

DARCNN: Domain Adaptive Region-based Convolutional Neural Network for Unsupervised Instance Segmentation in Biomedical Images

Joy Hsu
Stanford University
joycj@stanford.edu

Wah Chiu
Stanford University
wahc@stanford.edu

Serena Yeung
Stanford University
syeung@stanford.edu

Abstract

In the biomedical domain, there is an abundance of dense, complex data where objects of interest may be challenging to detect or constrained by limits of human knowledge. Labelled domain specific datasets for supervised tasks are often expensive to obtain, and furthermore discovery of novel distinct objects may be desirable for unbiased scientific discovery. Therefore, we propose leveraging the wealth of annotations in benchmark computer vision datasets to conduct unsupervised instance segmentation for diverse biomedical datasets. The key obstacle is thus overcoming the large domain shift from common to biomedical images. We propose a Domain Adaptive Region-based Convolutional Neural Network (DARCNN), that adapts knowledge of object definition from COCO, a large labelled vision dataset, to multiple biomedical datasets. We introduce a domain separation module, a self-supervised representation consistency loss, and an augmented pseudo-labelling stage within DARCNN to effectively perform domain adaptation across such large domain shifts. We showcase DARCNN's performance for unsupervised instance segmentation on numerous biomedical datasets.

1. Introduction

State-of-the-art machine learning methods have accomplished a wide variety of impressive tasks including instance segmentation, yet much of their progress in the real world is limited to supervised methods with large, labelled datasets. In areas such as the biomedical domain, this is particularly problematic, as the prerequisite labels that accompany the complex data are often time consuming to obtain. In addition, we may also be constrained by human knowledge — biomedical data often contains unknown objects that scientists have yet to uncover, and therefore cannot accurately annotate.

Thus, there exists a need for methods that can produce instance segmentation for unlabelled datasets. We tackle this problem through solving the unsupervised do-

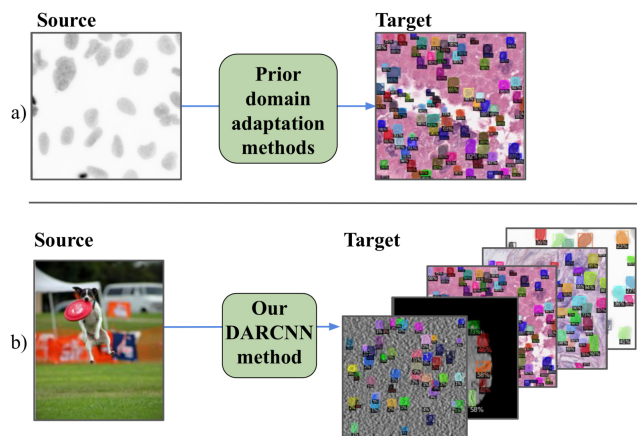


Figure 1. a) Prior domain adaptation methods for biomedical images tackle small domain shifts by using similar labelled biomedical datasets as sources to adapt to specific target datasets. b) DARCNN uses a common benchmark dataset as source and can adapt to a wide range of biomedical images.

main adaptation task, in which we use a source dataset with instance segmentation annotations to transfer knowledge and perform instance segmentation on target datasets.

Our choice of source dataset is motivated by the abundance of benchmark datasets in the vision field depicting common objects. We explore leveraging the large amount of labelled vision data in Common Objects in Context (COCO) [20] to achieve instance segmentation in diverse, natural biomedical images where annotations are difficult to obtain. Our main contributions include overcoming the large domain shift between natural images and biomedical images, and introducing a method for unsupervised instance segmentation on a wide range of biomedical datasets.

Past work tackling this problem in the biomedical field have depended on the availability of similar labelled biomedical datasets for the unsupervised instance segmentation task (see Figure 1), but it is not always feasible to find and annotate similar images. For these prior domain adaptation methods that focus on small domain shifts, joint image-level and feature-level adaptation approaches and object-

specific models have seen success [6, 7, 13, 17]. However, few works study unsupervised domain adaptation on large domain shifts such as from COCO to biomedical images, where such image-level adaptation fails. In addition, other past methods also design models specific to segmenting particular structures, which limits both application to other biomedical datasets as well as discovery [14, 21].

Hence we propose Domain Adaptive Region-based Convolutional Neural Network (DARCNN), a two stage class agnostic unsupervised domain adaptation model for instance segmentation of all distinct objects, capturing the notion of objectness. DARCNN first tackles feature-level adaptation, then refines segmentation masks through image-level pseudo-labelling. Our method can be applied to datasets with consistent background (e.g. of homogeneous cell background in microscopy) instead of split backgrounds (e.g. of the sky and grass as commonly seen in COCO). DARCNN leverages the success of the two step Mask R-CNN framework [12] and learns domain invariant and specific features for region proposal and segmentation mask prediction. The features are learned through a self-supervised background representation consistency loss based on predicted regions within an image.

