Microscopy image cell segmentation is one of the main fields in the area of medical image with the focus of studying the morphological properties of biological cells, i.e. geometrical shape and size along with other tasks such as cell detection [43], segmentation [5] and counting [1, 10]. Over the past years, research efforts have been devoted to automate microscopy image cell analysis, initially, with the support of classical image processing algorithms [42]. Subsequently, deep neural networks (DNNs), specifically, encoder-decoder architectures have dramatically evolved to become the state-of-the art automation approach in several microscopy image cell tasks, including microscopy image cell segmentation [5].

Earlier studies employed DNNs to learn a fully supervised cell segmentation model, which required the collection and pixel-level labelling of a large amount of microscopy image data to enable a robust training [41]. Recently, a more practical study [9] showed a cellsegmentation method that can be trained with a support set containing a few annotated microscopy training images, this method is known as few-shot microscopy image cell segmentation. In this setup, a deep neural network model is trained using source data, containing training images from various types of cell segmentation problems. Afterwards, the trained model is fine-tuned to the target images with cells of interest, using a support set containing few randomly selected and annotated microscopy images. Even though effective, random image selection can be improved since the informativeness of the support set may be low, which may result in a poor fine-tuning process that also leads to a low-accuracy performance in the unseen testing target images.

In this paper, we propose a new approach to optimise the selection of samples to annotate and include in the support set, so that it contains highly informative training samples that help improve the classification accuracy of the finetuned model, compared with random selection. In return, we offer a more efficient use of expert's time for annota-



Figure 1. Visual Result. We visually compare our scoring function (Ours) to random selection using the FCRN architecture at  $|\mathcal{B}| = 3$ -shots. The red colour corresponds to false positive, the green colour to false negative, the black colour to true negative, and the white colour to true positive. Best viewed in colour.

tion. Our selection approach relies on a scoring function to select the support set from the unlabelled target set by evaluating the consistency in the predictions of the segmentation model. In particular, our scoring function calculates the pixel-wise cross-entropy loss between the segmentation model prediction using an image from the target set and the predictions of its augmented versions. However, we notice that the segmentation model trained using only the source data does not produce a reliable scoring function for the images belonging to the target set. Therefore, we propose to fine tune the segmentation model to the unlabelled target set using novel pretext tasks. Specifically, we propose to learn cell segmentation in images belonging to the target set using pseudo-binary segmentation labels, which we generate using classical image processing segmentation operations. The fine-tuned model, obtained from this pseudosegmentation learning, is used to calculate the cross entropy scores of the target images. Then, the top few images with the least consistency scores are added to the support set for oracle (i.e., expert) annotation and later used to fine-tune the model to the target images. At last, we evaluate the finetuned model on the testing target images. We summarise our contribution as follows 1) We propose a novel pretext task of cell segmentation learning to fine-tune the segmentation model which we use afterwards for selecting samples to be annotated in a few-shot learning problem 2) We present a new scoring function for support set selection from the unlabelled target data that measures the performance consistency for the pretext tasks as a function of specific data augmentations 3) In our experiments we show promising results on five target sets involving different types of cell images compared to the random selection approach, in addition to other selection approaches, namely, Shannon's entropy and Monte-Carlo dropout. Our code and models are made publicly available. <sup>1</sup>

## 2. Related Work

## 2.1. Cell Segmentation

Early automatic microscopy cell segmentation methods have been developed with the aid of classical image processing and computer vision algorithms [11, 24, 40]. More recently, deep neural network architectures ranging from fully convolutional networks (FCN) to self-attention based have significantly evolved to become the state-of-the-art automation approach for several cell analysis tasks like nuclei segmentation [29], mitosis detection in histology images [6] and cardiac segmentation in MRI images [16]. Among FCN architectures, U-Net [36], was firstly developed for the task of neuronal structure segmentation in electron microscopy, but nowadays, it is applied in several types of medical image analysis problems. Another type of FCN architecture is the fully convolutional regression network (FCRN) [41], designed for cell counting and segmentation in microscopy images. In our work, we rely on FCRN architecture for automating microscopy cell segmentation.

<sup>&</sup>lt;sup>1</sup>https://github.com/Yussef93/KnowWhatToLabel/

#### 2.2. Few-Shot Segmentation

The availability of large annotated training sets enables a robust fully supervised training of models, but many realworld problems only contain small training sets, reducing the viability of supervised-learning. These problems are known as few-shot learning. To that end, several approaches have been developed to enable the learning of a generic model that can be adapted to different tasks using a limited amount of annotated training data [12, 31]. Only few studies exist on the problem of few-shot medical image segmentation [28, 33]. Among them is few-shot microscopy image cell segmentation [9] and organ segmentation [26], where a model is trained using source datasets and then fine-tuned to the target images using a support set containing a handful of randomly selected and annotated images. Nevertheless, the randomly selected images in the support set may not be informative for the fine-tuning process, which can result in poor training and low segmentation accuracy in the target microscopy testing images. Therefore, we argue that the selection of images to be included in the support set must be optimised, in terms of their information content, to achieve a better fine-tuning process that can result in a good cell segmentation performance in target microscopy testing images.

