



Published in final edited form as:

Med Image Anal. 2023 January ; 83: 102641. doi:10.1016/j.media.2022.102641.

Memory Consistent Unsupervised Off-the-Shelf Model Adaptation for Source-Relaxed Medical Image Segmentation

Xiaofeng Liu^a, Fangxu Xing^a, Georges El Fakhri^a, Jonghye Woo^a

^aGordon Center for Medical Imaging, Department of Radiology, Massachusetts General Hospital and Harvard Medical School, Boston, MA, 02114

Abstract

Unsupervised domain adaptation (UDA) has been a vital protocol for migrating information learned from a labeled source domain to facilitate the implementation in an unlabeled heterogeneous target domain. Although UDA is typically jointly trained on data from both domains, accessing the labeled source domain data is often restricted, due to concerns over patient data privacy or intellectual property. To sidestep this, we propose “off-the-shelf (OS)” UDA (OSUDA), aimed at image segmentation, by adapting an OS segmentor trained in a source domain to a target domain, in the absence of source domain data in adaptation. Toward this goal, we aim to develop a novel batch-wise normalization (BN) statistics adaptation framework. In particular, we gradually adapt the domain-specific low-order BN statistics, e.g., mean and variance, through an exponential momentum decay strategy, while explicitly enforcing the consistency of the domain shareable high-order BN statistics, e.g., scaling and shifting factors, via our optimization objective. We also adaptively quantify the channel-wise transferability to gauge the importance of each channel, via both low-order statistics divergence and a scaling factor. Furthermore, we incorporate unsupervised self-entropy minimization into our framework to boost performance alongside a novel queued, memory-consistent self-training strategy to utilize the reliable pseudo label for stable and efficient unsupervised adaptation. We evaluated our OSUDA-based framework on both cross-modality and cross-subtype brain tumor segmentation and cardiac MR to CT segmentation tasks. Our experimental results showed that our memory consistent OSUDA performs better than existing source-relaxed UDA methods and yields similar performance to UDA methods with source data.

Keywords

Unsupervised domain adaptation; Image segmentation; Batch Normalization; Self-training; Memory-based learning

1. Introduction

Accurate delineation of lesions or anatomical structures is a critical step for clinical intervention and treatment planning and has been markedly progressed over the past several years, mainly due to advances in deep neural networks (DNN) (Tajbakhsh et al., 2020).

A deep segmentor trained on source domain data, however, cannot generalize well in a heterogeneous target domain, e.g., different clinical centers, subtypes, scanner vendors, or imaging modalities (Ghafoorian et al., 2017; Liu et al., 2021e). Additionally, it often poses a great challenge to annotate labels for new target domain data (Che et al., 2019). To mitigate these issues, unsupervised domain adaptation (UDA) has been proposed, as a promising technique, to achieve knowledge transfer from a labeled source domain to an unlabeled target domain (Liu et al., 2021f,d).

Early attempts at UDA include statistic moment matching (Long et al., 2015), feature/pixel-level adversarial learning (Liu et al., 2021a), and self-training (Zou et al., 2019; Liu et al., 2021f), all of which are dependent on the joint training on both source and target domain data. Well-labeled source domain data, however, are often inaccessible, due to concerns over patient data privacy or intellectual property (IP) (Bateson et al., 2020). As such, it is of great interest and needs to develop an adaptation strategy using an “off-the-shelf (OS)” source domain model, without access to the source domain data.

In recent years, a source-free UDA approach for classification (Liang et al., 2020) was proposed to yield multiple predictions, but producing the diverse neighboring predictions is ill-suited for the segmentation purpose. A source relaxed UDA approach for segmentation (Bateson et al., 2020) was also proposed to train an auxiliary class-ratio prediction model with source domain data, by relying on the assumption that the pixel proportion in segmentation is consistent between source and target domains. There are two major limitations in that work. First, the class-ratio in two domains can be different, due to label shift (Kouw, 2018; Liu et al., 2021a,e,b), e.g., the incident rate of the disease and tumor size could vary depending on different tumor subtypes or populations. In addition, the class-ratio prediction model requires an extra training step with the source domain data.

To address the aforementioned limitations, in this work, we propose a novel source-free UDA approach for segmentation, without an additional auxiliary network trained on the source domain data or reliance on the assumption of the same class proportion between two domains. Fig. 1 highlights the characteristics of conventional transfer learning approaches and our “off-the-shelf” UDA (OSUDA). Our OSUDA does not need any label supervision in neither source nor target domain in adaptation, which works under a restrictive setting, compared with the conventional approaches (Long et al., 2015; Liu et al., 2021a,f). In our prior work (Liu et al., 2021g), we propose to leverage batch-wise normalization statistics that can be easily accessed and computed. Specifically, in modern deep learning backbones, such as ResNet (He et al., 2016) and U-Net (Zhou and Yang, 2019), Batch Normalization (BN) (Ioffe and Szegedy, 2015) has been widely used to achieve fast and stable training. After training, the BN statistics are typically stored alongside the network parameters. Recent literature on source-available UDA methods has indicated that the low-order BN statistics, including the mean and variance in BN, are domain-specific, owing to the discrepancy of input data (Chang et al., 2019). To achieve a gradual adaptation of the low-order BN statistics in a source-relaxed manner, a novel momentum-based progression strategy is presented, in which the momentum follows an exponential decay over the adaptation iteration. In addition, for a seamless transfer of the high-order BN statistics, including domain shareable scaling and shifting factors (Maria Carlucci et al., 2017), a high-

order BN statistics consistency loss is proposed to enforce the discrepancy minimization. To this end, the transferability of each channel is measured first in an adaptive manner, followed by gauging the channel-wise importance. Further, unsupervised self-entropy (SE) minimization is used to enhance adaptation performance.

In the present work, we build on our prior work (Liu et al., 2021g) as follows. First, we investigate and leverage the correlation between the high-order scaling factor and channel transferability for adaptive BN-based adaptation (in Sec. 3.2.3). Second, we develop a novel queued memory consistent OSUDA (MCO-SUDA) in order to leverage the reliable pseudo label for stable and efficient adaptation (Secs. 3.4 and 3.5). Third, we further provide a detailed interpretation of our framework in the context of recent work in this area (Liu et al., 2021h,a; Kundu et al., 2021; Chen et al., 2021; You et al., 2021). Finally, in order to validate our framework, we experiment on brain tumor segmentation as well as cardiac segmentation tasks. New comparisons are also made with recent UDA work for all of the tasks to demonstrate the validity and superiority of our framework.

Our contributions are summarized as follows:

- We propose a novel UDA segmentation framework in the absence of the source domain data, which thus only relies on an OS segmentor with BN trained on the source domain data. We do not need an additional auxiliary model trained in the source domain, or the assumption of the class-ratio consistency as in Bateson et al. (2020).
- We systematically explore the domain-specific and shareable BN statistics, by means of the low-order BN statistics progression with an exponential momentum decay (EMD) strategy and the quantified transferability adaptive high-order BN statistics consistency loss, respectively.
- In addition to unsupervised SE minimization, we propose a queued memory consistent self-training for stable and efficient progression of the pseudo label in the target domain, via the memory-based supervision signal, conditioned on the historical consistency.
- We evaluate our framework on both (cross-modality and cross-subtype) brain tumor segmentation and cardiac MR to CT segmentation tasks using BraTS2018 and MM-WHS databases, to demonstrate the general efficacy of our framework and its superiority to conventional source-free/source-available UDA segmentation methods.

2. Related Work

Unsupervised Domain Adaptation

Unsupervised Domain Adaptation has been widely used to migrate domain knowledge of one domain to another (Kouw, 2018; Liu et al., 2021a). Conventional approaches for UDA utilized the data in both domains for adaptation (Long et al., 2015; He et al., 2020; Zou et al., 2019). This setting, however, requires the sharing of labeled source domain data, which poses a data privacy concern in real-world implementations. Recently, source-free UDA for

classification (Li et al., 2020a,b; Liang et al., 2020; Wang et al., 2021) has been proposed to only use a pre-trained classifier, rather than co-training of the network with source and target domain data. Whereas these methods were applied only to classification or object detection, the present work focuses on segmentation tasks, i.e., fine-grained pixel-wise classification. CRUDA (Bateson et al., 2020) pre-trained a class-ratio prediction model in the source domain, and enforced the pixel proportion consistency between two domains for segmentation.

To the best of our knowledge, our prior work (Liu et al., 2021g) was the first attempt at OS source-relaxed UDA for segmentation, without the need for an additional auxiliary network, trained on the source domain data, or the assumption of the same class-ratio similar to CRUDA (Bateson et al., 2020). The follow-up works proposed methods based on knowledge distillation (Liu et al., 2021h) and source data hallucination (Ye et al., 2021), but an additional attention network or generative modules are required. In addition, Kundu et al. (2021) demanded a specific training strategy in the source domain, which does not utilize an OS segmentor, unlike our approach. Moreover, You et al. (2021) proposed a negative training strategy, designed specifically for classifying numerous classes, which is equivalent to the vanilla solution for a binary case. These works are also compared and discussed in relation to our proposed work.

Batch Normalization

Batch Normalization has been widely used to stabilize network training (Ioffe and Szegedy, 2015), by eliminating an internal covariate shift. Early attempts of applying BN to adaptation simply added BN to a target domain, and did not have an interaction with the source domain (Li et al., 2018). Recent work (Chang et al., 2019; Maria Carlucci et al., 2017; Wang et al., 2019a; Mancini et al., 2018) demonstrated that the low-order batch statistics, including the mean and variance in BN, are domain-specific, owing to the divergence of feature representations. Note that simply forcing the mean and variance in the target domain to be the same as the source domain can lead to a loss of expressiveness of the networks (Zhang et al., 2020). In addition, once the low-order BN statistics discrepancy has been partially mitigated, the high-order BN statistics can be shareable between two domains (Maria Carlucci et al., 2017; Wang et al., 2019a).

However, all of the aforementioned methods (Chang et al., 2019; Maria Carlucci et al., 2017; Zhang et al., 2020; Wang et al., 2019a; Mancini et al., 2018) need joint training on source domain data. In the present work, however, we rather opt to reduce the domain discrepancy, using a momentum-based adaptive low-order batch statistics progression strategy and explicit high-order BN statistics consistency loss for our source-free UDA for segmentation.

Memory-based Learning

Memory-based Learning was initially proposed to stabilize supervised training with external modules to store memory (Weston et al., 2014). Then, the idea of the memory mechanism has been generalized to semi-supervised learning (Tarvainen and Valpola, 2017; Laine and Aila, 2016), which utilized historical models as the regularization of the current network parameters. Therefore, we would expect more stable and competitive predictions. A moving-

average model was used as the smoothed teacher model (Tarvainen and Valpola, 2017) to regularize the UDA training (French et al., 2017; Zheng and Yang, 2019; Luo et al., 2021). However, that work relied on the data in both domains for adaptation. Instead, in this work, we share a similar idea of the memory mechanism for efficiently stabilizing the OS adaptation.

Self-training

Self-training has been proposed to address semi-supervised learning (Triguero et al., 2015). Deep self-training was further proposed to integrate deep embedding learning and classifier adaptation (Zou et al., 2019). Recently, several deep self-training methods were proposed to utilize the pseudo label in the target domain data for UDA (Busto et al., 2018; Zou et al., 2019; Liu et al., 2020, 2021f). However, the aforementioned works require co-training on source domain data, which is not applicable to our source-relaxed setting. Note that the conventional pseudo label generation can be highly unstable and unreliable, since it relies on the source domain supervision for correction (Zou et al., 2019; Liu et al., 2021c). Accordingly, in the present work, we used the self-training strategy with historical consistency-guided target data only learning.

3. Methodology

Let $\mathbf{x} \in \mathbb{R}^{H_0 \times W_0 \times C_0}$ be an input image with the height, width, and channel of H_0 , W_0 , and C_0 , respectively. In segmentation, we predict the label of each pixel $\mathbf{y}_{h,w} \in \{1, 2, \dots, N\}$ as one of N classes, yielding a segmentation map $\mathbf{y} \in \mathbb{R}^{H_0 \times W_0}$. There are a source domain \mathcal{D}^s and a target domain \mathcal{D}^t , where $\mathcal{D}^s \neq \mathcal{D}^t$ indicates that their inherited data distributions are different (Kouw, 2018; Liu et al., 2021a).