In the second stage of DARCNN, pseudo-labelling on augmented input is introduced as a strong supervisory image-level signal. Through pseudo-labelling we are able to attain stable image-level segmentation after feature-level adaptation. We discover that our sequential two stage process is able to solve the domain adaptation task with large concept shift, shown on several biomedical datasets. In addition, we demonstrate that our method achieves strong performance on tasks of smaller domain shift as well.

Our key contributions are the following:

- We introduce a domain separation module to learn domain invariant and domain specific features for the two step instance segmentation framework.
- We propose a self-supervised representation consistency loss based on predicted regions within an image for feature adaptation.
- We utilize pseudo-labelling with data augmentation within DARCNN for strong image-level supervision.
- We demonstrate the effectiveness of our approach through quantitative experiments on adapting from COCO to five diverse biomedical datasets and a qualitative experiment for object discovery on a cryogenic electron tomography dataset.

2. Related Work

2.1. Unsupervised Domain Adaptation

Prior unsupervised domain adaptation approaches can be categorized into feature-level adaptation, image-level adaptation, or a combination of both. Feature-level adaptation

includes minimizing distances between source and target features through extracting shared domain features [10, 31], minimizing maximum mean discrepancy [23], or adversarial approaches such as [30, 33].

Image-level adaptation, such as those that tackle pixel-to-pixel translation between source and target domains, are often evaluated on adaptations between similar domains with no concept shift, such as from Cityspaces [9] to GTA [27]. Common works include [16, 36], which conducts image-to-image translation through generative adversarial networks. However, approaches such as pixel-to-pixel translation are extremely limited by size of domain shift.

Several key domain adaptive methods formulates adaptation across both image-level and feature-level, including [13, 17], while others approach domain shift on the instance-level and image-level as well [7]. These methods conduct feature-level and image-level adaptation jointly or rely heavily on image-level adaptation, which has seen impressive results in adaptation with small domain shifts, but struggle with larger concept shifts. DARCNN conducts first feature-level adaptation then image-level refinement sequentially in a two stage process, overcoming limitations of prior works.

In addition, a line of prior work that tackles small domain shift across biomedical images has shown strong performance on specific datasets. [6] transforms appearance of MR and CT images through synergistic fusion of adaptations from both feature-level and image-level, and [14] generates synthesized nucleus segmentation masks with importance weighting. [21] similarly introduces a nuclei inpainting mechanism for unsupervised nucleus segmentation through domain adaptation with re-weighting. However, these works tackle specific biomedical datasets, where methods can be crafted for detection of specific objects, and are also limited to smaller domain shifts, where an additional labelled biomedical dataset is needed for adaptation.

Our work overcomes the limitations of small domain shift and object-specific techniques for biomedical datasets through a two stage feature-level adaptation and image-level pseudo-labelling that segments all objects of interests. Previous works such as [5] has shown success in domain separation networks in simpler classification settings; similarly, we introduce a domain adaptation module and integrate this with a self-supervised loss that learns feature discriminability for instance segmentation.

2.2. Unsupervised Background-Foreground Segmentation

Prior works on unsupervised background-foreground segmentation primarily use a combination of consistency constraints and domain-specific assumptions. For example, [34] focuses on consistency between generated image and outputs of edge detectors, [28] leverages salient pixels in

the foreground and matching foregrounds between different images, and [3] utilizes a multi-task formulation with need for clean background images.

DARCNN similarly takes a self-supervised approach by maintaining a background representation consistency constraint, leveraging proposed regions within each image. Through this objective, our approach is able to learn domain invariant and domain specific representation for segmentation.

2.3. Pseudo-labels

Pseudo-labelling has often been used as a technique for utilizing unlabelled data in semi-supervised training. Prior work have chosen maximum predicted probability labels [19], uncertainty weighted class predictions [29], and used group-based label propagation [15]. Similarly, co-training methods use an ensemble of models to find labels through consistency regularization [26]. Fewer methods have used this method in the unsupervised setting, though works such as [8] have used high density clusters as pseudo-labels, inferring high confidence without supervision.

After first stage feature-level adaptation, our unsupervised method uses pseudo-labels to gain stronger image-level supervision, where high confidence pseudo-labels comes from first stage DARCNN. In addition, to learn across invariances, we add data augmentation to the unlabelled target images such as in [4], allowing DARCNN to learn across different imaging conditions.

3. Methods

We propose methods for unsupervised instance segmentation through the task of domain adaptation with large concept shift. In this section, we describe our two stage DARCNN model. The initial stage of feature-level adaptation consists of a domain separation module and a self-supervised representation consistency loss, which can be found in Section 3.1 and Section 3.2. The second stage of image-level pseudo-labelling can be found in Section 3.3.

We train each stage separately to tackle our problem of large domain shift. Image-level adaptation such as pixel-to-pixel translation does not work on such extreme concept shift, thus sequentially using image-level pseudo-labelling as a second stage allows for features to first learn to adapt between domains, before augmenting training with stronger pixel-level supervisory signal. DARCNN is pre-trained with source dataset weights, and jointly trains with a batch of source and target inputs. See Figure 2 for an overview of our model.