# 2.3. Self-Supervised Learning in Medical Image Analysis

Learning prior representations from the unlabelled data, i.e. self-supervision, has proven to be an effective approach when fine-tuned on subsequent target tasks such as classification and segmentation. Numerous approaches have been developed for learning prior representations in deep neural networks, e.g., contrastive learning [4], jigSaw puzzles [32], rotation prediction [18] and image reconstruction [2]. A wide variety of generic self-supervised tasks have been adapted to various medical image analysis applications, like jigSaw and rotation prediction for 3D computed tomography (CT) scans [39], contrastive learning for spinal MRI [21] and context restoration [3] for CT and ultrasound scan analysis. However, other studies have proposed target specific pretext tasks [8], for instance, superpixels prediction for abdominal CT-scans [33]. In our work, we follow a similar path as in [33], where we propose a new target specific pretext task. More specifically, we use conventional image segmentation operations to extract binary segmentation pseudo-labels for the microscopy images in the target dataset. Afterwards, we fine-tune our generic segmentation model using the pseudo-labelled target dataset to obtain a good prior representation for the target data, which we leverage to design our scoring function.

#### 2.4. Support Set Selection Approaches

The selection of data samples has been thoroughly studied under different learning paradigms, but more noticeably in active learning [37]. Most studies rely on scoring functions for selecting images to be annotated from a pool of unlabelled images, where data samples fulfilling the selection criterion are annotated and appended to the labelled set under the constraint of a limited annotation budget. Then, deep neural networks are trained on a target learning task at hand, e.g. classification or segmentation, using this labelled set. Several scoring functions have been proposed over the past years, such as Shannon's entropy [38], variation ratio [13], and Monte-Carlo (MC) dropout [14]. Leveraging these scoring functions for the problem of data sample selection has been well studied for natural image classification [15] and semantic segmentation [19]. Nevertheless, it has been under explored for the problem of support set selection for few-shot microscopy image cell segmentation. In fact, to the best of our knowledge, we are the first to discuss a solution to this problem based on a new automatic selection mechanism to construct the support set of target microscopy image dataset that leverages consistency loss with respect to data augmentation as our scoring function. Our approach outperforms other selection functions, such as random, Shannon's entropy and MC-dropout approaches.

#### 3. Method

In this section, we start by the problem definition followed by our support set selection method which consists of three steps. First, we propose our new self-supervised pretext task using the target dataset. Next, we utilise this selfsupervised trained model together with our scoring function to select the images to include and annotate in the support set. Finally, we fine-tune the self-supervised trained model using the support set and then evaluate on the testing target images.

#### 3.1. Problem Definition

Assume a collection of microscopy image datasets  $\widehat{S} = \{S_i\}_{i=1}^{|\widehat{S}|}$ . Each dataset is denoted by  $S_i = \{(\mathbf{x}, \mathbf{y})_l\}_{l=1}^{|S_i|}$ , where  $(\mathbf{x}, \mathbf{y})_l$  is a pair of microscopy image  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{H \times W}$  ( $H \times W$  is the image size) and pixel-level binary cell segmentation ground-truth  $\mathbf{y} \in \mathcal{Y} \subset \{0, 1\}^{H \times W}$ . All datasets in  $\widehat{S}$  are referred to as the sources. Note that each dataset represents a different microscopy image domain with different image appearance and cell segmentation task. We rely on the source datasets to learn a generic binary cell segmentation function  $f_{\theta} : \mathcal{X} \to [0, 1]^{H \times W}$  approximated by a deep neural network with parameters  $\theta \in \Theta$ .

We also have the target dataset  $\mathcal{T} = \{(\mathbf{x})_j\}_{j=1}^{|\mathcal{T}|}$ , which contains images belonging to a different microscopy domain and a different cell segmentation task. Our main ob-

jective is to train a model to segment cells in microscopy images from the target set using a support set containing annotated images denoted by  $\mathcal{T}^{(|\mathcal{B}|)}$ . The model is trained under the constraint of a limited annotation budget  $|\mathcal{B}|$ , i.e. fewshot training. Towards our objective, we consider a pool of unlabelled training images  $\mathcal{T}^{(pool)} = \{(\mathbf{x})_n\}_{n=1}^{|\mathcal{T}^{(pool)}|}$ , where  $\mathcal{T}^{(|\mathcal{B}|)} \subset \mathcal{T}^{(pool)} \subset \mathcal{T}$ , with the target testing set formed by  $\hat{\mathcal{T}} = \mathcal{T} \setminus \mathcal{T}^{(pool)}$ . The samples in the support set are selected based on the pixel-wise binary cross entropy (BCE) scoring function  $s : \mathcal{X} \times \Theta \to \mathbb{R}$  for assessing images in  $\mathcal{T}^{(pool)}$ . In summary, only the images  $\mathbf{x} \in \mathcal{T}^{(pool)}$  that maximise the scoring function s(.) are inserted into  $\mathcal{T}^{(|\mathcal{B}|)}$  to be labelled by an oracle until the annotation budget is exhausted. Next we present our pretext task followed by our selection approach to construct  $\mathcal{T}^{(|\mathcal{B}|)}$ .