We assume that a segmentation model with BN, i.e., $f: \mathbf{x} \rightarrow \mathbf{y}$ parameterized with \mathbf{w} , is pre-trained with source domain samples $(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}^s$. Note that the BN statistics are stored in the network itself, following the typical BN protocol. At the adaptation stage, we adapt the trained OS segmentor with only the unlabeled target domain samples $\mathbf{x} \sim \mathcal{D}^t$.

The domain adaptation theory (Ben-David et al., 2010) states that, for a hypothesis h drawn from \mathcal{H} , the following condition is met: $\epsilon^t(h) \leq \epsilon^s(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}\{s, t\} + e$, where, $\epsilon^s(h)$ and $\epsilon^t(h)$ represent the expected losses with hypothesis h in source and target domains, respectively. The right side terms are considered an upper bound of the target loss. Of note, e is usually a small and negligible value (Ben-David et al., 2007), and therefore UDA attempts to minimize the cross-domain divergence, $d_{\mathcal{H}\Delta\mathcal{H}}\{s, t\}$. Because the batch statistics are closely associated with the domain characteristics (Li et al., 2016; Pan et al., 2018), we propose to leverage the transferability-aware batch-wise statistics for domain alignment, thereby offering a powerful inductive bias for target domain learning (in Subsec. 3.2). Specifically, in this work, the connection of the transferability between low and high-order statistics is investigated. In addition, the SE minimization of the target data prediction (in Subsec. 3.3) and the novel memory-consistence self-training (MCSF) strategy (in Subsec. 3.4) are further integrated as a unified framework (in Subsec. 3.5).

3.1. Revisiting the Batch Normalization

BN has been widely used in modern deep learning models (Ioffe and Szegedy, 2015). For a batch with B images, which have the height, width, and channel in the l -th layer of H_l , W_l , and C_l , the batch of input features in the l -th layer $f_l \in \mathbb{R}^{B \times H_l \times W_l \times C_l}$ is normalized for each channel.

We index the samples in a batch with $b \in \{1, \dots, B\}$, index the spatial position in the l -th layer with $m \in \{1, \dots, H_l \times W_l\}$, and index the channel in the l -th layer with $c \in \{1, \dots, C_l\}$, respectively. The channel-wise mean in BN is calculated by $\mu_{l,c} = \frac{1}{B \times H_l \times W_l} \sum_b \sum_m^{H_l \times W_l} f_{l,b,m,c}$, where $f_{l,b,m,c} \in \mathbb{R}$ is the feature value. Then, the channel-wise variance can be $\{\sigma^2\}_{l,c} = \frac{1}{B \times H_l \times W_l} \sum_b \sum_m^{H_l \times W_l} (f_{l,b,m,c} - \mu_{l,c})^2$. Note that the input feature is then normalized as $\bar{f}_{l,b,m,c}$, followed by applying a linear mapping:

$$\bar{f}_{l,b,m,c} = \frac{f_{l,b,m,c} - \mu_{l,c}}{\sqrt{\{\sigma^2\}_{l,c} + \epsilon}} \quad (1)$$

$$\tilde{f}_{l,b,m,c} = \gamma_{l,c} \bar{f}_{l,b,m,c} + \beta_{l,c}, \quad (2)$$

where $\gamma_{l,c}$ and $\beta_{l,c}$ are learnable scaling and shifting factors. $\epsilon \in \mathbb{R}^+$ is a small scalar for numerical stability. After the normalization, $\bar{f}_{l,b,m,c}$ has zero mean and unit variance. In backpropagation, $\gamma_{l,c}$ and $\beta_{l,c}$ are updated with the gradient as the network parameters.

Instead of B samples in a training batch, the testing input is usually a single sample. To bridge this gap, BN layer stores the weighted average of BN statistics in training and uses it for testing. Specifically, we use $k \in \{1, 2, \dots, K\}$ to index the training iteration, and the mean and variance in each iteration are tracked progressively following:

$$\bar{\mu}_{l,c}^k = (1 - \eta) \cdot \bar{\mu}_{l,c}^{k-1} + \eta \cdot \mu_{l,c}^k, \quad (3)$$

$$\{\bar{\sigma}^2\}_{l,c}^k = (1 - \eta) \cdot \{\bar{\sigma}^2\}_{l,c}^{k-1} + \eta \cdot \{\sigma^2\}_{l,c}^k, \quad (4)$$

where $\eta \in [0, 1]$ is used to balance between the current and historical values. After K training iteration, $\bar{\mu}_{l,c}^K$, $\{\bar{\sigma}^2\}_{l,c}^K$, $\gamma_{l,c}^K$, and $\beta_{l,c}^K$ are stored and used for the normalization in testing (Ioffe and Szegedy, 2015).

3.2. Adaptive Source-relaxed Batch-wise Statistics Adaptation

We propose to explore both the shared and domain-specific batch-normalization statistics with the quantified transferability to achieve the domain alignment (Liu et al., 2021g). Specifically, we propose to reduce the domain divergence using an adaptive low-order BN statistics progression with EMD, and explicitly enforce the consistency of the high-order BN statistics in a source-free manner. In addition, the role of the scaling factor in transferability is further investigated.

3.2.1. Low-order statistics progression with EMD—For the target domain-specific factors, including mean and variance (Chang et al., 2019; Maria Carlucci et al., 2017), we propose a gradually learning scheme with EMD for low-order batch statistics progression in source-free UDA. First, the target domain mean and variance are initialized using the tracked $\bar{\mu}_{l,c}^K$ and $\{\sigma^2\}_{l,c}^K$ in the source domain. In what follows, the target domain mean and variance in t -th adaptation iteration are progressively updated as

$$\bar{\mu}_{l,c}^t = (1 - \eta^t) \cdot \mu_{l,c}^t + \eta^t \cdot \bar{\mu}_{l,c}^K, \quad (5)$$

$$\{\bar{\sigma}^2\}_{l,c}^t = (1 - \eta^t) \cdot \{\sigma^2\}_{l,c}^t + \eta^t \cdot \{\bar{\sigma}^2\}_{l,c}^K, \quad (6)$$

where $\eta^t = \eta^0 \exp(-t)$ is a momentum parameter with an exponential decay over iteration t . Note that $\mu_{l,c}^t$ and $\{\sigma^2\}_{l,c}^t$ are the mean and variance in the current target batch. Thus, the proportion of source statistics, i.e., $\bar{\mu}_{l,c}^K$ and $\{\bar{\sigma}^2\}_{l,c}^K$, are smoothly decreased along with the training, while $\mu_{l,c}^t$ and $\{\sigma^2\}_{l,c}^t$ gradually represent the low-order BN statistics in the target domain.

3.2.2. Transferability adaptive high-order statistics consistency—For the domain shareable high-order BN statistics, including the learned scaling and shifting factors (Maria Carlucci et al., 2017; Wang et al., 2019a), we explicitly enforce their consistency across two domains in source-free UDA with a high-order BN statistics (HBS) loss, given by:

$$\mathcal{L}_{HBS} = \sum_l^L \sum_c^{C_l} (1 + \alpha_{l,c}) \{ |\gamma_{l,c}^K - \gamma_{l,c}^t| + |\beta_{l,c}^K - \beta_{l,c}^t| \}, \quad (7)$$

where $\gamma_{l,c}^t$ and $\beta_{l,c}^t$ are the scaling and shifting factors in t -th adaptation iteration. We note that $\gamma_{l,c}^K$ and $\beta_{l,c}^K$ are stored in the source domain model. $\alpha_{l,c}$ is used to adaptively balance among the channels.

The transferability of each channel can be different. For instance, Pan et al. (2018) demonstrated that channels with smaller low-order BN statistics divergence can be more transferable. As such, we propose a novel loss in a way that the channel with higher transferability contributes more to UDA. To this end, in order to quantify the channel-wise domain discrepancy, the difference between batch statistics is measured as an efficient surrogate. For source-free UDA, we define a novel channel-wise cross-domain low-order BN statistics divergence at t -th adaptation iteration as

$$d_{l,c} = \left| \frac{\bar{\mu}_{l,c}^K}{\sqrt{\{\bar{\sigma}^2\}_{l,c}^K + \epsilon}} - \frac{\mu_{l,c}^t}{\sqrt{\{\sigma^2\}_{l,c}^t + \epsilon}} \right|. \quad (8)$$

Then, the channel-wise low-order statistics divergence based transferability is quantified with $\alpha_{l,c} = \frac{L \times C \times (1 + d_{l,c})^{-1}}{\sum_l \sum_c (1 + d_{l,c})^{-1}}$. Therefore, the more transferable channels will have larger

weight $(1 + \alpha_{l,c})$ in \mathcal{L}_{HBS} , thereby contributing to the optimization with a higher importance.

3.2.3. Scaling factor for channel-wise transferability—In addition to the mean and variance investigated in our prior work (Liu et al., 2021g), we further explore the scaling factor $\gamma_{l,c}^K$ for quantifying the transferability. Conventionally, the scaling factor has been widely used as a criterion for the channel-wise importance in the channel pruning¹ operations (Liu et al., 2017; Kang and Han, 2020). The small scaling factor usually indicates the less effectiveness of this channel in a single domain classification task (Molchanov et al., 2019).

Although these pruning works only focus on a single domain, we propose to investigate whether the transferability of each channel is also associated with its scaling factor. Specifically, we use the segmentation U-Net trained on high-grade gliomas (HGG) datasets and adapt it for low-grade gliomas (LGG) datasets in the BraTS2018 database (Menze et al., 2014) with our framework. \mathcal{A} -distance is a typical measurement for the feature discrepancy between two domains (Ben-David et al., 2010), which has a smaller value for better alignment and 0 for the same feature distribution. From Fig. 3(a), we can see that the pruning of 10% of the channels with the smallest $\gamma_{l,c}^K$ has the similar \mathcal{A} -distance as the network adapted without pruning, which indicates that these channels with small $\gamma_{l,c}^K$ may have little impact on UDA. This phenomenon can be simply explained from a gradient perspective. We denote all of the optimization objectives at the adaptation stage as \mathcal{L} . Then, the partial derivative of \mathcal{L} w.r.t. $\bar{f}_{l,b,m,c}$ can be expressed as:

$$\frac{\partial \mathcal{L}}{\partial \bar{f}_{l,b,m,c}} = \frac{\partial \mathcal{L}}{\partial \bar{f}_{l,b,m,c}^K} \gamma_{l,c}^K \quad (9)$$

which has a negligible gradient of $\partial \mathcal{L} / \partial \bar{f}_{l,b,m,c}$, if $\gamma_{l,c}^K$ is approaching to 0. Thus, the adaptation loss cannot efficiently enforce the domain alignment for this channel at the adaptation stage.

We note that the channel dropping can be an efficient binary hard-weighting strategy to remove the less transferable channels, while it is difficult to define the reasonable threshold, thus rendering the network less expressive. Therefore, it can lead to a suboptimal solution. As shown in Fig. 3(b), simply dropping 10% of the channels will still lead to a performance drop in the target domain. Instead of simply dropping the channels with small $\gamma_{l,c}^K$, similar to prior work, we propose to leverage the scaling factor measured reliability via a soft-weighting strategy. Specifically, we utilize the scaling factor adjusted HBS loss:

$$\mathcal{L}_{HBS}^v = \sum_l^L \sum_c^{C_l} \exp(-\gamma_{l,c}^K) (1 + \alpha_{l,c}) \{ |\gamma_{l,c}^K - \gamma_{l,c}^t| + |\beta_{l,c}^K - \beta_{l,c}^t| \}, \quad (10)$$

where $\exp(-\gamma_{l,c}^K)$ has a smaller weight, if $\gamma_{l,c}^K$ is smaller. Thus, \mathcal{L}_{HBS}^v can take both low- and high-order BN statistics to achieve channel-wise transferability quantification for adaptive BN-based adaptation.

¹ <https://intellabs.github.io/distiller/tutorial-struct-pruning.html>

3.3. Self-entropy minimization in target domain

Although the label supervision in the target domain is not available, the unlabeled target domain can be guided by an unsupervised learning objective (Bateson et al., 2020; Liu et al., 2022).