3.1. Domain Separation Module

The Mask R-CNN framework [12] is a powerful instance segmentation model, and we leverage its two step framework for DARCNN as well as propose a domain separation

module designed for our task. The region proposal network from the first step finds potential bounding boxes of interest given features learned through convolutional layers, while the mask prediction head from the second step refines these boxes and produces a mask for each instance.

To tackle the problem of domain shift, we propose a domain separation module that learns domain invariant and domain specific features as input into the region proposal network and mask segmentation network. The domain invariant features encode objectness of the source and target domain in a joint representational subspace, while the domain specific features capture discriminability of each domain as well as contain additional unconstrained embedding space.

The losses of the DARCNN are: L_{sim} to encourage domain invariant features, L_{diff} to learn domain specific features, L_{source} which includes the original Mask R-CNN losses for supervised training of the source dataset, and our proposed L_{target} for segmentation through a self-supervised consistency loss. Weighting factors α , β , and γ are used to balance the loss. See Equation 1 below.

$$L_{DARCNN} = \alpha L_{sim} + \beta L_{diff} + \gamma L_{target} + L_{source} \quad (1)$$

3.1.1 Domain Invariant Features

Intuitively, region proposals should be based on high level definition of objectness in the input image shared between both domains. Hence we encourage source and target domain invariant features to move into a joint representational subspace. The similarity loss helps the unlabelled target features better encode the objectness learned from the labelled source features. We utilize the maximum mean discrepancy loss L_{sim} from Equation 2 below. The maximum mean discrepancy loss [23, 5] is a kernel-based distance function between pairs of samples; we can think of the loss as computing the difference in distribution s and t where source inputs are drawn from s and target inputs are drawn from t . We let κ be our kernel function, and h_c^s and h_c^t be our shared source and shared target features respectively.

$$L_{sim} = \frac{1}{(N^s)^2} \sum_{i,j=0}^{N^s} \kappa(h_c^{s,i}, h_c^{s,j}) - \frac{2}{N^s N^t} \sum_{i,j=0}^{N^s N^t} \kappa(h_c^{s,i}, h_c^{t,j}) + \frac{1}{(N^t)^2} \sum_{i,j=0}^{N^t} \kappa(h_c^{t,i}, h_c^{t,j}) \quad (2)$$

In our implementation of maximum mean discrepancy, we use a Gaussian kernel κ . We downsample h_c^s and h_c^t with a dimension reduction convolutional layer with 1 filter of 1x1 kernel as an additional projection head to learn source and target domain invariant features.