#### 3.2. Support Set Selection

As previously mentioned, we assume that images belonging to the target testing set come from a different microscopy image domain and cell segmentation task, as the ones present in the source training set. Under this assumption, we argue that it is ineffective to use the source training set alone to train the segmentation model  $f_{\theta}(.)$  to be used in the scoring function s(.). Alternatively, we seek to adapt the entire model, i.e. the encoder and decoder, to the task of cell segmentation learning using a pseudo-labelled from the target data. In order to achieve this, we propose to employ classical computer vision operators to extract binary cell segmentation pseudo-labels for all the images in the target testing set. Specifically, we exploit image thresholding, global contrast enhancement using histogram equalisation [35], and dilation [20] filters to create a pseudo-labelled target dataset  $\mathcal{T}^{(PL)} = \{(\mathbf{x}, \mathbf{y}^{(PL)})_j\}_{j=1}^{|\mathcal{T}^{(PL)}|}$  where  $\mathbf{y}_j^{(PL)} \in \mathcal{Y}$  is the corresponding pseudo-label of microscopy image  $\mathbf{x}_{i}$ . Note that we use all images of the target data  $\mathcal{T}$  to create  $\mathcal{T}^{(PL)}$ . To acquire  $\mathbf{y}^{(PL)}$ , first, an input image  $\mathbf{x}$ (from  $\mathcal{T}^{(PL)}$ ) is passed through a histogram equalisation filter to get  $\mathbf{x}^{(e)}$ . Afterwards, we use a threshold filter  $t_{thresh}(\mathbf{x}^{(e)}(\omega))$  defined as:

$$\mathbf{x}^{(\gamma)} = t_{thresh}(\mathbf{x}^{(e)}(\omega)) = \{ \begin{array}{cc} 1, & \mathbf{x}^{(e)}(\omega) \le \gamma \\ 0, & \mathbf{x}^{(e)}(\omega) > \gamma \end{array}, \forall \omega \in \Omega,$$
(1)

where  $\gamma \in \mathbb{R}$  is a threshold that approximately represents the pixel value of the cell of interest in the target data, and  $\omega \in \Omega$  is a pixel address in the image lattice of size  $H \times W$ . At last, we apply a dilation filter of size  $2 \times 2$  on the thresholded image  $\mathbf{x}^{(\gamma)}$  to get the pseudo-label  $\mathbf{y}^{(PL)}$ .

Our training process starts with a meta-learning process that minimises the average BCE loss on the source training sets, where the loss for each training set  $S_i \in \hat{S}$  is defined



Figure 2. We present an example of our binary cell segmentation pseudo-label generation pipeline using a microscopy image **x** from TNBC target dataset [30] in (a). First, we enhance the contrast of the image using histogram equalisation  $(\mathbf{x}^{(e)})$  in (b), followed by image thresholding  $(\mathbf{x}^{(\gamma)})$  in (c) and dilation filtering in (d) to form the pseudo-label  $\mathbf{y}^{(PL)}$ .

as:

$$\mathcal{L}_{BCE}(\theta, S_i) = -\frac{1}{|S_i|} \sum_{(\mathbf{x}, \mathbf{y}) \in S_i} \sum_{\omega \in \Omega} [\mathbf{w}(\omega) \mathbf{y}(\omega) \log(\hat{\mathbf{y}}(\omega)) + (1 - \mathbf{y}(\omega)) \log(1 - \hat{\mathbf{y}}(\omega))],$$
(2)

where  $\hat{\mathbf{y}}(\omega)$  denotes the prediction of the segmentation model (represented by  $\hat{\mathbf{y}} = f_{\theta}(\mathbf{x})$ ) at spatial position  $\omega$ , and  $\mathbf{w}(\omega)$  is a weighting factor defined by the ratio of foreground to background classes in the dataset. After this initial training with (2), we fine-tune the parameters  $\theta$  of our generic model  $f_{\theta}(.)$  by minimising  $\mathcal{L}_{BCE}(\theta, \mathcal{T}^{(PL)})$ . By using  $\mathcal{L}_{BCE}(\theta, \mathcal{T}^{(PL)})$ , we seek to adapt the parameters of the generic model using the following optimisation:

$$\theta' = \arg\min_{\alpha} [\mathcal{L}_{BCE}(\theta, \mathcal{T}^{(PL)})].$$
(3)

Once the model is trained on pseudo-segmentation learning, we fix its learned parameters  $\theta'$  and start our selection process. In Algorithm 1, we describe our complete support set selection approach. To begin with, each microscopy image  $\mathbf{x}_m \in \mathcal{T}^{(pool)}, m \in \{1, 2, \dots, |\mathcal{T}^{(pool)}|\}$  is augmented three times, each time with a different augmentation operation of magnitude  $\psi$ , namely, auto-contrast  $C(\mathbf{x})$ , brightness  $B(\mathbf{x}, \psi)$  and sharpness  $S(\mathbf{x}, \psi)$ . We show in our experiments that these specific augmentations are suitable to our selection method for the microscopy image domain. Then, we calculate the BCE score for every microscopy image  $\mathbf{x}$ in  $\mathcal{T}^{(pool)}$  between the prediction of the augmented images, i.e.  $\hat{\mathbf{y}}_A = f_{\theta'}(A(\mathbf{x})), A(.) \in \{C(\mathbf{x}), B(\mathbf{x}, \psi), S(\mathbf{x}, \psi)\}$ and the prediction of the non-augmented image  $\hat{\mathbf{y}} = f'_{\theta}(\mathbf{x})$ , as follows:

$$s(\mathbf{x}, \theta') = -\sum_{A(.) \in \{C(\mathbf{x}), B(\mathbf{x}, \psi), S(\mathbf{x}, \psi)\}} \frac{1}{H \times W} \sum_{\omega \in \Omega} [\mathbf{w}(\omega) \hat{\mathbf{y}}(\omega) \log(\hat{\mathbf{y}}_A(\omega)) + (1 - \hat{\mathbf{y}}(\omega)) \log(1 - \hat{\mathbf{y}}_A(\omega))].$$
(4)