A possible solution would be self-entropy (SE) minimization which has been a widely used objective in several deep learning models to enforce the confident prediction, i.e., the maximum softmax value can be high (Grandvalet and Bengio, 2005; Liang et al., 2020; Wang et al., 2021; Bateson et al., 2020). The pixel-wise SE for segmentation is formulated by the averaged entropy of the predictions of each pixel, given by

$$\mathcal{L}_{SE} = -\frac{1}{B \times H_0 \times W_0} \sum_b^B \sum_m^{H_0 \times W_0} \{p_{b,m} \log p_{b,m}\}, \quad (11)$$

where H_0 and W_0 denote the height and width of the input, and $p_{b,m} \in \mathbb{R}^N$ indicates the N -class softmax output of the m -th pixel of the b -th image in a batch. Optimizing \mathcal{L}_{SE} can encourage $p_{b,m}$ to approach to an one-hot vector. In addition, there are more alternative objectives to encourage confident predictions, e.g., the minimum class confusion (MCC) loss (Jin et al., 2020). We note that the MCC loss takes much more sophisticated computation than the classical SE loss, and is not scalable to the segmentation task (Jin et al., 2020).

However, either SE loss or MCC loss only takes the prediction at the current iteration into consideration, which can be highly unreliable and unstable along with the training (Zou et al., 2019). In the case of source-free UDA, since there is no source data supervision in each iteration for correction, the biased prediction could significantly mislead the training.

3.4. Queued memory-consistent Self-training

We further propose a novel queued memory-consistent self-training (MCSF) strategy for stable and efficient source-relaxed UDA. MCSF is able to stabilize the OSUDA from two perspectives. First, we only account for the pixel with high confident prediction in each iteration akin to conventional self-training (Zou et al., 2019). In addition, we calculate the supervision signal, conditioned on the historical consistency.

Specifically, for a pixel in a target domain sample $\mathbf{x}_{b,m}$, we have the corresponding network prediction $p_{b,m}^t$ in t -th training iteration at the adaptation stage. The pseudo label $\hat{\mathbf{y}}_{b,m} = (\hat{y}_{b,m}^1, \dots, \hat{y}_{b,m}^N)$ of a target sample $\mathbf{x}_{b,m}$ is a N -dimensional vector indexed by n . The n -th dimension has the value of 1, only if the histogram distribution $p_{b,m}$ takes the maximum probability in n -th class, and the corresponding probability is larger than a class-wise threshold λ^n (Zou et al., 2019). We note that λ^n works as a confidence threshold to only keep the relatively reliable pseudo labels. We usually resort to the maximum value of $p_{b,m}$ as a surrogate of the confidence (Zou et al., 2019), and rank all of the pixels in a batch w.r.t. their maximum value of $p_{b,m}$. Then, λ^n is set to select the top $\alpha\%$ of the most confident pixels in each class. Specifically, each class-wise bin of the pseudo label $\hat{\mathbf{y}}_{b,m}$ can be formulated as:

$$\hat{\mathbf{y}}_{b,m}^n = \begin{cases} 1, & \text{if } n = \arg \max_n \left\{ \frac{p_{b,m}(n | \mathbf{x}, \mathbf{w})}{\lambda^n} \right\} \text{ and } p_{b,m}(n | \mathbf{x}, \mathbf{w}) > \lambda^n, \\ 0, & \text{otherwise,} \end{cases} \quad (12)$$

where $p_{b,m}(n | \mathbf{x}, \mathbf{w})$ is the value of $p_{b,m}$ in the n -th bin, which indicates the predicted probability of n -th class. Therefore, $\hat{\mathbf{y}}_{b,m}^n$ can be a one-hot histogram for the reliable pixels, while $\hat{\mathbf{y}}_{b,m}^n$ can be 0 vector for the ones that are not confident. That is, the feasible set of $\hat{\mathbf{y}}_{b,m} \in \Delta^{N-1} \cup \{\mathbf{0}\}$ is a union of probability simplex Δ^{N-1} and $\{\mathbf{0}\}$. We note that the cross-entropy loss for $\hat{\mathbf{y}}_{b,m} = \{\mathbf{0}\}$ is always 0, which indicates that the corresponding pixel is not counted to the final loss.

Then, for each $\mathbf{x}_{b,m}^\tau$, we measure the pixel-wise historical consistency $\psi_{b,m}^\tau$ over H consecutive iteration. Note that traversing all of the historical data can be inefficient to process.

Therefore, we propose to incorporate the most recent H iteration with a dynamic queue module that evolves smoothly alongside the network update. As shown in Fig. 4, we utilize a ‘‘first-in-first-out’’ memory queue to store $p_{b,m}^\tau$ and its previous values, i.e., $p_{b,m}^{\tau-1}, \dots, p_{b,m}^{\tau-H}$, to calculate the pixel-wise historical consistency $\psi_{b,m}^\tau$ at τ -th iteration by:

$$\psi_{b,m}^\tau = 1 - \text{Sigmoid} \left(\frac{1}{H} \sum_{h=1}^H \|p_{b,m}^\tau - p_{b,m}^{\tau-h}\|_1 \right). \quad (13)$$

We propose to utilize $\psi_{b,m}^\tau$ to achieve pixel-wise adaptive reweighting of the self-training objective. In the case of a good agreement of the historical predictions, the target sample can be considered a well-learned one. Therefore, it is reasonable to rely on these samples and increase their contribution to the overall loss function, by assigning a large $\psi_{b,m}^\tau$. By contrast, we lower the weights for the largely historical inconsistent samples in the loss calculations. The training loss following the cross-entropy-based self-training can be formulated as:

$$\mathcal{L}_{MCST} = - \frac{1}{B \times H_0 \times W_0} \sum_b^B \sum_m^{H_0 \times W_0} \psi_{b,m}^\tau \times \hat{\mathbf{y}}_{b,m}^\tau \log p_{b,m}^\tau, \quad (14)$$

where $p_{b,m}^\tau$ is the prediction at τ -th iteration, and $\hat{\mathbf{y}}_{b,m}^\tau$ is its pseudo label generated with Eq. (12).

Compared with the self-entropy in Eq. (11), \mathcal{L}_{MCST} utilizes the pseudo label $\hat{\mathbf{y}}_{b,m}^\tau$ instead of $p_{b,m}^\tau$. Therefore, the low confident predictions will not be used to update the networks. In addition, the loss is adaptively adjusted with $\psi_{b,m}^\tau$.

3.5. Overall training protocol for OSUDA with MCSF

At the target domain adaptation stage, the OSUDA with EMD and \mathcal{L}_{HBS}^τ can be combined with either the conventional SE minimization or our novel queued memory-consistent self-training. Note that the training protocol of the source model, e.g., loss function and

hyperparameters, is not a prerequisite for our model, but the trained model with its stored BN statistics alone is required.

In the case of OSUDA with SE minimization, the overall training objective can be formulated as $\mathcal{L} = \mathcal{L}'_{HBS} + \lambda \mathcal{L}_{SE}$, where λ balances between these two terms. However, SE may lead to a trivial solution that every sample has the same one-hot prediction (Grandvalet and Bengio, 2005). For stabilization, we propose to change the contribution of the loss terms along with the training (Granger et al., 2020; Ganin et al., 2016; Tang et al., 2022). Specifically, we simply linearly decrease λ , e.g., from 10 to 0, along with the adaptation. Of note, the use of more sophisticated decreasing functions, e.g., the log or exponential function w.r.t. epoch (Granger et al., 2020; Ganin et al., 2016) could further enhance adaptation performance.

The optimization objective of the proposed OSUDA with MCSF can be expressed as

$$\begin{aligned} \min_{\mathbf{w}, \hat{\mathbf{y}}} \mathcal{L} &= \mathcal{L}'_{HBS} + \lambda \mathcal{L}_{SE} + \varphi \mathcal{L}_{MCSF}, \\ \text{s.t. } \hat{\mathbf{y}} &\in \Delta^{K-1} \cup \{\mathbf{0}\}, \forall \mathbf{x} \in \mathcal{D}', \text{ and EMD,} \end{aligned} \quad (15)$$

which is typically formulated as a classification maximum likelihood (CML) problem (Amini and Gallinari, 2002), and can be optimized with Classification Expectation Maximization (CEM) (Zou et al., 2019). Specifically, there are three steps in each iteration (Zou et al., 2019):

1. **Expectation:** Estimating $p_{b,m}^f$ for all of $\mathbf{x}_{b,m}$ in a target batch with the forward pass of the current model.
2. **Classification:** Calculating the pseudo label $\hat{\mathbf{y}}_{b,m}^f$ (in Eq. (12)) and the corresponding memory-consistency $\psi_{b,m}^f$ (in Eq. (13)) for all of $\mathbf{x}_{b,m}$ in a target batch.
3. **Maximization:** Updating the low-order BN statistics with EMD, and fine-tuning the network parameter \mathbf{w} with $\mathcal{L}'_{HBS} + \lambda \mathcal{L}_{SE} + \varphi \mathcal{L}_{MCSF}$ via backpropagation.

We note that solving for Eq. (12) in our Expectation-classification steps is a typical concave problem that has a globally optimal solution. In addition, the Maximization step is seen as supervised learning, which is usually convergent (Shalev-Shwartz and Ben-David, 2014; Cover, 1999). Thus, the overall training process can be convergent. In each iteration, we implement these three steps sequentially as the conventional self-training (Zou et al., 2019; Liu et al., 2021f).

4. Experiments and Results

To show the effectiveness of our framework, we experimented on both the brain tumor and cardiac segmentation tasks. We provided detailed comparisons against “partial” source-relaxed UDA methods as well as the contemporary or follow-up source-free UDA methods. The source available UDA methods with the same backbones are used as our “upper bounds.”

In addition, we provide ablation studies of the components in our framework, and the sensitivity analysis of hyperparameters involved. We denote our prior work (Liu et al., 2021g) as OSUDA. The OSUDA without the adaptive channel-wise weighting and SE minimization are denoted as OSUDA-AC and OSUDA-SE, respectively. In addition, OSUDA+ \mathcal{L}_{HBS}^s indicates using the scaling factor adjusted HBS loss for the high-order BN statistics consistency. We further integrate the memory-consistent self-training as our memory consistent OSUDA (MCOSUDA), i.e., OSUDA+ \mathcal{L}_{HBS}^s + MCSF, to achieve superior source-free adaptation performance.

We set $\epsilon = 1 \times 10^{-6}$ for batch normalization as in Ioffe and Szegedy (2015). Following the previous self-training works (Zou et al., 2019; Liu et al., 2020, 2021f), we empirically initialize $\alpha = 20$, and linearly increase it to 80 along with the training, since the pseudo-label is inherently more noisy at the start of training.

The training was performed with the PyTorch deep learning toolbox (Paszke et al., 2017) on an NVIDIA V100 GPU. For the evaluation metrics, we employed the widely accepted Dice similarity coefficient (DSC) and Hausdorff distance (HD) as in Zou et al. (2020); Bateson et al. (2020). The DSC, a.k.a. dice score (the higher, the better), measures the overlap between the predicted segmentation mask and the label. The HD (the lower, the better) is defined for two sets of points in the prediction and label in a metric space (Zou et al., 2020; Bateson et al., 2020). The standard deviation was reported over five runs.

4.1. Brain Tumor Segmentation

The BraTS2018 database contains a total of 285 patients (Menze et al., 2014), in which a total of 210 patients have glioblastoma, i.e., high-grade gliomas (HGG), while the remaining 75 patients have low-grade gliomas (LGG).

Each patient has 4 registered Magnetic Resonance (MR) Imaging (MRI) modalities, i.e., T1-weighted (T1), T1-contrast enhanced (T1ce), T2-weighted (T2), and T2 Fluid Attenuated Inversion Recovery (FLAIR) MRI. The MRI voxels have four class labels, i.e., enhancing tumor (EnhT), peritumoral edema (ED), tumor core (CoreT), and background. The union of CoreT, ED, and EnhT represents the whole tumor (Shanis et al., 2019). To demonstrate the validity and generality of the proposed OSUDA, we carried out two cross-domain protocols, including cross-modality (Zou et al., 2020) and cross-subtype UDA (Shanis et al., 2019). Fig. 5 shows example samples with the four MRI modalities from HGG or LGG datasets. Of note, because there are imaging artifacts and some volumes are of low resolution (e.g., LGG T1 slices), some of the structures are incomplete. In addition, volumes in the BraTS database have different resolution, and so co-registration was carried out (Menze et al., 2014), which also involves interpolation to a standard resolution.