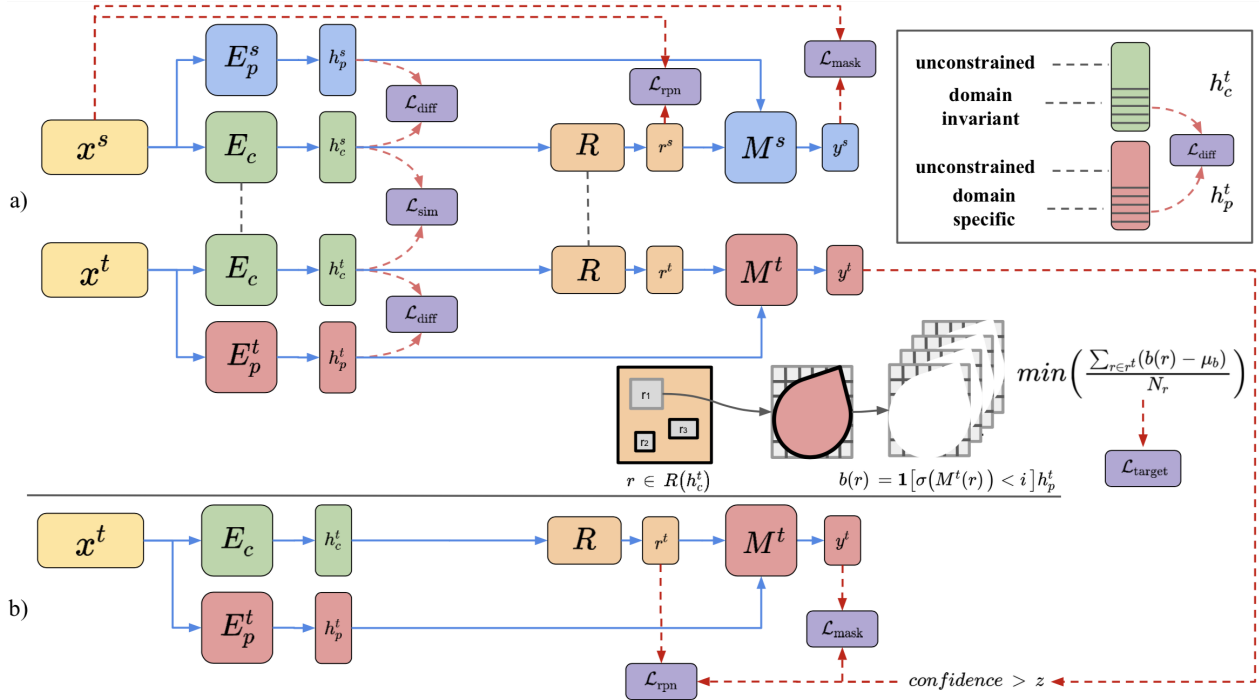


Figure 2. a) First stage DARCNN model with a domain separation module and self-supervised representation consistency loss. Let s and t represent source and target, and h_c and h_p represent common domain invariant features and private domain specific features respectively. E_c is the shared encoder, E_p^s and E_p^t are domain specific encoders, R is the shared region proposal network, and M^s and M^t are domain specific mask prediction heads. We let $b(r)$ be extracted background features for each region r . Top right corner showcases our soft orthogonality constraint on half of the domain specific features. b) Second stage DARCNN model with pseudo-labels of augmented input from first stage DARCNN, with annotations chosen over a confidence threshold, continuing training of DARCNN’s target branch.

However, as we focus on scenarios where domain shift between our source and target domains is large, the initial distance between s and t as calculated by the maximum mean discrepancy loss is also large. [32] shows that minimizing the maximum mean discrepancy loss equates to maximizing the source and target intra-class distances respectively, but doing so also jointly minimizes their variance with some implicit weights, such that feature discriminability degrades. Therefore, if L_{sim} too quickly overwhelms other losses in DARCNN that retain semantic features, the discriminability needed for instance segmentation is lost.

Hence we propose a maximum mean discrepancy loss that uses a warmup weighting scheduler. Our approach increases the weight α of L_{sim} from α_0 to α over n epochs, where α_0 is smaller when the domain shift is larger.

3.1.2 Domain Specific Features

We next consider input needed for mask predictions. Domain specific features captures feature discriminability for the target and source domains as well as granularity of the background representation for our self-supervised loss. Hence we use L_{diff} to separate information that is unique to

each domain as well as learn specificity, and define the loss through the soft subspace orthogonality constraint from [5] between the domain invariant and domain specific features of both domains.

However, in order to let part of the mask feature representations learn semantically relevant embeddings that may potentially be domain invariant, we utilize the orthogonality difference loss only between parts of the domain invariant and domain specific features. We give DARCNN the freedom to learn features necessary for segmentation in an unconstrained embedding space, whether it be domain invariant or specific. In our implementation, we use half of the feature depth. See top right corner of Figure 2. We let H be matrices whose rows are half of hidden representations h in depth, where h_c^s and h_p^s are the invariant and specific features of the source, and h_c^t and h_p^t of the target.

$$L_{diff} = \|H_c^{s\top} H_p^s\|_F^2 + \|H_c^{t\top} H_p^t\|_F^2 \quad (3)$$

To additionally give signal to the unlabelled target domain specific features, we propose a self-supervised representation consistency loss instead of the reconstruction loss used for classification in [5]. The source domain specific features are supervised by the original Mask R-CNN bounding box and mask losses.

3.2. Self-Supervised Representation Consistency

Biomedical images commonly contain homogeneous backgrounds, therefore we leverage this assumption between region proposals of the same image to self-supervise our feature representations. In contrast to approaches that define a global background consistency across images of a dataset, we use independent background consistency for each image, which allows for variation in backgrounds within the dataset. We leverage the region proposal network and minimize the differences between background representations of each predicted instance. See Figure 2.

We accomplish this through the two step framework of Mask R-CNN. DARCNN utilizes self-supervision during training through first finding the top region proposals with confidence over threshold k from the region proposal network. It passes each high confidence region proposal through the class agnostic mask head and determines which parts of the predicted instance are background. To do this, outputs from the mask head are passed through a sigmoid activation σ and all values less than threshold i are taken as background.

Then, we retrieve features output from our convolutional encoder E_p^t that corresponds to background predictions. We minimize the differences between these background representations. Letting r be a predicted region from shared region proposal network R and M^t be our target mask prediction head, we define the background features $b(r)$ in Equation 4, where the indicator function extracts parts of h_p^t , domain specific features, that spatially corresponds to background mask predictions, threshold by value i after sigmoid function σ . We can then define μ_b , computed per image, as the mean of background representations across all regions in an image. These regions are predicted by region proposal network R after taking domain invariant features h_c^t as inputs. Finally, we define L_{target} to minimize differences between background features.

$$b(r) = \mathbb{1}[\sigma(M^t(r)) < i]h_p^t \quad (4)$$

$$\mu_b = \frac{1}{N_p} \sum_{p \in R(h_c^t)} b(p) \quad (5)$$

$$L_{\text{target}} = \frac{1}{N_r} \sum_{r \in R(h_c^t)} |b(r) - \mu_b| \quad (6)$$

In our training process, the combination of the fully supervised and self-supervised losses from the source and target dataset respectively allows DARCNN to learn semantically relevant proposals and mask predictions.

3.3. Augmented Pseudo-Labeling

Our first stage DARCNN utilizes feature-level adaptations for unsupervised domain adaptation and leads to initial coarse mask predictions that overcome large domain

shift. However, it lacks strong image-level supervisory signal as in [13] or [7]; the lack of a pixel-level signal leads to more unstable and unrefined segmentations. Therefore we propose a second stage image-level pseudo-labelling with our first stage DARCNN’s output as pseudo-labels in order to gain this image-level supervision. See part b) in Figure 2. Only the target branch of DARCNN is trained during the second stage pseudo-labelling process, while the source branch is frozen and no longer needed.