The last step of the algorithm consists of forming the support set  $\mathcal{T}^{(|\mathcal{B}|)}$  using the following optimisation:

$$\mathcal{T}^{(|\mathcal{B}|)} = \arg \max_{\widehat{\mathcal{T}}^{(|\mathcal{B}|)} \subset \mathcal{T}^{(pool)}} \sum_{\mathbf{x} \in \widehat{\mathcal{T}}^{(|\mathcal{B}|)}} s(\mathbf{x}, \theta')$$

$$s.t. |\mathcal{T}^{(|\mathcal{B}|)}| = |\mathcal{B}|,$$
(5)

where  $|\mathcal{B}|$  is the size of the support set.

#### 3.3. Support Set Fine-tuning

After selecting the support set images, we request an oracle (i.e., human expert) to manually annotate them. Next, we adapt  $f_{\theta'}(\mathbf{x})$  using  $\mathcal{T}^{(|\mathcal{B}|)} \subset \mathcal{T}^{(pool)}$  and the binary cross entropy loss  $\mathcal{L}_{BCE}(\theta', \mathcal{T}^{(|\mathcal{B}|)})$ , as in:

$$\theta^* = \arg\min_{\theta'} [\mathcal{L}_{BCE}(\theta', \mathcal{T}^{(|\mathcal{B}|)})].$$
(6)

We evaluate the fine-tuned model with  $f_{\theta^*}(\mathbf{x})$  on the target test set  $\widehat{\mathcal{T}} = \mathcal{T} \setminus \mathcal{T}^{(pool)}$ .

Algorithm 1: Dynamic Support Set Selection based on Data Augmentation

- 1: Input: Segmentation model  $f_{\theta}(\mathbf{x})$  trained using  $\mathcal{L}_{BCE}(\theta, \widehat{S})$  from (2), pseudo-labelled target dataset  $\mathcal{T}^{(PL)} = \{(\mathbf{x}, \mathbf{y}^{(PL)})_j\}_{j=1}^{|\mathcal{T}|}$ , unlabelled target training set  $\mathcal{T}^{(pool)}$ , and annotation budget  $|\mathcal{B}|$
- 2: Estimate  $\theta'$  from (3) using  $\mathcal{T}^{(PL)}$
- 3: for  $m = 1, 2, ..., |\mathcal{T}^{(pool)}|$  do
- 4: Augment  $\mathbf{x}_m \in \mathcal{T}^{(pool)}$  with auto-contrast  $C(\mathbf{x}_m)$ , brightness  $B(\mathbf{x}_m, \psi)$  and sharpness  $S(\mathbf{x}_m, \psi)$
- 5: Get prediction  $\hat{\mathbf{y}}_m = f_{\theta'}(\mathbf{x}_m)$
- 6: Get predictions  $\hat{\mathbf{y}}_{C,m} = f_{\theta'}(C(\mathbf{x}_m)),$  $\hat{\mathbf{y}}_{B,m} = f_{\theta'}(B(\mathbf{x}_m, \psi)),$  $\hat{\mathbf{y}}_{S,m} = f_{\theta'}(S(\mathbf{x}_m, \psi))$
- 7: Calculate score  $s(\mathbf{x}_m, \theta')$  from (4)
- 8: end for
- 9: Select  $\mathcal{T}^{(|\mathcal{B}|)}$  from (5) using  $\{s(\mathbf{x}_m, \theta')\}_{m=1}^{|\mathcal{T}^{(pool)}|}$
- 10: Estimate  $\theta^*$  from (6) using  $\mathcal{T}^{(|\hat{\mathcal{B}}|)}$
- 11: Output:  $f_{\theta^*}(.)$  to be tested on  $\widehat{\mathcal{T}} = \mathcal{T} \setminus \mathcal{T}^{(pool)}$ .

## 4. Experiments

In our experiments, we use the same microscopy image cell segmentation benchmark of [9], which consists of five microscopy image datasets. More specifically, we have B5 and B39 datasets from the Broad Bioimage Benchmark Collection (BBBC) [23]. The former contains 1200 fluorescent synthetic stain cells images, while the latter contains 200 fluorescent synthetic stain cells. We also have Serial Section Transmission Electron Microscopy (ssTEM) [17] and Electron Microscopy (EM) [25] datasets, containing 165 and 20 electron microscopy images, respectively, of mitochondria cells. The final dataset is Triple Negative Breast Cancer (TNBC) that consists of 50 histology images of breast biopsy [30]. We compare our scoring function against Shannon's entropy, MC-dropout, and random selection. Furthermore, we define our evaluation protocol since, to the best of our knowledge, we are the first to address support set selection problem.