4.1.1. Cross-MR-modality UDA: T2-weighted MRI to T1-weighted/T1ce/FLAIR MRI

—In the cross-modality UDA task, there can be relatively large appearance discrepancies among the MRI modalities. Considering the clinical manual annotation of the brain tumor usually works on T2-weighted MRI, the conventional cross-MR-modality UDA focuses on using T2-weighted MRI as the labeled source domain, while T1-weighted/T1ce/FLAIR MRI are used as the target domain (Zou et al., 2020). Following the standard

protocol, we used 80% subjects for training and 20% subjects for testing (Zou et al., 2020), and used the same network backbone. All of the samples were used in a subject-independent and unpaired manner (Zou et al., 2020).

We followed the backbone as in Zou et al. (2020). We set the batch size to 12 for segmentation. For all the adaptation models, a model trained on the source data with cross-entropy over 100 epochs was used as an initialization. The network was trained over 100 adaptation epochs. The consecutive iteration in the memory H was set as 5. The training took about 1 hour. In practice, segmenting one image in testing only took about 0.1 seconds.

In Table 1, we provide the quantitative evaluation results. The proposed OSUDA-based methods outperformed CRUDA (Bateson et al., 2020) and the follow-up source-free UDA methods, e.g., GTA (Kundu et al., 2021), DPL (Chen et al., 2021), and SFKT (Liu et al., 2021h). Note that CRUDA requires training an additional label proportion network in the source domain, and GTA requires training a special domain generalization module in the source domain, both of which are not source-free settings. In addition, our OSUDA-based methods with \mathcal{L}'_{HBS} and MCSF outperformed the source-available UDA methods, e.g., CycleGAN (Zhu et al., 2017) and SIFA (Chen et al., 2019), for T2-weighted MRI to T1-weighted/FLAIR MRI transfer tasks w.r.t. DSC and HD. We listed the state-of-the-art source available UDA methods as our “upper bounds”, and did not manage to beat all of them. Fig. 6 illustrates the visual comparisons of 3 target MRI modalities, which shows the promising results of OSUDA-based methods over the other comparison methods.

OSUDA-AC and OSUDA-SE yielded inferior performance than OSUDA, demonstrating the effectiveness of the adaptive channel-wise weighting and the SE minimization. In addition, the scaling factor adjusted \mathcal{L}'_{HBS} and MCSF can further improve the performance.

4.1.2. Cross-subtype UDA: HGG to LGG—The subtypes of HGG and LGG can take different tumor sizes and positions. Following the standard cross-subtype data split protocol in Shanis et al. (2019) to adapt the model trained on HGG to LGG subjects. We chose the same backbone as in Shanis et al. (2019) with 15 layers in U-Net (Ronneberger et al., 2015). Networks were trained with four-channel sliced 2D axial MRI slices to perform pixel-wise four-class segmentation, including background, EnhT, CoreT, and ED. The input samples have the size of $128 \times 128 \times 4$, which is a slice-wise concatenation of four MRI modalities. The networks were trained with a batch size of 12. The consecutive iteration in the memory H was set as 3. Following Shanis et al. (2019), both HGG and LGG volumes were split into training and testing datasets. We pre-trained our network on the source domain data over 150 epochs as in Shanis et al. (2019). The adaptation training took about 2 hours.

In Table 2, we provide the results of our quantitative evaluations. Due to different class-wise pixel proportions in the two subtypes, the class-ratio-based CRUDA (Bateson et al., 2020) only achieved marginal improvements. Especially, the DSC of CoreT was inferior to the source model without adaptation, which can be the case of the negative transfer (Wang et al., 2019b). Our proposed OSUDA with \mathcal{L}'_{HBS} and MCSF was able to achieve superior performance for source-free UDA segmentation, and approached the results of source-

available SEAT (Shanis et al., 2019). The ablation studies also confirmed the effectiveness of the adaptive channel-wise weighting, SE minimization, \mathcal{L}_{HBS}^i , and MCSF.

The qualitative comparisons are shown in Fig. 7. Source-free UDA methods were able to substantially improve the performance over the source domain model without adaptation. However, CRUDA (Bateson et al., 2020) tended to predict a larger area for the tumor.

4.2. Segmentation of Cardiac Structures

We used the MM-WHS database for the whole heart segmentation (Zhuang and Shen, 2016), which contains a total of 40 datasets acquired from multiple clinical sites. There are a total of 20 subjects who have MRI scans, each of which has a total of 128 MRI slices. Another 20 subjects have CT scans, each of which has about a total of 256 slices. The segmentation label of each slice in the cardiac datasets has five classes, including left ventricle blood cavity (LVC), left atrium blood cavity (LAC), the myocardium of the left ventricle (MYO), ascending aorta (AA), and background. Note that the subjects with MRI and CT scans are unpaired. We followed the previous evaluation protocols in Chen et al. (2019); Zou et al. (2020); Chanti and Mateus (2021) to use the coronal view slices, and chose 8/2 subject split for training and testing, respectively. We chose MRI and CT as the source and target domains, respectively. As a preprocessing step, the slice was cropped to 256×256 . For a fair comparison, we adopted the same backbone in Chen et al. (2019); Zou et al. (2020). We set $\varphi = 5$ and linearly decreased λ from 10 to 0. In all of our cardiac segmentation tasks, we set the batch size B as 12, and the consecutive iteration in the memory H as 5. We trained the source model over 100 epochs, followed by the source-free adaptation over 100 epochs. The training took about 1.5 hours.

The numerical comparisons are provided in Table 3. Consistent with the brain tumor segmentation tasks, our framework performed better than CRUDA (Bateson et al., 2020) and SFKT (Liu et al., 2021h), by more than 4% w.r.t. DSC, and 6mm w.r.t. HD. When compared with the methods with the source data, our framework performed better than the classical cycleGAN (Zhu et al., 2017). In Fig. 8, we provide the qualitative comparisons.

In addition, we followed the evaluation protocol in Chanti and Mateus (2021) to evaluate the performance with only one subject in the target domain. As in Chanti and Mateus (2021), we took three adjacent slices ($256 \times 256 \times 3$) as input and predicted the segmentation mask of the middle slice to explore the 2.5D information. The segmentation network has 5 convolutional and 5 de-convolutional layers (Chanti and Mateus, 2021). We set $H = 3$ and $\varphi = 5$, and linearly decreased λ from 10 to 0.

The numerical comparisons are provided in Table 4. We can see that the performance improvements over CRUDA (Bateson et al., 2020) and SFKT (Liu et al., 2021h) are more appealing. Specifically, the average DSC of MCOSUDA was about 15% higher than CRUDA and SFKT.

4.2. Sensitive analysis of hyper-parameters

There are several hyper-parameters involved in our framework. In this subsection, we provide a systematical analysis of these hyper-parameters. Firstly, the stability of the BN

statistics can be highly associated with the batch size. A larger batch size can provide more unbiased statistics, while introducing more computation and memory costs. As shown in Fig. 9 left, the performance can be benefited, by increasing the batch size from 1 to 10, while the performance is almost stable for the batch size larger than 10 for all of the tasks. Therefore, we simply set all of the batch size to 12 in this work.

We used λ to weight the SE minimization, which is linearly changed from a relatively large positive value to 0 at the adaptation stage. As shown in Table 5, we can see that the linearly decreasing λ can be better than fixed λ . To select the proper initial value of λ , we provide the sensitivity analysis in Fig. 10. The range of linearly decreasing λ can be relatively stable from $5 \rightarrow 0$ to $15 \rightarrow 0$. We simply set λ to $10 \rightarrow 0$ for all of our experiments. We note that setting λ to 0 is equivalent to OSUDA-SE, i.e., without the SE minimization term.

In MCSF, we utilized a memory queue to measure the historical consistency. As shown in Fig. 9 right, we analyzed the relationship between H and the performance. In both the cross-modality segmentation and the cardiac MR to CT segmentation tasks, the performances were stable, since H had a value larger than 5. For the HGG to LGG task, $H = 3$ was sufficient to achieve the DSC of 62.87%.

In order to balance the MCSF loss, we used φ to weight the objectives. For the sensitivity analysis, we provide the comparisons with different φ as shown in Fig. 11. By setting φ larger than 5, we were able to achieve relatively stable performance for all of the tasks, e.g., the average DSC of 59.17% in cross-modality whole tumor segmentation using the BraTS2018 dataset.

In addition, λ^n is set to select the top α % of the most confident pixels in each class. Therefore, α is a to-be tuned hyperparameter, and λ^n adaptively changes in each iteration based on α . Increasing α along with the training is a typical solution to accommodate the noisy level change of pseudo-label (Zou et al., 2019; Liu et al., 2020, 2021f). In Table 6, we compared the model with different start and end α for both tasks. We can see that too large end α can lead to worse performance. Since the pseudo-label can still be noisy at the late training epochs, assuming more than 80% of them are correct can mislead the training. In addition, too small start α can lead to slower convergence, since the pseudo label is not sufficiently utilized.

5. Discussion

The problem of domain shift is prevalent, when applying deep learning models trained on source domain data to carry out a variety of tasks in the target domain. As a result, the performance degradation has been clearly observed in tasks using data from different centers (Liu et al., 2021e), scanners (Ghafoorian et al., 2017), populations, subtypes (Liu et al., 2021e), and modalities (Liu et al., 2021f). Among these cases, cross-modality UDA can be the most challenging task, due to different imaging principles involved, resulting in different image appearances. The general validity and efficacy of our proposed framework were demonstrated using cross-MR-modality and cross-subtype brain tumor segmentation tasks, and cardiac MR to CT segmentation task. While these databases presented different

challenges, our framework robustly achieved superior performance consistently compared with the conventional approaches. In addition, different backbones were adopted, similar to the previous works. Notably, our proposed framework is not dependent on prior knowledge of the specific imaging modality, and can be easily applied to other UDA tasks.

Due to concerns over patient data privacy and IP, the restriction of the source data sharing in clinical practice can be a significant obstacle for many UDA approaches. To address the issue, in this work, we proposed a novel UDA segmentation framework in the absence of source domain samples. To the best of our knowledge, this is one of the first attempts at source-relaxed UDA for image segmentation, which does not need an auxiliary network, or the unreliable assumption of the same label proportion (Bateson et al., 2020). Our framework is only reliant on a pre-trained OS segmentor, with the widely used BN.

As shown in Tables 2–4 and Figs. 6–8, our framework yields superior performance over the other comparison methods, when qualitatively and quantitatively evaluated. There are a few important differences that give insights into their performances. We can see that the performance degradation of the source domain model widely exists, especially when we apply our framework to a different target domain. Due to a relatively large domain shift in the cardiac MR to CT segmentation task, the DSC of the source domain model was only 17.2%. Since the multi-modality segmentation task is widely used in clinical practice, and the dual annotation on multiple modalities can be a large burden, the domain adaptation methods can be a viable solution (Chen et al., 2019).

CRUDA (Bateson et al., 2020) assumes that the pixel proportion is consistent between two domains. In adaptation, this prior knowledge was used as the only transferable information. In Table 2, we show that the performance of CRUDA is largely inferior to our framework, since the tumor pixel proportion is largely different between HGG and LGG subjects. In addition, the class-ratio prediction model is usually not “off-the-shelf,” since it is not typically used in clinical practice and needs specialized training in the source domain. Recently, there are several contemporary or follow-up works, aimed at source-free UDA for natural image segmentation, which are re-implemented and compared in Tables 1, 2, and 3. Although the relatively good performance was achieved by Kundu et al. (2021), its domain generalization training in the source domain limits the model to be used for source-free UDA. The additional attention network used in SFKT may require relatively more data for training, thus yielding additional parameters. For the cardiac MR to CT segmentation task, with only one target training sample, the DSC of SFKT is 7% lower than our framework. The denoising strategy used in Chen et al. (2021) utilized the Monte Carlo method, which is challenging for Bayesian uncertainty approximation and training (Liu et al., 2021f).