Canonical use of pseudo-labels depends on some amount of labelled data, however, as our method is unsupervised, we instead use high confidence predictions from our first stage DARCNN. We use confidence threshold z to determine which labels to retrieve from the predictions.

In addition, to better learn invariances of labels despite imaging conditions, including quality and noise, we apply data augmentation procedures to strengthen pseudo-labels from the first stage DARCNN. The augmentations ensure that the same instance segmentations will be predicted of a given input regardless of lighting, contrast, and blur.

The target branch of DARCNN is the final model to be used for unsupervised instance segmentation.

4. Experiments

In our experiments, we show that we are able to overcome limitations of past work to adapt between large domain shifts, as well generalize across many biomedical datasets. We quantitatively demonstrate DARCNN’s performance on a large domain shift from COCO [20] to multiple biomedical datasets, bypassing drawbacks of previous literature that focuses on small domain shifts and specific datasets. In addition, in order to directly compare against existing work, we also show comparable performance to prior methods on tasks with small domain shifts following canonical biomedical adaptations of [21].

Our work is not limited to that of biomedical datasets, and is designed for all datasets with consistent backgrounds. We use biomedical datasets due to prevalence of homogeneous backgrounds in the biomedical field, for evaluation following previous unsupervised instance segmentation work [21], and to better illustrate our approach on large domain shifts from generalized COCO to diverse, biomedical datasets.

4.1. Implementation details

We evaluate our experiments on Aggregated Jaccard Index (AJI) [18]. AJI is used by prior work to evaluate the performance of instance segmentation; it computes an aggregated intersection cardinality numerator, and an aggregated union cardinality denominator for all ground truth and segmented predictions under consideration. It is a unified metric that measures both object-level and pixel-level performance, and is more stringent than other canonical metrics

such as IOU. In addition, we also show pixel F1 score and object F1 score to measure performance in specific aspects. The annotations for all five target biomedical datasets are not used during unsupervised training of DARCNN, only for evaluation. For our experiments, we stop our model training 0.1 epochs before loss plateaus.

We use Pytorch and build on the Detectron2 [35] framework, using the ResNet backbone. For our loss function L_{DARCNN} , we set α to increase over the first 0.1 epochs to 1, $\beta = 1$ and $\gamma = 0.1$. We use $k = 0.5$ as the confidence threshold for top predicted regions in our self-supervised loss, and $i = 0.5$ as our threshold for background. Confidence threshold z is set to be 0.5 for pseudo-labelling. Augmentation parameters used are Gaussian blur with sigma as 1, and contrast and brightness are changed through scaling and delta factors 1.5 and -150 respectively. We use learning rate 0.0001, and vary maximum number of detections to return per image during inference through initial coarse inspection of training images. We choose the number to be 100 or 50 accordingly for each dataset.

4.2. Adaptation from Microscopy to Histopathology

To compare against prior unsupervised domain adaptation methods tackling small domain shifts between biomedical images, we first quantitatively evaluate adaptation from a fluorescence microscopy dataset, BBBC [22], to two histopathology datasets, Kumar [18] and TNBC [25]. This comparison follows that of Liu et al. [21], and we follow the same implementations of prior work and evaluation.

Importantly, we also demonstrate DARCNN’s strong performance when adapting from a common dataset, COCO [20], to the same two histopathology datasets. We demonstrate that even without a similar source biomedical dataset that may be difficult to obtain, we are still able to conduct unsupervised instance segmentation adapting from COCO.

We first preprocess our source dataset, BBBC. A total of 100 training images and 50 validation images from BBBC are used, following the official data split. 10,000 patches of BBBC in size 256x256 are randomly cropped from the 100 training images, and pixel values are inverted to better synthesize histopathology images following [21]. Both Kumar [18] and TNBC [25], our target datasets, are trained with 10,000 patches of size 256x256 without any labels, and evaluated on the specified test set. To compare our method against the nucleus specific methods of [14] and [21], we utilize the standalone, non-deep learning based unsupervised synthesis module of [14] as additional input.

We can see in Table 1 that for the TNBC dataset in the scenario of a small domain shift, our method with fluorescence microscopy as source outperforms all prior work aside from [21]. [21] utilizes specific nuclei inpainting, which yields impressive performance, but can only be used for nuclei segmentation tasks.

| Method | AJI | Pixel-F1 | Object-F1 |
|-----------------------|--------|----------|-----------|
| Chen et al. [7] | 0.4407 | 0.6405 | 0.6289 |
| DDMRL [17] | 0.4642 | 0.7000 | 0.6872 |
| SIFA [6] | 0.4662 | 0.6994 | 0.6698 |
| CyCADA [13] | 0.4721 | 0.7048 | 0.6866 |
| Hou et al. [14] | 0.4775 | 0.7029 | 0.6779 |
| Liu et al. [21] | 0.5672 | 0.7593 | 0.7478 |
| Ours from BBBC | 0.5120 | 0.7175 | 0.6436 |
| Ours from COCO | 0.4906 | 0.6998 | 0.6396 |