#### 4.1. Implementation Details

We employ the FCRN architecture [41] as our cell segmentation backbone model. First, we train the segmentation model  $f_{\theta}(\mathbf{x})$  using the source datasets. However, instead of training the segmentation model from scratch, we exploit the readily trained model using a gradient-based metalearning reptile algorithm provided by [9].

**Pseudo-label segmentation** We generate pseudo-labels using all target images as described in Sec. 3.2, where the threshold value  $\gamma$  is heuristically defined per target dataset by visually inspecting the pixel values of cells of interest and setting the threshold value accordingly. Also, we augment  $\mathcal{T}^{(PL)}$  by extracting image patches of size 256 × 256 from every microscopy image and its corresponding pseudo-label. We initialise the segmentation model with the parameter  $\theta$  learned using source data, and train it for 100 epochs using  $\mathcal{T}^{(PL)}$  and Adam optimiser [22] with learning rate 0.0001 and weight decay 0.0005.

Support set selection We experiment with support set sizes  $|\mathcal{B}| \in \{1, 3, 5, 7, 10\}$  shots. For this stage, we noticed a better performance when selecting image patches instead of full resolution images. Every image in  $\mathcal{T}^{(pool)}$  is then cropped to patches of size  $256 \times 256$  pixels. The number of patches per image depends on the size of the original image, the crop step size, and the crop window size. We report the number of image patches per image in Table 4.1. Accordingly, we replace the support sizes mentioned earlier with the corresponding number of image patches per full resolution image times the number of shots. During the selection stage, we fix the trained model parameters  $\theta'$ . For Shannon's entropy score, we do one forward pass for each image

Target	TNBC	EM	ssTEM	B5	B39
Patches/Image	100	400	500	100	100

Table 1. Number of image patches extracted per full resolution image for each target dataset.

patch and calculate the prediction uncertainty using Shannon's entropy formula [38]. As for the MC-dropout scoring approach, we observed that adding a dropout layer before the last convolutional layer yields the best results. Accordingly, we set the dropout probability to 0.5. We make ten forward passes for each image patch, average out the output probability distribution and then we calculate Shannon's entropy [14]. Note that we only use MC-dropout during the selection stage, i.e., we do not train or fine-tune using MCdropout following [27]. For the random baseline, we randomly select image patches from  $\mathcal{T}^{(pool)}$ , where each image has an equal probability of selection. In our approach, we rely on auto-contrast, brightness, and sharpness augmentations provided in [7]. Moreover, we empirically set the magnitude of the distortion i.e.  $\psi$  of brightness and sharpness augmentations to 1.3<sup>2</sup> for targets TNBC, EM, ssTEM and B39. As for B5, we increase the distortion magnitude to 1.6, otherwise, the added distortion would be insufficient and hence, the BCE loss values would be insignificant. The selection criterion is the same for Shannon's entropy, MCdropout, and our approach, i.e., images that correspond to the top  $|\mathcal{B}|$  loss scores are inserted into the support set.

**Fine-tuning** We fine-tune the segmentation model  $f_{\theta'}(\mathbf{x})$  using the resulting support set and Adam optimiser with learning rate 0.0001 and weight decay 0.0005 for 20 epochs. At last, we test the model using  $\hat{\mathcal{T}}$ . Note that we use the test images in full resolution.

#### **4.2. Evaluation Protocol**

Throughout our experiments, we rely on the mean intersection over union (mIoU) for quantifying the segmentation performance. We follow the same protocol as [9] by conducting leave-one-dataset-out-cross-validation to split the microscopy datasets to source and target sets. In particular, we select four datasets as sources and treat the remaining dataset as target. The target set is randomly split to  $\mathcal{T}^{(pool)}$ and  $\hat{\mathcal{T}}$ . We repeat the random split ten times and report the mean and standard deviation of the mIoU over ten runs.

### 4.3. Results and Discussion

Figure 3 shows a comparison of mIoU results averaged over all target datasets, where we observe a consistently better performance using our scoring function. Additionally, we report the numerical results of TNBC, EM and ssTEM



Figure 3. Comparison of mIoU results averaged over all datasets using our approach as well as Shannon's entropy, MC-dropout and random selection approaches.



Figure 4. Impact on mIoU results when  $|\mathcal{B}|$  extracting more image patches for 1-shot for target data ssTEM.

in Tab. 2. Visual results are shown in Fig. 1. For all support set sizes, we notice that our selection approach (Ours) performs better than the random selection baseline at target cell segmentation tasks of complex microscopy domain involving different structural membranes other than the cell of interest, namely, histology (TNBC) and electron microscopy (ssTEM and EM). As for B5 and B39 (see supplementary material) we notice a high performance for both our approach and random selection, since both datasets come from a less complex microscopy domain consisting of synthetic cells and background only. As for MC-dropout and entropy approaches, we observe that our scoring function yields an overall better and more consistent performance in all target datasets. In very few cases, due to increasing support set size entropy and MC-dropout perform slightly better. We highlight the significance of our approach using Wilcoxon test (See Tab. 1 in supplementary material). Next, we conduct a series of ablation studies to analyse our approach under different configurations in Sec. 4.4.

#### 4.4. Ablation Study

**Support set size.** We examine the effect of limiting the support set size to only 100 image patches per target image. This only impacts the target datasets ssTEM and EM, while the results of the TNBC, B5 and B39 datasets remain unchanged. As reported in Fig. 5, we observe a significant performance drop in all selection approaches due to a smaller support set size in the fine-tuning process. On

 $<sup>^{2}\</sup>psi > 1$  means higher distortion.