Several state-of-the-art source-available UDA segmentation works are also compared (Liu et al., 2021a; Zou et al., 2020), which are used as the “upper bounds.” Note that we did not manage to outperform all of them, since we have different settings when training on the source data, and it can be applied to different applications. Nonetheless, the fact that our framework approached their performance indicates the superior performance of our source-free UDA. In some of the tasks, our framework was able to outperform even the

source data available methods (Zhu et al., 2017; Shanis et al., 2019; Chen et al., 2019), which further demonstrates the superiority of our framework.

We proposed to explore the BN statistics for source-free UDA using a systematical framework with low-order statistics progression and high-order statistics consistency loss. The transferability of each channel was exploited for the channel-wise adaptive alignment. We further proposed to utilize both the low-order statistics discrepancy and scaling factor for the transferability quantification, which are evidenced by the ablation studies for OSUDA-AC and OSUDA+ $\mathcal{L}_{HBS}^{\prime}$, respectively. In addition to the BN statistics alignment, we further integrated the SE minimization and MCSF into a unified framework to exploit the unsupervised learning and memory-based learning, respectively. The unsupervised SE minimization and pseudo label used in self-training can be noisy and unstable (Zou et al., 2019; Liu et al., 2021f, 2020). In our source-free setting, no source domain data are available at the adaptation stage to correct the biased update similar to the source available self-training (Zou et al., 2019), which makes the task more challenging. To address this issue, we proposed a novel MCSF strategy to efficiently regularize the adaptation with memory-based learning. In Fig. 12, we show that MCSF was able to effectively stabilize the training, and achieve better performance.

This work aimed to develop a source-free UDA approach, in which source domain data are not available at the adaptation stage, which is considered a more restricted setting, compared with UDA with source data. This setting, in turn, could help avoid any potential issues over patient data sharing or IP. As such, the evaluations against UDA approaches, in which source data are available, are used for performance comparison purposes only.

There are a few important aspects that are not fully studied in the present work. First, we proposed to exploit the BN statistics. Recent studies (Zhou et al., 2021) used instance-wise statistics for domain generalization, which can be orthogonal to our framework and could be potentially added to our framework. Second, we only investigate the case where the source and target domain models use the same backbone, which is a typical setting in UDA. The different backbones may be achievable by initializing the target model with knowledge distillation (Hinton et al., 2015). Third, we have taken the correlation between neighboring MRI slices into consideration (i.e., 2.5D segmentation) as in Chanti and Mateus (2021) in Table 4. Yet, the 3D segmentation backbones can hardly be applied to the datasets, due to the relatively limited number of 3D datasets to provide reliable BN statistics. In addition, while our experiments showed that many of the weights were not sensitive, they needed to be carefully tuned to balance among the optimization objectives.

6. Conclusion

In this paper, we have proposed a novel and practical source-free UDA framework, aimed at image segmentation. Our framework was only reliant on an OS pre-trained segmentation network with BN in the source domain, which could thereby sidestep the concerns over the patient data sharing and IP inherent in conventional UDA. The BN statistics were systematically investigated for domain alignment. Specifically, the low-order BN statistics progression with EMD was proposed to gradually learn the target domain-specific mean

and variance. The domain shareable high-order BN statistics consistency was encouraged by the HBS loss, which was adaptively adjusted, according to the channel-wise transferability. In addition to quantifying the transferability based on low-order statistics discrepancy, the high-order scaling factor was further explored. An unsupervised learning objective, i.e., SE minimization, was incorporated into our framework, and the novel queued memory-consistent self-training was further proposed to achieve stable memory learning with the pseudo label. Extensive experiments, ablation studies as well as sensitivity analysis on the cross-MR-modality and cross-subtype brain tumor segmentation tasks and cardiac MR to CT segmentation task demonstrated the effectiveness of our OSUDA, which can be potentially applied in a clinically meaningful setting.

Acknowledgments

We gratefully acknowledge funding support from NIH R01DC018511, R01DE027989, and P41EB022544.

References

- Amini MR, Gallinari P, 2002. Semi-supervised logistic regression, in: ECAI.
- Bateson M, Kervadec H, Dolz J, Lombaert H, Ayed IB, 2020. Source-relaxed domain adaptation for image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 490–499.
- Ben-David S, Blitzer J, Crammer K, Kulesza A, Pereira F, Vaughan JW, 2010. A theory of learning from different domains. *Machine learning* 79, 151–175. doi:10.1007/s10994-009-5152-4.
- Ben-David S, Blitzer J, Crammer K, Pereira F, 2007. Analysis of representations for domain adaptation, in: NIPS.
- Busto PP, Iqbal A, Gall J, 2018. Open set domain adaptation for image and action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Chang WG, You T, Seo S, Kwak S, Han B, 2019. Domain-specific batch normalization for unsupervised domain adaptation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7354–7362.
- Chanti DA, Mateus D, 2021. Olva: Optimal latent vector alignment for unsupervised domain adaptation in medical image segmentation. *MICCAI*.
- Che T, Liu X, Li S, Ge Y, Zhang R, Xiong C, Bengio Y, 2019. Deep verifier networks: Verification of deep discriminative models with deep generative models, in: ArXiv.
- Chen C, Dou Q, Chen H, Qin J, Heng PA, 2019. Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 865–872.
- Chen C, Liu Q, Jin Y, Dou Q, Heng PA, 2021. Source-free domain adaptive fundus image segmentation with denoised pseudo-labeling, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 225–235.
- Cover TM, 1999. Elements of information theory. John Wiley & Sons.
- French G, Mackiewicz M, Fisher M, 2017. Self-ensembling for visual domain adaptation. arXiv preprint arXiv:1706.05208.
- Ganin Y, Ustinova E, Ajakan H, Germain P, Larochelle H, Laviolette F, Marchand M, Lempitsky V, 2016. Domain-adversarial training of neural networks. *JMLR*.
- Ghafoorian M, Mehrtash A, Kapur T, Karssemeijer N, Marchiori E, Pesteie M, Guttman CR, de Leeuw FE, Tempny CM, Van Ginneken B, et al. , 2017. Transfer learning for domain adaptation in mri: Application in brain lesion segmentation, in: International conference on medical image computing and computer-assisted intervention, Springer. pp. 516–524.
- Grandvalet Y, Bengio Y, 2005. Semi-supervised learning by entropy minimization, in: NIPS.

- Granger E, Kiran M, Dolz J, Blais-Morin LA, et al. , 2020. Joint progressive knowledge distillation and unsupervised domain adaptation, in: 2020 International Joint Conference on Neural Networks (IJCNN), IEEE. pp. 1–8.
- He G, Liu X, Fan F, You J, 2020. Classification-aware semi-supervised domain adaptation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 964–965.
- He K, Zhang X, Ren S, Sun J, 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
- Hinton G, Vinyals O, Dean J, 2015. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531.
- Ioffe S, Szegedy C, 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: International conference on machine learning, PMLR. pp. 448–456.
- Jin Y, Wang X, Long M, Wang J, 2020. Minimum class confusion for versatile domain adaptation, in: European Conference on Computer Vision, Springer. pp. 464–480.
- Kang M, Han B, 2020. Operation-aware soft channel pruning using differentiable masks, in: International Conference on Machine Learning, PMLR. pp. 5122–5131.
- Kouw WM, 2018. An introduction to domain adaptation and transfer learning. arXiv.
- Kundu JN, Kulkarni A, Singh A, Jampani V, Babu RV, 2021. Generalize then adapt: Source-free domain adaptive semantic segmentation, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 7046–7056.
- Laine S, Aila T, 2016. Temporal ensembling for semi-supervised learning, in: ICLR.
- Li R, Jiao Q, Cao W, Wong HS, Wu S, 2020a. Model adaptation: Unsupervised domain adaptation without source data, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9641–9650.
- Li X, Chen W, Xie D, Yang S, Yuan P, Pu S, Zhuang Y, 2020b. A free lunch for unsupervised domain adaptive object detection without source data. arXiv preprint arXiv:2012.05400.
- Li Y, Wang N, Shi J, Hou X, Liu J, 2018. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition* 80, 109–117.
- Li Y, Wang N, Shi J, Liu J, Hou X, 2016. Revisiting batch normalization for practical domain adaptation. arXiv preprint arXiv:1603.04779.
- Liang J, Hu D, Feng J, 2020. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation, in: International Conference on Machine Learning, PMLR. pp. 6028–6039.
- Liu X, Guo Z, Li S, Xing F, You J, Kuo CCJ, El Fakhri G, Woo J, 2021a. Adversarial unsupervised domain adaptation with conditional and label shift: Infer, align and iterate, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 10367–10376.
- Liu X, Hu B, Jin L, Han X, Xing F, Ouyang J, Lu J, El Fakhri G, Woo J, 2021b. Domain generalization under conditional and label shifts via variational bayesian inference, in: IJCAI.
- Liu X, Hu B, Liu X, Lu J, You J, Kong L, 2020. Energy-constrained self-training for unsupervised domain adaptation. *ICPR*.
- Liu X, Hu B, Liu X, Lu J, You J, Kong L, 2021c. Energy-constrained self-training for unsupervised domain adaptation, in: 2020 25th International Conference on Pattern Recognition (ICPR), IEEE. pp. 7515–7520.
- Liu X, Li S, Ge Y, Ye P, You J, Lu J, 2021d. Recursively conditional gaussian for ordinal unsupervised domain adaptation, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 764–773.
- Liu X, Liu X, Hu B, Ji W, Xing F, Lu J, You J, Kuo CCJ, Fakhri GE, Woo J, 2021e. Subtype-aware unsupervised domain adaptation for medical diagnosis. *AAAI*.
- Liu X, Xing F, Stone M, Zhuo J, Reese T, Prince JL, El Fakhri G, Woo J, 2021f. Generative self-training for cross-domain unsupervised tagged-to-cine mri synthesis, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 138–148.

- Liu X, Xing F, Yang C, El Fakhri G, Woo J, 2021g. Adapting off-the-shelf source segmenter for target medical image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 549–559.
- Liu X, Yoo C, Xing F, Kuo CCJ, El Fakhri G, Woo J, 2022. Unsupervised domain adaptation for segmentation with black-box source model. SPIE Medical Imaging 2022: Image Processing.
- Liu Y, Zhang W, Wang J, 2021h. Source-free domain adaptation for semantic segmentation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1215–1224.
- Liu Z, Li J, Shen Z, Huang G, Yan S, Zhang C, 2017. Learning efficient convolutional networks through network slimming, in: Proceedings of the IEEE international conference on computer vision, pp. 2736–2744.
- Long J, Shelhamer E, Darrell T, 2015. Fully convolutional networks for semantic segmentation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3431–3440.
- Luo Z, Cai Z, Zhou C, Zhang G, Zhao H, Yi S, Lu S, Li H, Zhang S, Liu Z, 2021. Unsupervised domain adaptive 3d detection with multi-level consistency, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8866–8875.
- Mancini M, Porzi L, Bulò SR, Caputo B, Ricci E, 2018. Boosting domain adaptation by discovering latent domains, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3771–3780.
- Maria Carlucci F, Porzi L, Caputo B, Ricci E, Rota Bulò S, 2017. Autodial: Automatic domain alignment layers, in: Proceedings of the IEEE International Conference on Computer Vision, pp. 5067–5075.
- Menze BH, Jakab A, Bauer S, Kalpathy-Cramer J, Farahani K, Kirby J, Burren Y, Porz N, Slotboom J, Wiest R, et al. , 2014. The multimodal brain tumor image segmentation benchmark (BRATS). IEEE transactions on medical imaging 34, 1993–2024. [PubMed: 25494501]
- Molchanov P, Mallya A, Tyree S, Frosio I, Kautz J, 2019. Importance estimation for neural network pruning, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11264–11272.
- Pan X, Luo P, Shi J, Tang X, 2018. Two at once: Enhancing learning and generalization capacities via ibn-net, in: Proceedings of the European Conference on Computer Vision (ECCV), pp. 464–479.
- Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, Lin Z, Desmaison A, Antiga L, Lerer A, 2017. Automatic differentiation in pytorch.
- Ronneberger O, Fischer P, Brox T, 2015. U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer. pp. 234–241.
- Shalev-Shwartz S, Ben-David S, 2014. Understanding machine learning: From theory to algorithms. Cambridge university press.
- Shanis Z, Gerber S, Gao M, Enquobahrie A, 2019. Intramodality domain adaptation using self ensembling and adversarial training, in: Domain Adaptation and Representation Transfer and Medical Image Learning with Less Labels and Imperfect Data. Springer, pp. 28–36.
- Tajbakhsh N, Jeyaseelan L, Li Q, Chiang JN, Wu Z, Ding X, 2020. Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. Medical Image Analysis 63, 101693. [PubMed: 32289663]
- Tang H, Wang Y, Jia K, 2022. Unsupervised domain adaptation via distilled discriminative clustering. Pattern Recognition 127, 108638.
- Tarvainen A, Valpola H, 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results, in: NIPS.
- Triguero I, García S, Herrera F, 2015. Self-labeled techniques for semi-supervised learning: taxonomy, software and empirical study. Knowledge and Information Systems 42, 245–284.
- Wang D, Shelhamer E, Liu S, Olshausen B, Darrell T, 2021. Tent: Fully test-time adaptation by entropy minimization. ICLR.
- Wang X, Jin Y, Long M, Wang J, Jordan M, 2019a. Transferable normalization: Towards improving transferability of deep neural networks. arXiv preprint arXiv:2019.