Table 1. Comparison of unsupervised methods adapting BBBC to TNBC. We see that DARCNN shows strong performance against prior work both with BBBC as source and with COCO as source, even against methods designed for nucleus-specific segmentation [14, 21] and small domain shifts [7, 17, 6, 13].

| Method | AJI | Pixel-F1 | Object-F1 |
|-----------------------|--------|----------|-----------|
| Chen et al. [7] | 0.3756 | 0.6337 | 0.5737 |
| SIFA [6] | 0.3924 | 0.6880 | 0.6008 |
| CyCADA [13] | 0.4447 | 0.7220 | 0.6567 |
| DDMRL [17] | 0.4860 | 0.7109 | 0.6833 |
| Hou et al. [14] | 0.4980 | 0.7500 | 0.6890 |
| Liu et al. [21] | 0.5610 | 0.7882 | 0.7483 |
| Ours from BBBC | 0.4461 | 0.6619 | 0.5410 |
| Ours from COCO | 0.4421 | 0.6549 | 0.5104 |

Table 2. Comparison of unsupervised methods adapting BBBC to Kumar. DARCNN shows comparable performance to prior works on the nucleus segmentation task though it is designed to over-segment and retrieve all instances of interest.

More importantly, we show that DARCNN with COCO as the source dataset is able to achieve similar performance, also outperforming other methods aside from [21]. Without labels from BBBC, we are still able to adapt from COCO to histopathology datasets, which is essential in cases where similar labelled biomedical datasets do not exist.

In Table 2, we see that for the Kumar dataset, due to DARCNN predicting more objects than are classified as nucleus, our scores are comparable but do not beat all of prior methods. We hypothesis that this is because past work focuses on image-level adaptation at small domain shifts, and their models are able to better learn what the object of focus is for segmentation. DARCNN is designed for instance segmentation of all distinct objects that exist within an image, hence predicts false positives from the nucleus perspective. See below Figure 3 for an example of DARCNN adapted from COCO to Kumar; note that the red objects in the middle are not considered nucleus from the ground truth, but are segmented by DARCNN as objects of interest. DARCNN is useful for discovering objects in data when no labels are available, and produces class agnostic instance segmentations that can help scientists uncover new objects in complex data. Postprocessing methods could allow us to gain more insight into objects found by DARCNN. Through




Figure 3. Qualitative results on DARCNN, adaptation from COCO to TNBC (top) & Kumar (bottom) respectively. Left is input, middle is ground truth, and right shows instance segmentation result.

rule-based filtering with a coarse prior (filtering for masks with average pixel value less than a threshold, over a set of thresholds), our method obtains AJI = 0.5026, Pixel-F1 = 0.7272, Object-F1 = 0.6481, outperforming all prior work except [21], which is designed for nuclei segmentation.

4.3. Ablation Studies

We conduct ablation studies on DARCNN to showcase the effectiveness of the domain similarity loss, the background representation consistency loss, and the augmented pseudo-labelling stage. The DARCNN here is shown with COCO [20] as the source dataset, tackling our stated problem of large domain shift to biomedical datasets.

We perform this ablation study for two separate setups — first we study DARCNN adapted to TNBC, trained with the unsupervised, standalone synthesis module of [14], which allows for its comparable performance to previous nucleus-specific methods. Then we observe DARCNN’s ablation performance for the BBBC dataset, showcasing each component’s contribution to DARCNN without initial object-specific synthesis.

In Table 3 and Table 4, we observe that class agnostic Mask R-CNN trained on COCO images performs extremely poorly on biomedical datasets. We can also see that for the first stage of DARCNN, both the domain similarity loss and representation consistency loss improved model performance. Especially in the case of BBBC without the unsupervised synthesis module, our representation consistency loss dramatically improves the performance of DARCNN due to ability to gain self-supervised signal for features. The full first stage DARCNN also shows significantly higher performance in both datasets than the initial baselines.

In the second stage of DARCNN, pseudo-labelling also improved our unsupervised instance segmentation performance. Similarly, the performance improvement from pseudo-labelling in the BBBC case without the synthesis

| Method | AJI | Pixel-F1 | Object-F1 |
|------------------------------|--------|----------|-----------|
| Mask R-CNN | | | |
| w/ COCO pre-trained | 0.0060 | 0.2769 | 0.0181 |
| w/ synthesized images | 0.3332 | 0.5782 | 0.6061 |
| First stage DARCNN | | | |
| Domain sim. only | 0.3687 | 0.6023 | 0.6099 |
| Bg. consistency only | 0.3808 | 0.6120 | 0.5470 |
| Full 1st stage DARCNN | 0.4071 | 0.6353 | 0.5986 |
| Second stage DARCNN | | | |
| Pseudo-label w/o aug | 0.4463 | 0.6781 | 0.6339 |
| Full 2nd stage DARCNN | 0.4906 | 0.6998 | 0.6396 |