Target: TNBC							
Mathad	1-shot	3-shot	5-shot	7-shot	10-shot		
Wiethou	$( \mathcal{B}  = 100)$	$( \mathcal{B}  = 300)$	$( \mathcal{B}  = 500)$	$( \mathcal{B}  = 700)$	$( \mathcal{B}  = 1000)$		
Entropy	40.8% ±3.9	44.1% ±3.1	43.8% ±4.7	45.6% ±5.4	46.7% ±5.8		
MC-dropout	40.8% ±4.0	43.9% ±3.1	43.7% ±4.7	45.4% ±5.4	46.8% ±5.8		
Random	37.0% ±9.4	44.7% ±4.7	42.1% ±2.3	45.7% ±4.7	46.8% ±4.8		
Ours	47.1% ±3.5	47.8% ±5.8	47.7% ±5.0	48.0% ±5.2	49.2% ±4.9		
Target: EM							
Mathad	1-shot	3-shot	5-shot	7-shot	10-shot		
Wiethou	$( \mathcal{B}  = 400)$	$( \mathcal{B}  = 1200)$	$( \mathcal{B}  = 2000)$	$( \mathcal{B}  = 2800)$	$( \mathcal{B}  = 4000)$		
Entropy	61.0% ±2.9	67.1% ±2.7	68.6% ±2.3	70.1% ±1.9	72.9% ±1.6		
MC-dropout	61.7% ±2.6	63.4% ±2.4	66.1% ±2.1	68.0% ±1.8	70.0% ±2.4		
Random	58.9% ±5.0	65.2% ±2.9	68.6% ±3.4	71.2% ±2.9	72.2% ±3.2		
Ours	62.0% ±3.2	69.8% ±2.7	70.6% ±3.3	73.1% ±2.9	73.7% ±3.2		
Target: ssTEM							
Mathad	1-shot	3-shot	5-shot	7-shot	10-shot		
Wiethou	$( \mathcal{B}  = 500)$	$( \mathcal{B}  = 1500)$	$( \mathcal{B}  = 2500)$	$( \mathcal{B}  = 3200)$	$( \mathcal{B}  = 5000)$		
Entropy	49.5% ±3.6	62.6% ±2.6	65.4% ±3.7	67.7% ±2.7	68.6% ±3.6		
MC-dropout	49.3% ±3.0	62.2% ±3.1	64.6% ±3.5	67.9% ±3.6	68.2% ±4.5		
Random	47.0% ±4.8	60.6% ±4.0	63.1% ±2.8	66.4% ±4.1	67.1% ±4.1		
Ours	51.3% ±3.2	63.3% ±3.2	64.2% ±3.2	67.3% ±2.9	68.7% ±2.6		

Table 2. mIoU results for the target testing sets of TNBC, EM, ssTEM. Best results are highlighted.









10-shot

Pseudo-label Generation	1-shot	3-shot	5-shot	7-shot	10-shot
K-means	54.4% ±2.7	63.1% ±2.8	65.8% ±3.0	69.5% ±2.4	71.7% ±2.6
Equalisation+Threshold+Dilation	56.7% ±3.2	66.6% ±3.0	67.4% ±3.2	70.2% ±2.9	71.2% ±2.9

Table 3. Effect on mIoU results for using different pseudo-cell segmentation tasks averaged over two target datasets i.e. ssTEM and EM.

Target	TNBC	EM	ssTEM	B5	B39
$f_{\theta'}$	31±4.7	11±0.5	8±0.5	99±0	92±0.4

Table 4. Effect on mIoU results (%) for testing the pseudo-label segmentation model  $f_{\theta'}(\mathbf{x})$  on the target datasets without support set fine-tuning.

the other hand, our selection mechanism is still robust and yields better performance. This implies that our proposed scoring function remains more accurate compared to random, entropy and MC-dropout. We also report the results of using 200, 300, and 400 images per target image for dataset ssTEM in Fig. 4 which highlights the significant increase in performance as more image patches are extracted.

## Pseudo-label cell segmentation learning and fine-tuning.

We claimed that the segmentation model  $f_{\theta}(\mathbf{x})$  trained using source data only is not robust enough for scoring the target image predictions, which is the main motivation for the pretext task of pseudo-label cell segmentation learning. To support our claim, we conduct an experiment where we show the poorer performance of  $f_{\theta}(\mathbf{x})$  in comparison to the pseudo-label segmentation model  $f_{\theta'}(\mathbf{x})$ . The mIoU results in Fig. 6 averaged over TNBC, EM and ssTEM target data sets show that our pretext task generally improves the performance due to better scoring and selection of support set. Numerical results for each target data set could be viewed in supplementary material. Moreover, we motivate the importance of support set fine-tuning by evaluating the pseudo-label segmentation model on the target test sets. The results in Table 4 clearly show that relying on pseudolabel segmentation learning alone is insufficient for targets TNBC, EM, and ssTEM, so it is necessary to perform support set fine-tuning with expert annotation of these target microscopy images. However, B5 and B39 yield fairly accurate results since their microscopy domain comprises only cells of interest and background, therefore, our pseudolabel segmentation pipeline results in a useful pseudo labels that are close to expert level annotation. Finally, we study a different technique to generate the pseudo-labels  $(\mathbf{y}^{(PL)})$  for the task of pseudo-label cell segmentation learning. Namely, we use K-means [34] to cluster the pixels of each unlabelled target image to foreground and background classes i.e. binary segmentation map, hence, K = 2. The study is conducted using target datasets EM and ssTEM. It can be noticed that the cells of interest (mitochondria) for both target datasets attain dark pixel values relative to other cells/membranes in the same image (see Fig. 1 a), therefore, pixels belonging to cluster center of darker pixel value are assigned to the foreground class while the pixels belonging to cluster center of brighter pixel value are treated as background. We report the results in Table 3. We clearly notice a better performance using the model fine-tuned on