- Wang Z, Dai Z, Póczos B, Carbonell J, 2019b. Characterizing and avoiding negative transfer, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11293–11302.
- Weston J, Chopra S, Bordes A, 2014. Memory networks. arXiv preprint arXiv:1410.3916.
- Wilson G, Cook DJ, 2020. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)* 11, 1–46.
- Ye M, Zhang J, Ouyang J, Yuan D, 2021. Source data-free unsupervised domain adaptation for semantic segmentation, in: Proceedings of the 29th ACM International Conference on Multimedia, pp. 2233–2242.
- You F, Li J, Zhu L, Chen Z, Huang Z, 2021. Domain adaptive semantic segmentation without source data, in: Proceedings of the 29th ACM International Conference on Multimedia, pp. 3293–3302.
- Zhang J, Qi L, Shi Y, Gao Y, 2020. Generalizable semantic segmentation via model-agnostic learning and target-specific normalization. arXiv preprint arXiv:2003.12296.
- Zheng Z, Yang Y, 2019. Unsupervised scene adaptation with memory regularization in vivo. arXiv preprint arXiv:1912.11164.
- Zhou K, Yang Y, Qiao Y, Xiang T, 2021. Domain generalization with mixstyle. arXiv preprint arXiv:2104.02008.
- Zhou XY, Yang GZ, 2019. Normalization in training u-net for 2-D biomedical semantic segmentation. *IEEE Robotics and Automation Letters* 4, 17921799.
- Zhu JY, Park T, Isola P, Efros AA, 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks, in: ICCV.
- Zhuang X, Shen J, 2016. Multi-scale patch and multi-modality atlases for whole heart segmentation of mri. *Medical image analysis* 31, 77–87. [PubMed: 26999615]
- Zou D, Zhu Q, Yan P, 2020. Unsupervised domain adaptation with dualscheme fusion network for medical image segmentation, in: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20, International Joint Conferences on Artificial Intelligence Organization, pp. 3291–3298.
- Zou Y, Yu Z, Liu X, Kumar B, Wang J, 2019. Confidence regularized self-training, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 5982–5991.

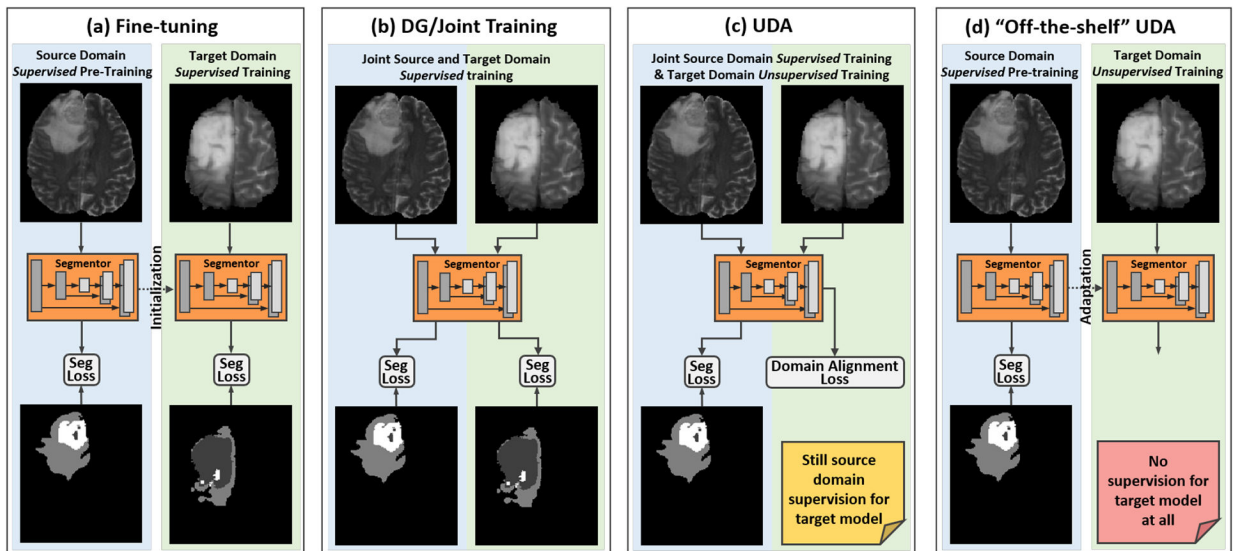


Fig. 1:

Comparisons of our framework with different related knowledge transfer methods: (a) fine-tuning makes use of labels in both domains via two stages, i.e., supervised pre-training in source domain and supervised re-training in target domain; (b) domain generalization (DG) (Liu et al., 2021b) relies on joint training and expects generalization in unseen domains; (c) conventional UDA (Wilson and Cook, 2020; Liu et al., 2021a) is trained jointly on labeled source domain and unlabeled target domain data; and (d) our source-relaxed “off-the-shelf” model adaptation for segmentation is based on adaptation, without source domain data.

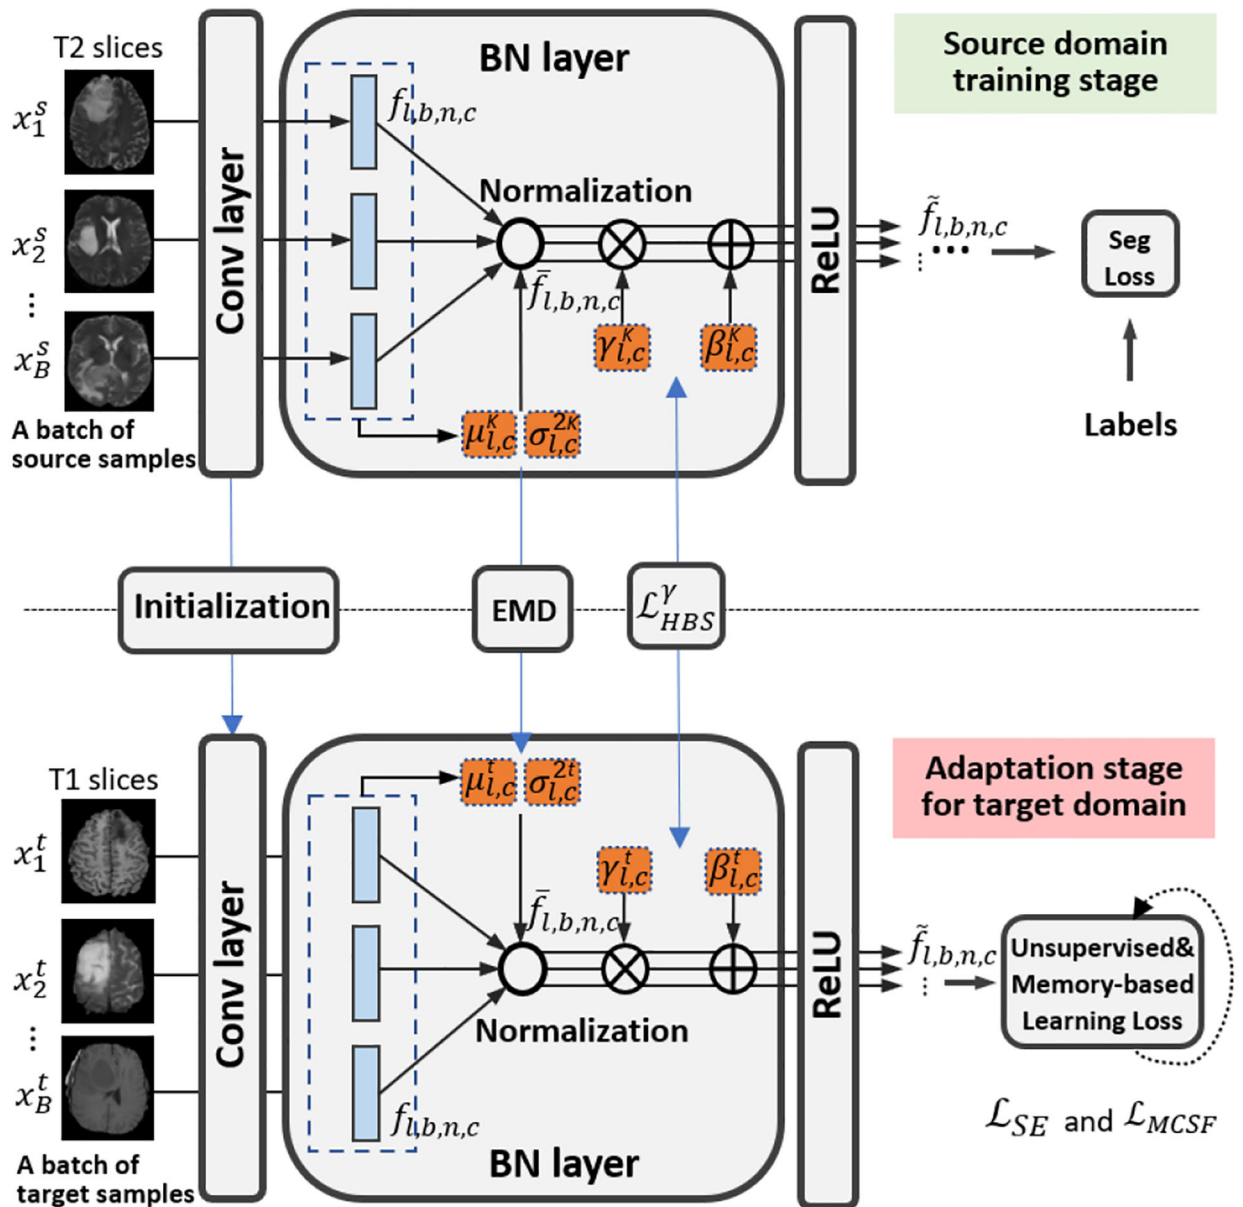


Fig. 2: Illustration of a channel in our OSUDA framework, based on the pre-trained “off-the-shelf (OS)” model with BN. We mitigate the domain discrepancy with the adaptive BN statistics in each channel.