Table 3. Ablation study adapting from COCO as source to TNBC.

| Method | AJI | Pixel-F1 | Object-F1 |
|------------------------------|--------|----------|-----------|
| Mask R-CNN | | | |
| w/ COCO pre-trained | 0.0315 | 0.3144 | 0.0818 |
| First stage DARCNN | | | |
| Domain sim. only | 0.1414 | 0.4905 | 0.4295 |
| Bg. consistency only | 0.3250 | 0.7128 | 0.5720 |
| Full 1st stage DARCNN | 0.3371 | 0.6409 | 0.5904 |
| Second stage DARCNN | | | |
| Pseudo-label w/o aug | 0.4349 | 0.6914 | 0.7151 |
| Full 2nd stage DARCNN | 0.4725 | 0.6586 | 0.6733 |

Table 4. Ablation study adapting from COCO as source to BBBC.

module is larger than in TNBC, as for BBBC it is the first time DARCNN’s target branch receives image-level supervision. Pseudo-labelling with augmented data demonstrates even better performance, helping remove difficulties in segmentation under various imaging conditions of biomedical datasets. We also show in Figure 3 a qualitative example of DARCNN adapting from COCO to TNBC, capturing nuclei as objects of interests in the image.

4.4. Adaptation from COCO to Additional Biomedical Datasets

Most importantly, we demonstrate DARCNN’s ability to generalize across datasets by comparing performance on COCO adapted to three diverse biomedical datasets – the fluorescence microscopy dataset [22], cryogenic electron tomography dataset [11], and brain MRI dataset [24, 1, 2]. Through this, we demonstrate potential for object discovery without the need for similar labelled datasets to the target domain, and without need for designing object-specific models for segmentation of particular biomedical datasets.

We compare our results with the prior methods that can be used outside of specific biomedical datasets. [6, 14, 21] depend on specific biomedical images and nuclei synthesis methods, hence we do not use them as comparison.

DARCNN significantly outperforms all other methods when adapting with a large domain shift from COCO. As [13, 17] depend on CycleGAN [36], and [7] depends on image-level shift, we hypothesize that these methods per-

| Dataset | Method | AJI | Pixel-F1 | Object-F1 |
|---------|-----------------|---------|----------|-----------|
| BBBC | Chen et al. [7] | 0.1500 | 0.5251 | 0.2111 |
| | CyCADA [13] | 0.1231 | 0.5173 | 0.1917 |
| | DDMRL [17] | 0.0928 | 0.5188 | 0.1145 |
| | Ours | 0.4725 | 0.6586 | 0.6733 |
| SHREC | Chen et al. [7] | 0.0 | 0.0 | 0.0 |
| | CyCADA [13] | 0.0051 | 0.0064 | 0.0025 |
| | DDMRL [17] | 0.0039 | 0.0182 | 0.0 |
| | Ours | 0.1268 | 0.3007 | 0.3371 |
| BraTS | Chen et al. [7] | 0.2868* | - | - |
| | CyCADA [13] | 0.3485* | - | - |
| | DDMRL [17] | 0.3951* | - | - |
| | Ours | 0.5577* | - | - |

Table 5. Comparison of unsupervised methods adapting from COCO to fluorescence microscopy (BBBC), cryogenic electron tomography (SHREC), and radiology (BraTS) datasets. *Indicates metric measuring maximum intersection over union.

form poorly when tasked with a large domain shift from COCO to natural biomedical images. DARCNN is able to generalize across diverse biomedical datasets and adapt between common objects to objects in microscopy, tomography, and MRI due to its two stage sequential feature-level adaptation and image-level pseudo-labelling.

For BBBC, DARCNN is able to significantly outperform prior work as the background consistency assumption is strong. In addition, even for a more challenging task like instance segmentation in SHREC where the signal to noise ratio is low, DARCNN is still able to capture objects of interests given noisy background. Prior work, CyCADA [13], Chen et al. [7], and DDMRL [17], are all unable to learn meaningful adaptations from COCO to SHREC due to the difficulties of image-level adaptation when even human recognition is limited. In the BraTS dataset, due to the specificity of the detected object, tumor, we measure performance by maximum intersection over union from closest predicted object. Though this does not account for unbounded false positives, we provide a qualitative example of performance to supplement. In Figure 4, we see that DARCNN also picks up on ridges and darker spots in the BraTS MRI, which could potentially be useful in understanding additional structures of interest. We show qualitative examples on our three biomedical datasets in Figure 4.

4.5. Adaptation from COCO to CryoET

Finally, we showcase the promise of our unsupervised instance segmentation model for adapting from COCO to an unlabelled cryogenic electron tomography (cryoET) dataset, collected by Dr. Wah Chiu’s group at SLAC National Accelerator Laboratory. This cryoET dataset contains tomograms of crowded cellular environments in which objects are too dense and of too underexplored a subject area to be annotated. We qualitatively evaluate the performance of DARCNN instance segmentations in this dataset.




Figure 4. Qualitative results on DARCNN, adaptation from COCO to BBBC (top), SHREC (middle), and BraTS (bottom) datasets.

In Figure 5, our unsupervised algorithm segments known biological objects such as the autophagosome and granules. In addition, DARCNN also discovers distinct vesicles and organelles inside the amphisome, difficult for humans to annotate, and representing potential new objects of interest.




Figure 5. Qualitative results on DARCNN, adaptation from COCO to cryoET, demonstrating discovery of distinct objects of interests.

5. Conclusion

We propose DARCNN, a two stage feature-level adaptation and image-level pseudo-labelling method for unsupervised instance segmentation. We leverage the abundance of labelled benchmark datasets for domain adaptation to unlabelled biomedical images. DARCNN tackles large domain shifts between common and biomedical objects, and can be used across diverse datasets with consistent background. Through a domain separation module, a representation consistency loss, and augmented pseudo-labelling, we achieve strong performance in multiple experiments as well as show potential for object discovery within biomedical datasets.

Acknowledgments. We thank support from the Chan-Zuckerberg Initiative (to J.H., S.Y. and W.C.) and National Institutes of Health (grant number P01NS092525 to W.C.).

References

- [1] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4:170117, 2017.
- [2] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018.
- [3] Yaniv Benny and Lior Wolf. Onegan: Simultaneous unsupervised learning of conditional image generation, foreground segmentation, and fine-grained clustering. *arXiv preprint arXiv:1912.13471*, 2019.
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 5049–5059, 2019.
- [5] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *Advances in neural information processing systems*, pages 343–351, 2016.
- [6] Cheng Chen, Qi Dou, Hao Chen, Jing Qin, and Pheng-Ann Heng. Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 865–872, 2019.
- [7] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018.
- [8] Jaehoon Choi, Minki Jeong, Taekyung Kim, and Changick Kim. Pseudo-labeling curriculum for unsupervised domain adaptation. *arXiv preprint arXiv:1908.00262*, 2019.
- [9] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [10] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015.
- [11] Ilja Gubins, Marten L Chaillet, Gijs van der Schot, Remco C Veltkamp, Friedrich Förster, Yu Hao, Xiaohua Wan, Xuefeng Cui, Fa Zhang, Emmanuel Moebel, et al. Shrec 2020: Classification in cryo-electron tomograms. *Computers & Graphics*, 91:279–289, 2020.
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [13] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018.
- [14] Le Hou, Ayush Agarwal, Dimitris Samaras, Tahsin M Kurc, Rajarsi R Gupta, and Joel H Saltz. Robust histopathology image analysis: to label or to synthesize? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8533–8542, 2019.
- [15] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5070–5079, 2019.
- [16] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [17] Taekyung Kim, Minki Jeong, Seunghyeon Kim, Seokeon Choi, and Changick Kim. Diversify and match: A domain adaptive representation learning paradigm for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12456–12465, 2019.
- [18] Neeraj Kumar, Ruchika Verma, Sanuj Sharma, Surabhi Bhargava, Abhishek Vahadane, and Amit Sethi. A dataset and a technique for generalized nuclear segmentation for computational pathology. *IEEE transactions on medical imaging*, 36(7):1550–1560, 2017.
- [19] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, 2013.
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [21] Dongnan Liu, Donghao Zhang, Yang Song, Fan Zhang, Lauren O’Donnell, Heng Huang, Mei Chen, and Weidong Cai. Unsupervised instance segmentation in microscopy images via panoptic domain adaptation and task re-weighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4243–4252, 2020.
- [22] Vebjorn Ljosa, Katherine L Sokolnicki, and Anne E Carpenter. Annotated high-throughput microscopy image sets for validation. *Nature methods*, 9(7):637–637, 2012.
- [23] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pages 97–105. PMLR, 2015.
- [24] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.

- [25] 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.
- [26] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan Yuille. Deep co-training for semi-supervised image recognition. In *Proceedings of the european conference on computer vision (eccv)*, pages 135–152, 2018.
- [27] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, pages 102–118. Springer, 2016.
- [28] Michael Rubinstein, Armand Joulin, Johannes Kopf, and Ce Liu. Unsupervised joint object discovery and segmentation in internet images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1939–1946, 2013.
- [29] Weiwei Shi, Yihong Gong, Chris Ding, Zhiheng MaXiaoyu Tao, and Nanning Zheng. Transductive semi-supervised deep learning using min-max features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 299–315, 2018.
- [30] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. *arXiv preprint arXiv:1511.05547*, 2015.
- [31] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017.
- [32] Wei Wang, Haojie Li, Zhengming Ding, and Zhihui Wang. Rethink maximum mean discrepancy for domain adaptation. *arXiv preprint arXiv:2007.00689*, 2020.
- [33] Jun Wen, Risheng Liu, Nenggan Zheng, Qian Zheng, Zhefeng Gong, and Junsong Yuan. Exploiting local feature patterns for unsupervised domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5401–5408, 2019.
- [34] John Winn and Nebojsa Jojic. Locus: Learning object classes with unsupervised segmentation. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 1, pages 756–763. IEEE, 2005.
- [35] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2, 2019.
- [36] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.