the pseudo-labels generated using the proposed pipeline in Sec. 3.2 compared to the ones generated using K-means. However, the results also show that the support set selection algorithm is robust to the underlying pseudo-label generation technique and it can still perform well using both models fine-tuned to pseudo-labels of both techniques.

**Data augmentation** Data augmentation may impact the consistency in the model predictions due to data distribution shifts. However, task specific augmentations may introduce a noticeable inconsistent predictions compared to random augmentations. In this work, we claim that pixel-level augmentations are more meaningful for the scoring function calculation than affine transformations. To this end, we compare our chosen augmentations with affine transformations, comprising image rotation (by 30°) and translation along the positive x (30% of H) and y axes (30% of W). We report our averaged mIoU results over TNBC, EM and ssTEM target data sets in Fig. 6 b. Clearly, the proposed augmentations of contrast, brightness and sharpness have a positive impact in the BCE calculation, providing a better support set selection. Numerical results are listed in supplementary material.

## 5. Conclusion

We presented an approach to optimise support set selection for an effective fine-tuning process, hence, a better performance in few-shot microscopy image cell segmentation. Throughout our experiments, we demonstrated that by relying on our novel pretext task of pseudo-label cell segmentation learning and our scoring function, consistent and overall better results are achieved outperforming Shannon's entropy, MC-dropout and more importantly, random selection. Moreover, a series of ablation studies highlighted the important factors of our approach, which are the support set size, the impact of pseudo-label segmentation learning on support set selection, fine-tuning, and the effect of data augmentation on the selection process. Our work can be extended by combining other selection techniques such as diversity-based selection. Also, we plan to combine semisupervised learning with support set fine-tuning.

#### Acknowledgments

This work was partially funded by Deutsche Forschungsgemeinschaft (DFG), Research Training Group GRK 2203: Micro- and nano-scale sensor technologies for the lung (PULMOSENS), by the German Federal Ministry for Economic Affairs and Energy within the project "KI Delta Learning" (Forderkennzeichen "19A19013A) and the Australian Research Council through grant FT190100525. G.C. acknowledges the support by the Alexander von Humboldt-Stiftung for the renewed research stay sponsorship.

## References

- Carlos Arteta, Victor Lempitsky, and Andrew Zisserman. Counting in the wild. In *European conference on computer vision*, pages 483–498. Springer, 2016.
- [2] Pierre Baldi. Autoencoders, unsupervised learning, and deep architectures. In *Proceedings of ICML workshop on unsuper*vised and transfer learning, pages 37–49. JMLR Workshop and Conference Proceedings, 2012.
- [3] Liang Chen, Paul Bentley, Kensaku Mori, Kazunari Misawa, Michitaka Fujiwara, and Daniel Rueckert. Self-supervised learning for medical image analysis using image context restoration. *Medical image analysis*, 58:101539, 2019.
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [5] Dan Ciresan, Alessandro Giusti, Luca Gambardella, and Jürgen Schmidhuber. Deep neural networks segment neuronal membranes in electron microscopy images. *Advances in neural information processing systems*, 25:2843–2851, 2012.
- [6] Dan C Cireşan, Alessandro Giusti, Luca M Gambardella, and Jürgen Schmidhuber. Mitosis detection in breast cancer histology images with deep neural networks. In *International conference on medical image computing and computer-assisted intervention*, pages 411–418. Springer, 2013.
- [7] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020.
- [8] Youssef Dawoud, Katharina Ernst, Gustavo Carneiro, and Vasileios Belagiannis. Edge-based self-supervision for semisupervised few-shot microscopy image cell segmentation. In *International Workshop on Medical Optical Imaging and Virtual Microscopy Image Analysis*, pages 22–31. Springer, 2022.
- [9] Youssef Dawoud, Julia Hornauer, Gustavo Carneiro, and Vasileios Belagiannis. Few-shot microscopy image cell segmentation. In Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track -European Conference, ECML PKDD 2020, Ghent, Belgium, September 14-18, 2020, Proceedings, Part V, volume 12461, pages 139–154. Springer, 2020.
- [10] Klaas Dijkstra, J Loosdrecht, LRB Schomaker, and Marco A Wiering. Centroidnet: A deep neural network for joint object localization and counting. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 585–601. Springer, 2018.
- [11] Geisa M Faustino, Marcelo Gattass, Stevens Rehen, and Carlos JP de Lucena. Automatic embryonic stem cells detection and counting method in fluorescence microscopy images. In 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pages 799–802. IEEE, 2009.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Modelagnostic meta-learning for fast adaptation of deep networks.