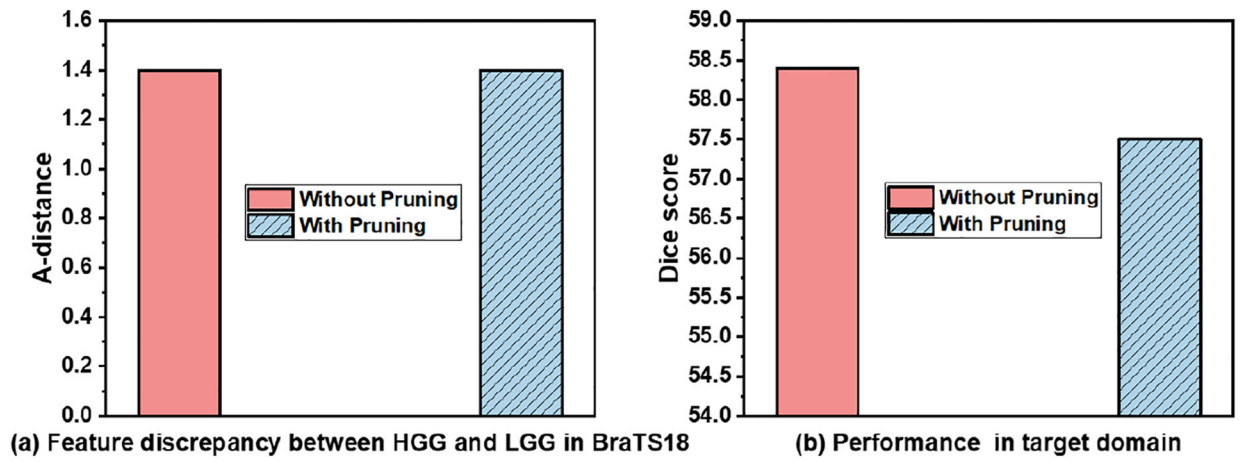


Fig. 3: Comparison of (a) \mathcal{A} -distance (Ben-David et al., 2010) and (b) the target domain DSC of the source domain model w/wo pruning of channels with a small scaling factor in the HGG to LGG task.

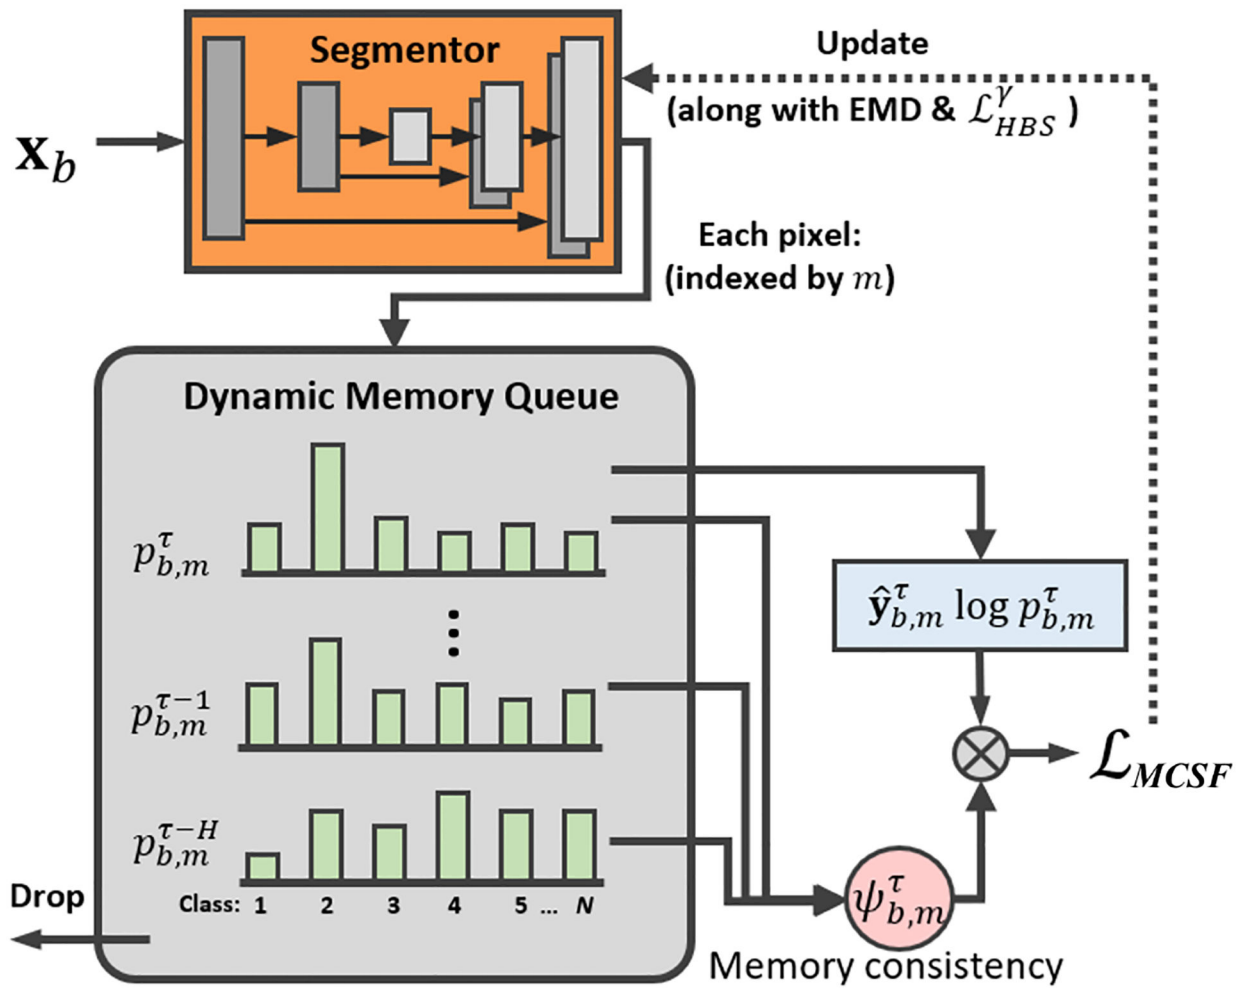


Fig. 4: Illustration of our proposed queued dynamic memory-consistent self-training strategy.




Fig. 5: Illustration of the tumor size variability in the BraTS2018 database: the top row shows axial slices of LGG (top row) and HGG (bottom row) tumors with four MRI modalities and the corresponding segmentation label used in this work.

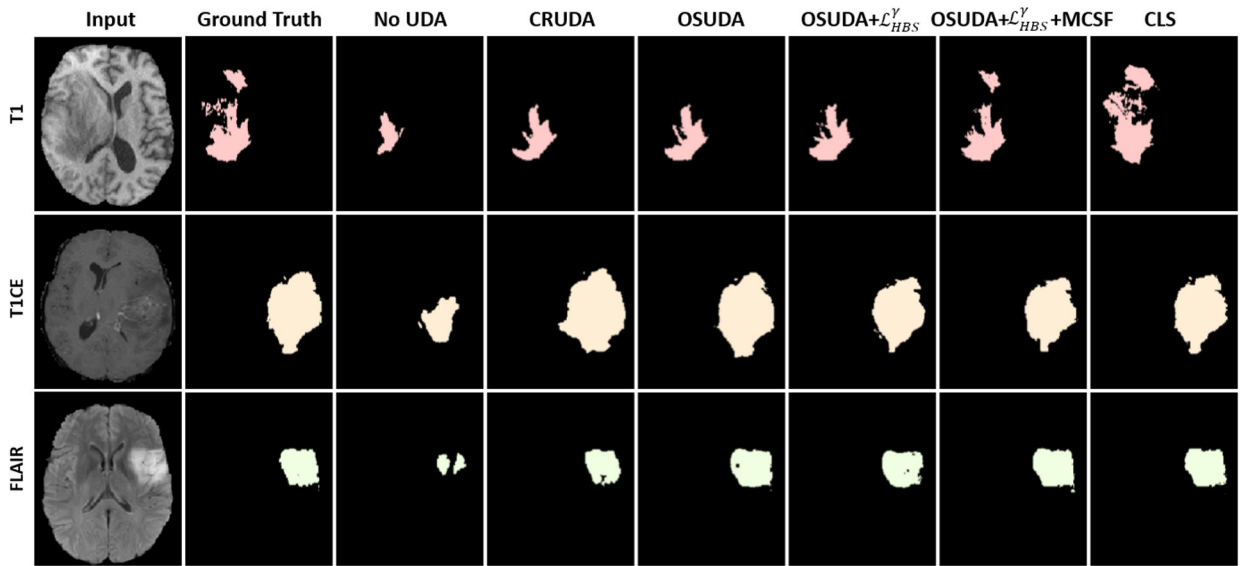


Fig. 6: Comparisons of our framework against other UDA methods, and ablation studies of T2-weighted MRI-to-T1-weighted/T1ce/FLAIR MRI UDA for whole tumor segmentation using BratS2018. Note that CLS (Liu et al., 2021a) with source data for training is regarded as an “upper bound.”




Fig. 7: Comparisons of our framework against other UDA methods, and ablation studies for HGG to LGG cross-subtype brain tumor UDA segmentation. The source-available CLS (Liu et al., 2021a) is regarded as an “upper bound.”

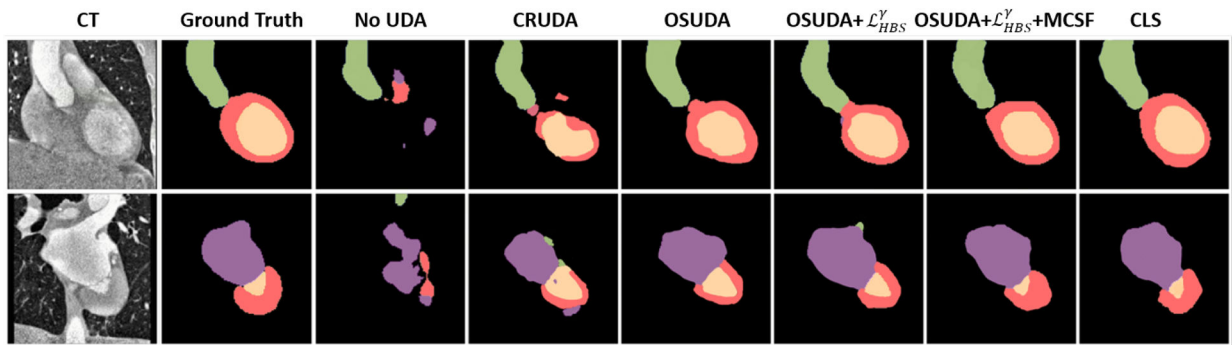


Fig. 8: Comparisons of our framework against other UDA methods, and ablation studies for cardiac MR to CT segmentation. The source-available CLS (Liu et al., 2021a) is regarded as an “upper bound.”




Fig. 9: The average DSC of OSUDA+ \mathcal{L}_{HBS}^i + MCSF(MCOSUDA) with different batch size (left) and H in the memory queue for the cross-modality (blue) and the HGG to LGG (red) brain tumor segmentation task, and the cardiac MR to CT segmentation task (green).




Fig. 10:
The DSC of OSUDA+ $\mathcal{L}_{HBS}^{\gamma}$ + MCSF(MCOSUDA) in cross-modality tumor segmentation (left) and cardiac MR to CT segmentation (right) with different λ .

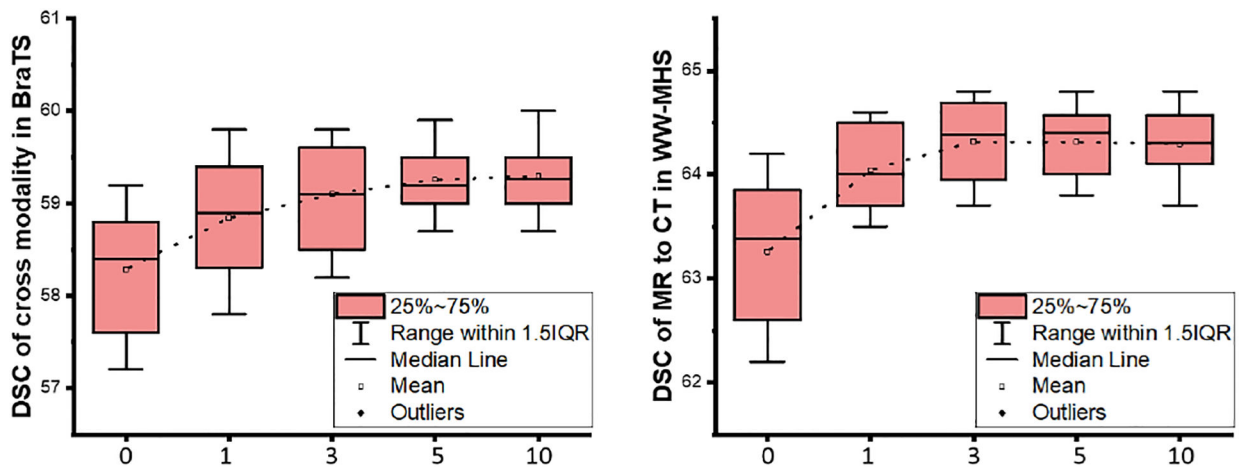


Fig. 11: The DSC of OSUDA+ \mathcal{L}'_{HBS} + MCSF(MCOSUDA) in cross-modality tumor segmentation (left) and cardiac MR to CT segmentation (right) with different ϕ .