In International Conference on Machine Learning, pages 1126–1135. PMLR, 2017.

- [13] Linton C Freeman. Elementary applied statistics: for students in behavioral science. New York: Wiley, 1965.
- [14] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.
- [15] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192. PMLR, 2017.
- [16] Yunhe Gao, Mu Zhou, and Dimitris N Metaxas. Utnet: a hybrid transformer architecture for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 61–71. Springer, 2021.
- [17] Stephan Gerhard, Jan Funke, Julien Martel, Albert Cardona, and Richard Fetter. Segmented anisotropic sstem dataset of neural tissue. *figshare*, pages 0–0, 2013.
- [18] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. arXiv preprint arXiv:1803.07728, 2018.
- [19] Marc Gorriz, Axel Carlier, Emmanuel Faure, and Xavier Giro-i Nieto. Cost-effective active learning for melanoma segmentation. arXiv preprint arXiv:1711.09168, 2017.
- [20] Robert M Haralick, Stanley R Sternberg, and Xinhua Zhuang. Image analysis using mathematical morphology. *IEEE transactions on pattern analysis and machine intelli*gence, (4):532–550, 1987.
- [21] Amir Jamaludin, Timor Kadir, and Andrew Zisserman. Selfsupervised learning for spinal mris. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 294–302. Springer, 2017.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [23] Antti Lehmussola, Pekka Ruusuvuori, Jyrki Selinummi, Heikki Huttunen, and Olli Yli-Harja. Computational framework for simulating fluorescence microscope images with cell populations. *IEEE transactions on medical imaging*, 26(7):1010–1016, 2007.
- [24] Zhi Lu, Gustavo Carneiro, and Andrew P Bradley. An improved joint optimization of multiple level set functions for the segmentation of overlapping cervical cells. *IEEE Transactions on Image Processing*, 24(4):1261–1272, 2015.
- [25] Aurélien Lucchi, Yunpeng Li, and Pascal Fua. Learning for structured prediction using approximate subgradient descent with working sets. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 1987– 1994, 2013.
- [26] Anastasia Makarevich, Azade Farshad, Vasileios Belagiannis, and Nassir Navab. Metamedseg: Volumetric metalearning for few-shot organ segmentation. arXiv preprint arXiv:2109.09734, 2021.
- [27] Lu Mi, Hao Wang, Yonglong Tian, Hao He, and Nir Shavit. Training-free uncertainty estimation for dense regression:

Sensitivity as a surrogate. *arXiv preprint arXiv:1910.04858*, 2019.

- [28] Arnab Kumar Mondal, Jose Dolz, and Christian Desrosiers. Few-shot 3d multi-modal medical image segmentation using generative adversarial learning. *arXiv preprint arXiv:1810.12241*, 2018.
- [29] Peter Naylor, Marick Laé, Fabien Reyal, and Thomas Walter. Nuclei segmentation in histopathology images using deep neural networks. In 2017 IEEE 14th international symposium on biomedical imaging (ISBI 2017), pages 933–936. IEEE, 2017.
- [30] Peter Naylor, Marick Laé, Fabien Reyal, and Thomas Walter. Segmentation of nuclei in histopathology images by deep regression of the distance map. *IEEE transactions on medical imaging*, 38(2):448–459, 2018.
- [31] Alex Nichol, Joshua Achiam, and John Schulman. On firstorder meta-learning algorithms. *CoRR*, abs/1803.02999, 2018.
- [32] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016.
- [33] Cheng Ouyang, Carlo Biffi, Chen Chen, Turkay Kart, Huaqi Qiu, and Daniel Rueckert. Self-supervision with superpixels: Training few-shot medical image segmentation without annotation. In *European Conference on Computer Vision*, pages 762–780. Springer, 2020.
- [34] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal* of machine Learning research, 12:2825–2830, 2011.
- [35] Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter Haar Romeny, John B Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its variations. *Computer vision, graphics, and image processing*, 39(3):355– 368, 1987.
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. Unet: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [37] Burr Settles. Active learning literature survey. 2009.
- [38] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [39] Aiham Taleb, Winfried Loetzsch, Noel Danz, Julius Severin, Thomas Gaertner, Benjamin Bergner, and Christoph Lippert. 3d self-supervised methods for medical imaging, 2020.
- [40] Carolina Wählby, I-M Sintorn, Fredrik Erlandsson, Gunilla Borgefors, and Ewert Bengtsson. Combining intensity, edge and shape information for 2d and 3d segmentation of cell nuclei in tissue sections. *Journal of microscopy*, 215(1):67– 76, 2004.
- [41] Weidi Xie, J Alison Noble, and Andrew Zisserman. Microscopy cell counting and detection with fully convolu-

tional regression networks. *Computer methods in biomechanics and biomedical engineering: Imaging & Visualization*, 6(3):283–292, 2018.

- [42] Fuyong Xing and Lin Yang. Robust nucleus/cell detection and segmentation in digital pathology and microscopy images: a comprehensive review. *IEEE reviews in biomedical engineering*, 9:234–263, 2016.
- [43] Xianglilan Zhang, Hongnan Wang, Tony J Collins, Zhigang Luo, and Ming Li. Classifying stem cell differentiation images by information distance. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 269–282. Springer, 2012.