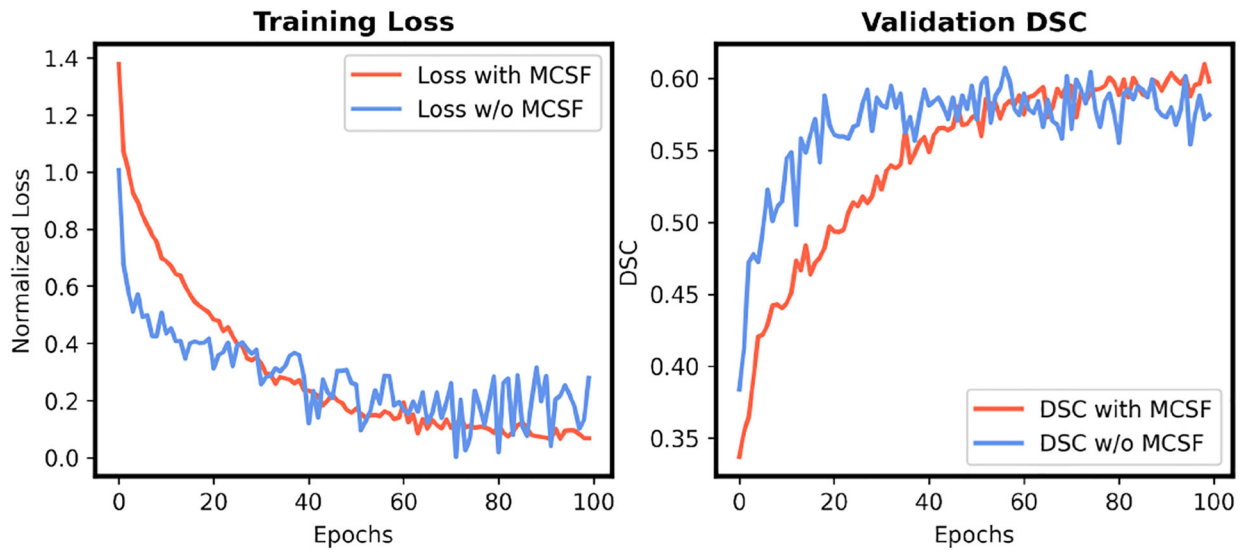


Fig. 12: The training loss (left) and validation DSC (right) of our OSUDA, with or without MCSF in cross-MR-modality tumor segmentation.

Table 1: Comparisons of our framework against other UDA methods on the cross-MR-modality whole brain tumor UDA segmentation.

Method	Source data	Dice			Score			[mm] ↓		
		T1	FLAIR	T1CE	T1CE	T1CE	Average	T1	FLAIR	T1CE
Source only (no adaption)	no UDA	6.78	54.36	6.71	22.62±0.176	58.7	21.5	60.2	46.8±0.15	
CRUDA (Bateson et al., 2020)	Partial	47.25	65.63	49.47	54.12±0.160	22.1	17.5	24.4	21.3±0.10	
GTA (Kundu et al., 2021)	Partial	52.24	64.03	52.46	56.24±0.135	23.2	16.7	22.8	20.9±0.11	
DPL (Chen et al., 2021)	no	50.35	63.17	51.78	55.10±0.121	23.4	17.0	22.6	21.0±0.09	
SFKT (Liu et al., 2021h)	no	51.71	61.50	52.42	55.31±0.197	22.5	16.9	24.1	21.2±0.07	
OSUDA (Liu et al., 2021g)	no	52.71	67.60	53.22	57.84±0.153	20.4	16.6	22.8	19.9±0.08	
OSUDA-AC	no	51.58	66.45	52.12	56.72±0.164	21.5	17.8	23.6	21.0±0.12	
OSUDA-SE	no	51.14	65.79	52.80	56.58±0.142	21.6	17.3	23.3	20.7±0.10	
OSUDA + \mathcal{L}_{HBS}^i	no	53.36	67.94	53.58	58.29±0.120	20.5	16.4	21.7	19.5±0.08	
OSUDA + \mathcal{L}_{HBS}^i + MCSF(MCOSUDA)	no	54.51	68.37	54.62	59.17±0.135	19.4	15.8	21.0	18.7±0.09	
CycleGAN (Zhu et al., 2017)	Yes	38.1	63.3	42.1	47.8	25.4	17.2	23.2	21.9	
SIFA (Chen et al., 2019)	Yes	51.7	68.0	58.2	59.3	19.6	16.9	15.01	17.1	
DSFN (Zou et al., 2020)	Yes	57.3	78.9	62.2	66.1	17.5	13.8	15.5	15.6	
CLS (Liu et al., 2021a)	Yes	56.92	78.75	63.24	66.30±0.108	17.2	13.5	15.8	15.5±0.13	

T2-weighted MRI is used as the source domain, while T1-weighted, FLAIR, and T1ce MRI are used as the unlabeled target domains. The source-available UDA methods are regarded as “upper bounds.”

Table 2:

Comparisons of our framework against other UDA methods on cross-subtype, i.e., HGG to LGG UDA segmentation.

Method	Source data	DSC [%] ↑			HD [mm] ↓				
		WholeT	EnhT	CoreT	WholeT	EnhT	CoreT	Overall	
Source only (Shanis et al., 2019)	no UDA	79.29	30.09	44.11	58.44±0.435	38.7	46.1	40.2	41.7±0.14
CRUDA (Bateson et al., 2020)	Partial	79.85	31.05	43.92	58.51±0.123	31.7	29.5	30.2	30.6±0.15
GTA (Kundu et al., 2021)	Partial	83.12	31.92	46.25	61.38±0.152	27.6	24.8	27.4	26.2±0.16
DPL (Chen et al., 2021)	no	80.24	31.26	45.13	59.72±0.136	29.2	26.5	28.9	28.0±0.13
SFKT (Liu et al., 2021h)	no	81.79	31.86	46.42	60.70±0.141	29.0	26.2	28.1	27.3±0.15
OSUDA (Liu et al., 2021g)	no	83.62	32.15	46.88	61.94±0.108	27.2	23.4	26.3	25.6±0.14
OSUDA-AC	no	82.74	32.04	46.62	60.75±0.145	27.8	25.5	27.3	26.5±0.16
OSUDA-SE	no	82.45	31.95	46.59	60.78±0.120	27.8	25.3	27.1	26.4±0.14
OSUDA + \mathcal{L}'_{HBS}	no	84.08	32.51	47.40	62.36±0.132	26.9	23.1	25.8	25.0±0.15
OSUDA + \mathcal{L}'_{HBS} + MCSF(MCOSUDA)	no	84.29	32.86	47.63	62.87±0.101	26.4	22.3	23.8	24.1±0.09
SEAT (Shanis et al., 2019)	Yes	84.11	32.67	47.11	62.17±0.153	26.4	21.7	23.5	23.8±0.16
CLS (Liu et al., 2021a)	Yes	85.37	34.94	49.25	64.02±0.135	25.6	21.0	22.3	22.9±0.15

The source-available methods, e.g., SEAT (Shanis et al., 2019), are regarded as the “upper bounds.”

Table 3: Comparisons of our framework against other UDA methods on cardiac MR to CT segmentation.

Method	Source data	DSC [%] ↑					HD [mm] ↓				
		AA	LAC	LVC	MYO	Average	AA	LAC	LVC	MYO	Average
Source only (Zou et al., 2020)	no UDA	28.4	27.7	4.0	8.7	17.2	32.6	35.1	54.2	57.4	44.8
CRUDA (Bateson et al., 2020)	Partial	74.94	72.33	58.26	30.40	58.98±0.109	21.7	22.1	24.3	30.7	24.7±0.14
SFKT (Liu et al., 2021h)	no	75.83	71.49	59.76	34.58	60.41±0.152	19.4	18.5	23.2	29.1	22.5±0.11
OSUDA(Liu et al., 2021g)	no	78.02	73.55	60.18	40.68	63.11±0.160	11.6	15.1	19.4	20.7	16.7±0.13
OSUDA-AC	no	72.89	73.43	60.02	40.41	62.94±0.145	12.0	16.4	19.9	21.3	17.4±0.16
OSUDA-SE	no	77.85	73.36	60.10	40.17	62.87±0.142	12.8	16.5	19.4	21.7	17.6±0.15
OSUDA + $\mathcal{L}_{HBS}^{\prime}$	no	78.16	73.82	60.39	41.46	63.46±0.117	11.4	14.2	18.4	20.1	16.0±0.12
OSUDA + $\mathcal{L}_{HBS}^{\prime}$ + MCSF(MCOSUDA)	no	78.80	75.63	61.45	42.38	64.57±0.195	10.7	11.9	18.3	19.5	15.1±0.10
CycleGAN (Zhu et al., 2017)	Yes	73.8	75.7	52.3	28.7	57.6	16.2	15.4	20.8	27.7	20.0
SIFA (Chen et al., 2019)	Yes	81.1	76.4	75.7	58.7	73.0	8.2	10.5	12.2	17.9	12.2
DSFN (Zou et al., 2020)	Yes	84.7	76.9	79.1	62.4	75.8	7.4	11.9	10.6	15.7	11.4
CLS (Liu et al., 2021a)	Yes	84.82	77.30	79.04	62.65	75.95±0.138	7.6	11.7	10.4	15.5	11.3±0.15

± indicates standard deviation. The source available UDA methods are regarded as "upper bounds."

Table 4:

Comparisons of cardiac MR to CT segmentation with only one target domain subject for training as in OLVA (Chanti and Mateus, 2021).

Method	Source data	DSC [%] \uparrow				
		AA	LAC	LVC	MYO	Average
CRUDA	Partial	53.45	59.20	73.32	35.64	55.40 \pm 0.13
SFKT	no	51.38	56.93	71.30	33.76	53.34 \pm 0.14
OSUDA	no	57.96	63.37	77.85	39.82	59.75 \pm 0.11
-AC	no	57.63	62.03	77.45	39.74	59.21 \pm 0.15
-SE	no	57.13	62.41	77.32	39.26	59.03 \pm 0.13
+ \mathcal{L}_{HBS}^y	no	58.06	63.49	78.12	39.94	59.90 \pm 0.12
+ \mathcal{L}_{HBS}^y + MCSF	no	58.27	64.18	78.44	40.51	60.35\pm0.08
SIFA	Yes	62	53	80	39	62
OLVA	Yes	60	70	78	68	69

\pm indicates standard deviation. The UDA methods with source domain data are regarded as "upper bounds."

Table 5:

The average DSC of OSUDA with fixed or linear changed λ for the BraTS cross modality/subtype and MR to CT segmentation tasks.

Cross modality segmentation task using BraTS18						
λ	10 \rightarrow 0	10	5	3	1	0
DSC	57.84	56.98	57.15	57.26	57.04	56.58
HGG to LGG segmentation task using BraTS18						
λ	10 \rightarrow 0	10	5	3	1	0
DSC	61.94	60.97	61.02	61.14	61.08	60.78
MR to CT segmentation task using MM-WHS						
λ	10 \rightarrow 0	10	5	3	1	0
DSC	63.11	62.95	62.96	62.95	62.90	62.87

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 6:

The average DSC of MCOSUDA with different α for three tasks.

Cross modality segmentation task using BraTS18					
α	20→80	10→80	30→80	20→70	20→90
DSC	59.17	59.02	58.87	59.13	58.74
HGG to LGG segmentation task using BraTS18					
α	20→80	10→80	30→80	20→70	20→90
DSC	62.87	62.74	62.65	62.72	62.46
MR to CT segmentation task using MM-WHS					
α	20→80	10→80	30→80	20→70	20→90
DSC	64.57	64.50	64.48	64.49	64.42

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript