

Constraining pseudo-label in self-training unsupervised domain adaptation with energy-based model

Lingsheng Kong¹ | Bo Hu² | Xiongchang Liu³ | Jun Lu⁴ | Jane You⁵ | Xiaofeng Liu⁶ 

¹Changchun Institute of Optics, Fine Mechanics and Physics, of the Chinese Academy of Sciences, Changchun, China

²Department of Accounting, National University of Singapore, Singapore

³Department of Information and Electrical Engineering, China University of Mining and Technology, Beijing, China

⁴Beth Israel Deaconess Medical Center and Harvard Medical School, Boston, MA, USA

⁵Department of Computing, The Hong Kong Polytechnic University, Hong Kong, China

⁶Gordon Center for Medical Imaging, Harvard University, Cambridge, MA, USA

Correspondence

Xiaofeng Liu, Gordon Center for Medical Imaging, Harvard University, Cambridge 02138, MA, USA.
Email: xliu61@mg.harvard.edu and liuxiaofeng@cmu.edu

Abstract

Deep learning is usually data starved, and the unsupervised domain adaptation (UDA) is developed to introduce the knowledge in the labeled source domain to the unlabeled target domain. Recently, deep self-training presents a powerful means for UDA, involving an iterative process of predicting the target domain and then taking the confident predictions as hard pseudo-labels for retraining. However, the pseudo-labels are usually unreliable, thus easily leading to deviated solutions with propagated errors. In this paper, we resort to the energy-based model and constrain the training of the unlabeled target sample with an energy function minimization objective. It can be achieved via a simple additional regularization or an energy-based loss. This framework allows us to gain the benefits of the energy-based model, while retaining strong discriminative performance following a plug-and-play fashion. The convergence property and its connection with classification expectation minimization are investigated. We deliver extensive experiments on the most popular and large-scale UDA benchmarks of image classification as well as semantic segmentation to demonstrate its generality and effectiveness.

KEYWORDS

domain adaptation, energy-based model, self-training

1 | INTRODUCTION

Deep learning (DL) has demonstrated its effectiveness in many areas, while it is usually data-starved and relies on the *i. i. d* assumption of training and testing data. However, there usually exists a discrepancy when we deploy the model in a new environment. The collection of a large number of target domain data for DL is challenging for many applications (e.g., medical image analysis), thereby hindering the wide adoption of DL methods.^{1,2} For example, densely annotating a Cityscapes image on average takes about 90 min,³ severely hindering the generalization of an autonomous driving system in different cities. Therefore, the unsupervised domain adaptation (UDA) is proposed to transfer knowledge from one labeled source domain to another target domain with the aid of unlabeled target data.^{4–6}

To this end, several solutions have been explored, for example, feature/image-level adversarial training and maximum mean discrepancy.^{7,8} Recently, one of the promising methods in UDA is the self-training,^{9,10} which iteratively generates a set of one-hot pseudo-labels in the target domain, and then retrains the network according to these pseudo-labels of the unlabeled target data. The typical choice of the recent deep self-training usually be the one-hot pseudo-label (i.e., hard label), and demonstrating that it is essentially an entropy minimization process,⁹ which enforces the output of the network to be as sharp as the one-hot hard pseudo-label.

However, the correctness of the generated pseudo-label is largely uncontrolled and can be hard to guarantee. Actually, the state-of-the-art accuracy of many UDA benchmarks is approximately 50%. Especially in the first few epochs, it is hard to produce a reliable pseudo-label. In addition, the labels of natural images can also be highly ambiguous. Taking a sample image from VisDA17¹¹ (see Figure 1) for illustration. We can see that both the person and car dominate significant portions in this image. Enforcing a model to be very confident in only one of the classes during training can also hurt the learning behavior,¹² particularly within the context of no ground truth label for the target samples. Encoding the pseudo-label as the one hot hard label vector is essentially trusting all of the pseudo-label as the “ground truth,” which can lead to overconfident mistakes and propagated errors.¹⁰

Targeting for this issue, our previous work¹⁰ proposes to make the network more conservative w.r.t. the pseudo-label of target sample via the manually defined pseudo label smoothing or entropy minimization. Specifically,¹⁰ proposes to modify the possibly inaccurate pseudo-label histogram distribution to be more smooth following a fixed smoothing operation for any data sample. For instance, the conventional self-training only transform the three-classes prediction [0.3, 0.6, 0.2] to the one-hot pseudo-label [0, 1, 0] with an argmax operation. Therefore, we can further apply the label smoothing to revise [0, 1, 0] to [0.05, 0.9, 0.05] for the subsequent conservative self-training. With the conventional CE loss, the confidence of the label can be reduced, which aligns with the practical reliability of pseudo-label. However, these constraints for pseudo-label can not be adaptively adjusted for different inputs or network parameters. It is clear that the reliability of the pseudo-label of the sample in Figure 1 can be significantly lower than the normal samples. Besides, the pseudo-label produced in the later stage can be more trustable than the early stage if the training goes smoothly. Moreover, the label smoothing is explicitly dependent on the inaccurate pseudo-label, since it only revises the weight of histogram distribution. Actually, if the model with the current parameters is sufficiently confident to a specific sample, it is more reasonable to trust its pseudo label.

The difficulties mentioned above motivate us to develop an adaptive regularizer rather than the pre-defined pseudo label smoothing rules. We would expect our regularization signal for the target sample can adaptively reflect the current input and the network parameters in this training iteration rather than depend on the inaccurate pseudo label.

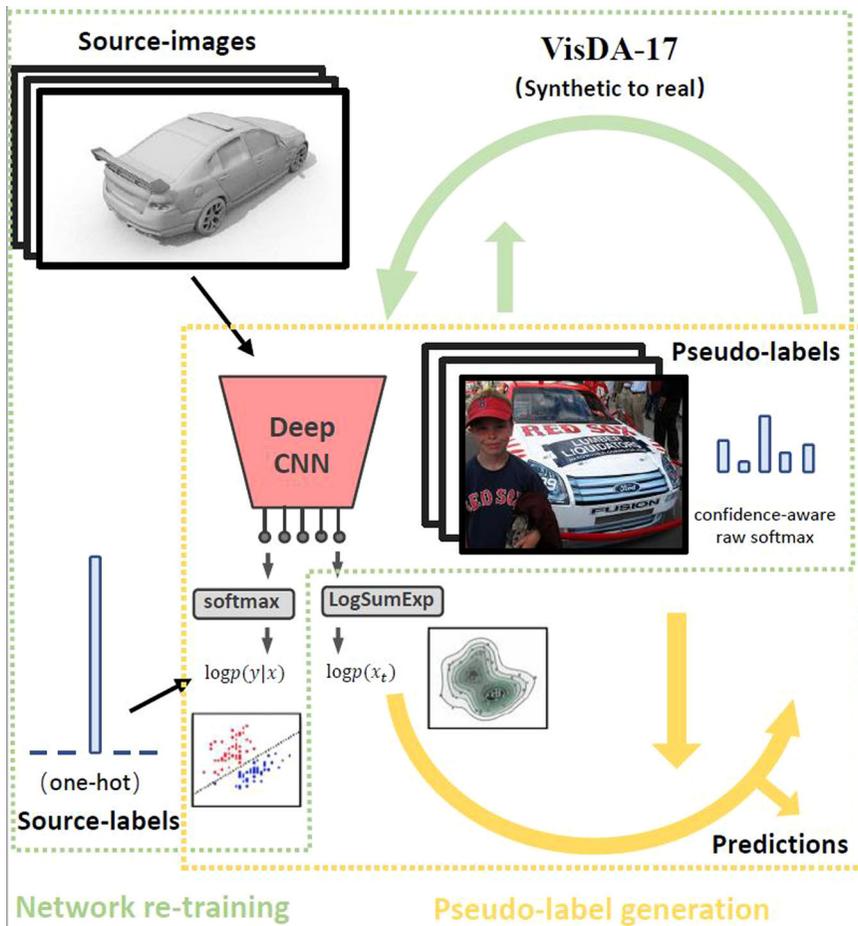


FIGURE 1 The illustration of our energy-constrained self-training framework for UDA. Minimizing the pseudo-label-irrelevant energy of $E_w(x_t)$ is introduced as additional objective for the target sample. UDA, unsupervised domain adaptation. [Color figure can be viewed at wileyonlinelibrary.com]

Recently, the correlation between the standard discriminative classifier and the energy-based model (EBM) is revealed.¹³ Essentially, the standard discriminative classifier can be reinterpreted as an EBM for the joint distribution $p(\mathbf{x}, \mathbf{y})$. Both the EBM and supervised discriminative model can be benefited from simultaneously optimizing both $p(\mathbf{y}|\mathbf{x})$ with cross-entropy (CE) loss and optimizing $\log p(\mathbf{x})$ with EBM,¹³ respectively. Here, we demonstrate that optimizing $\log p(\mathbf{x})$ can be even more promising on self-training-based UDA. Since we do not have reliable labels for C.E. loss training on the target domain, and $\log p(\mathbf{x})$ can be an idea regularization signal. We note that $\log p(\mathbf{x})$ is explicitly correlated with the input \mathbf{x} and network parameter w.r.t. $p(\cdot)$, and independent to pseudo label \mathbf{y} .

Therefore, we propose a simple and straightforward idea that configures the energy minimization of data point \mathbf{x} (i.e., $\log p(\mathbf{x})$) as an additional regularization term. The target domain examples are optimized with pseudo-label C.E. loss and EBM objective. With the help of the EBM, our self-training is expected to be more controllable. We give a thorough convergence analysis and theoretically connect it with classification expectation-maximization (CEM).

In addition to the EBM-based regularization, we further propose an efficient EBM loss for the target sample, which is able to integrate the confidence-aware pseudo-label into the energy minimization objective. Based on the insights that the EBMs capture dependencies by associating a scalar energy value (i.e., a measure of compatibility) to each configuration of the variables (e.g., \mathbf{y} and \mathbf{x}), we should not minimize the energy of all possible \mathbf{y} and \mathbf{x} configuration. It is more reasonable to focus on minimizing the correct energy w.r.t. configuration of \mathbf{y} and \mathbf{x} . In the context of self-training UDA, the unprocessed soft pseudo-label can be a good signal as the weights of energy at each class. The histogram distribution value can also inherit the confidence of the discriminator, which provides more information for the later round training. Besides, the pseudo-label of Figure 1 can be a multi-modal distribution which is also able to alleviate the adverse effect of the confusing images. It not only speeds up the training, removes the hyper-parameter validation but also achieves a comparable or even better performance than EBM regularization.

This study extends our previous paper¹⁴ in the following significant ways:

- We propose a practical and scalable EBM loss for the target sample with the confidence-aware soft pseudo-label. It can efficiently speed up the training and remove the hyper-parameter validation, while not sacrificing the performance.
- We theoretically analyze the convergence property of our energy-constrained self-training. Besides, its connection with the classification expectation maximization (CEM) is investigated from the probabilistic perspective.

In addition, we further demonstrate the generality of our methods (both EBM regularizer and EBM loss) in more UDA benchmarks (e.g., Office-31 and SYNTHIA2Cityscapes) with more visualization and analysis.

The main contributions of this paper are summarized as

- We propose a novel and intuitive framework to incorporate the EBM into the self-training UDA as a pseudo-label-irrelevant regularization signal, which benefits from both the generative and the discriminative model.
- It can simply be formulated as a regularization term of hard pseudo-label supervised training, and can also be an EBM loss for the target sample with the confidence-aware soft pseudo-label.
- The convergence property of our energy-constrained self-training is investigated. In addition, it can be explained from the probabilistic perspective. The throughout theoretical analysis bridge its connection with the CEM.

We empirically validate the effectiveness and generality of the proposed method on multiple challenging benchmarks (both classification and semantic segmentation) and achieve state-of-the-art performance.

2 | RELATED WORKS

2.1 | UDA

DL has achieved tremendous milestones in computer vision recently.¹⁵ Instead of designing features by hand and then feeding the features to a prediction model, DL suggests learning specific

features and the prediction model simultaneously from raw image following an end-to-end fashion.^{16,17} However, DL is usually data-starved and relies on the *i. i. d* assumption of training and testing data. The model trained on an annotated source domain does not generalize well on a different target domain due to the domain shift (or drift).^{18,19} In reality, collecting the large-scale labeled data in the new target domain is expensive or even prohibitive. Therefore, we would expect to migrate the knowledge from a labeled source domain to a different unlabeled target domain.^{20,21}

UDA with deep networks targets to learn domain invariant embeddings by minimizing the cross-domain difference of feature distributions with certain criteria. Examples of these methods include maximum mean discrepancy (MMD), deep Correlation Alignment (CORAL), sliced Wasserstein discrepancy, adversarial learning at input-level, feature level, output space level, and so forth.^{22,23} Despite the underlying difference among these techniques, there exists an interesting connection between some of these methods with conditional forms¹ and Self-training, as they can be broadly considered as E.M. algorithms,²⁴ and such conditional formulation has been widely proved to benefit the adaptation.¹⁰

Other than the self-training, there are some of the other UDA methods that also rely on the pseudo label. For instance,²⁵ utilizes pseudo labels to estimate target class centers, which are used to match source class centers. CAN²⁶ utilizes target pseudo-labels to estimate contrastive domain discrepancy. Since the pseudo labels can be noisy,²⁷ proposes to measure its correctness by the Gaussian-uniform mixture model. Instead, we rely on our energy model for regularization.

2.2 | Self-training

Self-training was initially proposed for the semi-supervised learning.²⁸ The deep self-training is different from the self-training with fixed hand-crafted feature input in that the deep self-training involves the learning of deep embedding, which renders greater flexibility toward domain alignment than classifier-level adaptation only. With the development of DL and UDA, there have been several deep self-training/pseudo-label-based works that are developed for UDA.^{24,29,30}

Targeting the noisy and unreliable pseudo-label, our preliminary work¹⁰ proposes to construct a more conservative pseudo-label that smoothes the one-hot hard label to a soft label vector.

The EBM regularizer proposed in this study can be orthogonal to the recent progress of self-training.¹⁰ We resort to the additional supervision signal of EBM, which is independent of the pseudo-label. In contrast with the pre-defined label smoothing as Zou et al.,¹⁰ our proposed energy constraint is able to adaptively regularize the training with respect to the sample and the current network parameters. More appealingly, the EBM-based regularization can be simply added on state-of-the-art self-training methods following a plug-and-play fashion.

2.3 | EBMs

EBMs is developed to capture the dependencies between variables via associating scalar energy to each configuration of the variables.³¹ The recent studies^{13,32} propose to approximate the expectation of the log-likelihood for a single example \mathbf{x} w.r.t. the network parameter with a sampler based on the Stochastic Gradient Langevin Dynamics (SGLD).³³

EBM is a typical unsupervised generative model. To combine the classifier and EBMs,^{34,35} reinterpreted the logits (i.e., network prediction before the softmax) to define a class-conditional EBM $p(\mathbf{x}|\mathbf{y})$, which requires additional parameters to be learned to derive a

classifier and an unconditional model.³⁶ inherits a similar idea as well, while it is trained with a GAN-like generator and targets different applications. In contrast, we do not focus on the generative model of $p(\mathbf{x}|\mathbf{y})$, and propose to utilize the probability density of $p(\mathbf{x})$ as regularizer. Based on Nijkamp et al.,³² the recent study¹³ scales the training of EBMs to high-dimensional data with the Contrastive Divergence and SGLD.

In this paper, we propose to regard EBMs as the pseudo-label-irrelevant optimization constraint, and explore their potentials on UDA. This helps realize the potential of EBMs on downstream supervised discriminative problems. In addition, thorough convergence analysis and theoretical connection with CEM are provided.

3 | METHODOLOGY

In the scenario of UDA, we can access to the labeled source samples $(\mathbf{x}_s, \mathbf{y}_s)$ from source domain $\{\mathbf{X}_S, \mathbf{Y}_S\}$, and target samples \mathbf{x}_t from unlabeled target domain data \mathbf{X}_T . Any target label $\hat{\mathbf{y}}_t = (\hat{y}_t^{(1)}, \dots, \hat{y}_t^{(K)})$ from $\hat{\mathbf{Y}}_T$ is unknown. K is the total number of classes. The parametric network $f_{\mathbf{w}} : \mathbb{R}^D \rightarrow \mathbb{R}^K$ with the weights \mathbf{w} is used to process the D -dim input sample. The output K real-valued numbers known as logits and usually followed by a softmax normalization to produce $p_{\mathbf{w}}(k|\mathbf{x}) = \exp(f_{\mathbf{w}}(\mathbf{x})[k]) / \sum_{k \in K} \exp(f_{\mathbf{w}}(\mathbf{x})[k])$ as the classifier's softmax probability for class k . $f_{\mathbf{w}}(\mathbf{x})[k]$ indicates the k th index of $f_{\mathbf{w}}(\mathbf{x})$, that is, the logit corresponding the the k th class label.

3.1 | Preliminary of EBM

The energy based function $E_{\mathbf{w}}(\mathbf{x}) : \mathbb{R}^D \rightarrow \mathbb{R}$ maps each point of an input space to a single scalar, which is called energy. One can parameterize an EBM using any function that takes \mathbf{x} as the input and returns a scalar. The learning is designed in such a way that $E_{\mathbf{w}}$ can assign low energies to observed configurations of variables while give high energies to unobserved ones.³² With $E_{\mathbf{w}}$, the probability density $p(\mathbf{x})$ for $\mathbf{x} \in \mathbb{R}^D$ in the EBMs is formulated as

$$p_{\mathbf{w}}(\mathbf{x}) = \frac{\exp(-E_{\mathbf{w}}(\mathbf{x}))}{Z(\mathbf{w})}, \quad (1)$$

where $Z(\mathbf{w}) = \int_{\mathbf{x}} \exp(-E_{\mathbf{w}}(\mathbf{x}))$ is the normalizing constant (with respect to \mathbf{x}) also known as the partition function. Usually, the normalized densities $p_{\mathbf{w}}(\mathbf{x})$ are intractable, since we cannot reliably estimate $Z(\mathbf{w})$ for the most choice of $E_{\mathbf{w}}(\mathbf{x})$. The typical solutions rely on the sophisticate Markov Chain Monte Carlo (MCMC) sampler to train EBMs.³² In our work, we use the deep neural networks as $E_{\mathbf{w}}$. The gradient estimation is usually used for EBM optimization, and sample data from it by MCMC.³²

3.2 | EBM for discriminative model

The recent study¹³ reveals that a standard classifier can be interpreted as an EBM of joint distribution $p_{\mathbf{w}}(\mathbf{x}, \mathbf{y})$, and optimizing the likelihood $\log p_{\mathbf{w}}(\mathbf{x}, \mathbf{y}) = \log p_{\mathbf{w}}(\mathbf{x}) + \log p_{\mathbf{w}}(\mathbf{y}|\mathbf{x})$

can be helpful for both the discrimination and generation task. Specifically, optimizing $p(\mathbf{y}|\mathbf{x})$ is simply achieved by using the standard CE loss as conventional classification. The additional optimization objective $\log p_{\mathbf{w}}(\mathbf{x})$ has been proven and demonstrating that can improve the confidence calibration and robustness for conventional classification task.¹³

Considering the target samples do not have a ground truth label, the self-training methods¹⁰ utilize the inaccurate pseudo label to calculate the CE loss. Therefore, optimizing $\log p_{\mathbf{w}}(\mathbf{x})$ can potentially be more helpful for UDA setting. Actually, $\log p_{\mathbf{w}}(\mathbf{x})$ is adaptive w.r.t. the input \mathbf{x} and network parameter \mathbf{w} , and irrelevant to the inaccurate pseudo label, which can be an ideal regularizer of self-training based UDA.

However, how to modeling $\log p_{\mathbf{w}}(\mathbf{x})$ can be a challenging task. Considering $\frac{\partial \log p_{\mathbf{w}}(\mathbf{x})}{\partial \mathbf{w}}$ can be approximated with $-\frac{\partial E_{\mathbf{w}}(\mathbf{x})}{\partial \mathbf{w}}$,¹³ it is possible to modeling the energy function $E_{\mathbf{w}}(\mathbf{x})$ instead of $\log p_{\mathbf{w}}(\mathbf{x})$.

As detailed in Equation (1), in EBMs, the probability density $p_{\mathbf{w}}(\mathbf{x})$ for $\mathbf{x} \in \mathbb{R}^D$ can be formulated as $p_{\mathbf{w}}(\mathbf{x}) = \exp(-E_{\mathbf{w}}(\mathbf{x}))/Z(\mathbf{w})$. Following Grathwohl et al.,¹³ we can define an EBM of the joint distribution $p_{\mathbf{w}}(\mathbf{x}, \mathbf{y}) = \exp(f_{\mathbf{w}}(\mathbf{x})[k])/Z(\mathbf{w})$, by defining $E_{\mathbf{w}}(\mathbf{x}, \mathbf{y}) = -f_{\mathbf{w}}(\mathbf{x})[k]$. By marginalizing out \mathbf{y} , we have $p_{\mathbf{w}}(\mathbf{x}) = \frac{\sum_k \exp(f_{\mathbf{w}}(\mathbf{x})[k])}{Z(\mathbf{w})}$.¹³ Considering $p_{\mathbf{w}}(\mathbf{x}) = \exp(-E_{\mathbf{w}}(\mathbf{x}))/Z(\mathbf{w})$, the energy function of \mathbf{x} can be

$$E_{\mathbf{w}}(\mathbf{x}) = -\text{LogSumExp}_k(f_{\mathbf{w}}(\mathbf{x})[k]) = -\log \sum_k \exp(f_{\mathbf{w}}(\mathbf{x})[k]). \quad (2)$$

In this setting, $p_{\mathbf{w}}(\mathbf{y}|\mathbf{x}) = \frac{p_{\mathbf{w}}(\mathbf{x}, \mathbf{y})}{p_{\mathbf{w}}(\mathbf{x})} = \frac{\exp(f_{\mathbf{w}}(\mathbf{x})[k])/Z(\mathbf{w})}{\sum_k \exp(f_{\mathbf{w}}(\mathbf{x})[k])/Z(\mathbf{w})}$. The normalization constant $Z(\mathbf{w})$ will be canceled out and yielding the standard softmax function, which bridges the EBMs and conventional classifiers. Since we can factorize the likelihood as

$$\log p_{\mathbf{w}}(\mathbf{x}, \mathbf{y}) = \log p_{\mathbf{w}}(\mathbf{x}) + \log p_{\mathbf{w}}(\mathbf{y}|\mathbf{x}) \quad (3)$$

we can simultaneously optimize $p(\mathbf{y}|\mathbf{x})$ with standard CE loss and $\log p_{\mathbf{w}}(\mathbf{x})$ with SGLD (where gradients are taken w.r.t. $\text{LogSumExp}_k(f_{\mathbf{w}}(\mathbf{x})[k])$) in the supervised training setting.¹³

Compared with conventional classifier which will not be affected by shifting the logits $f_{\mathbf{w}}(\mathbf{x})$ with a scalar, the $\log p_{\mathbf{w}}(\mathbf{x})$ is sensitive to this kind of perturbation. It introduces an extra degree of freedom w.r.t. logits to define the density function over input examples and the joint density among examples and labels, therefore further constraining the training.

3.3 | EBM as regularization for UDA

The self-training-based UDA proposes to utilize the iterative loss minimization scheme to gradually refine the model for the target domain. In self-training, the pseudo-labels are regarded as the learnable latent variables.¹⁰ Considering the target domain samples do not have the ground truth label, the optimization objective in self-training relies on the CE loss with the pseudo label, which is the relatively confident prediction in the previous round.²⁴ Considering the inaccuracy of pseudo label, the smoothed soft label¹⁰ is proposed as an alternative to the one-hot hard pseudo label.²⁴

Therefore, a straightforward solution for adapting EBM to UDA is to integrate Equation (2) as a regularization term. For the target sample, the network is enforced to achieve a good

prediction of pseudo-label and minimize the energy of $E_{\mathbf{w}}(\mathbf{x}_t)$. The energy constraint is inherently dependent on \mathbf{x} and \mathbf{w} . Therefore, it is more flexible than the pre-defined label smoothing or entropy regularization.¹⁰ Given that the pseudo-label is usually noisy, the latter objective is expected to be even more helpful than the supervised learning task.

Following the notations in our CRST,¹⁰ the self-training with EBM regularization (R-EBM) for target sample, that is, $E_{\mathbf{w}}(\mathbf{x}_t)$, can be formulated as

$$\begin{aligned} \min_{\mathbf{w}, \hat{\mathbf{Y}}_T} \mathcal{L}_{R-EBM}(\mathbf{w}, \hat{\mathbf{Y}}) &= - \sum_{s \in S} \sum_{k=1}^K y_s^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_s) \\ &- \sum_{t \in T} \left\{ \sum_{k=1}^K \left[\hat{y}_t^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_t) - \hat{y}_t^{(k)} \log \lambda_k \right] - \alpha E_{\mathbf{w}}(\mathbf{x}_t) \right\} \\ \text{s. t. } \hat{\mathbf{y}}_t &\in \Delta^{K-1} \cup \{\mathbf{0}\}, \forall t. \end{aligned} \quad (4)$$

For each class k , λ_k is determined by the confidence value selecting the most confident p portion of class k predictions in the entire target set.²⁴ If a sample's predication is relatively confident with $p_{\mathbf{w}}(k^*|\mathbf{x}_t) > \lambda_{k^*}$, it is selected and labeled as class $k^* = \operatorname{argmax}_k \left\{ \frac{p_{\mathbf{w}}(k|\mathbf{x}_t)}{\lambda_k} \right\}$. The less confident ones with $p_{\mathbf{w}}(k^*|\mathbf{x}_t) \leq \lambda_{k^*}$ are not selected. The same class-balanced λ_k strategy introduced in Zou et al.²⁴ is adopted for all self-training methods in this study.

The feasible set is the union of $\{\mathbf{0}\}$ and a probability simplex Δ^{K-1} .²⁴ α is a balancing hyperparameter of the regularization term, which does not directly relate to the label \hat{y}_t . The self-training can be solved by an alternating optimization scheme.

Step 1) Pseudo-label generation Fix \mathbf{w} and solve:

$$\begin{aligned} \min_{\hat{\mathbf{Y}}_T} &- \sum_{t \in T} \left\{ \sum_{k=1}^K \hat{y}_t^{(k)} [\log p_{\mathbf{w}}(k|\mathbf{x}_t) - \log \lambda_k] - \alpha E_{\mathbf{w}}(\mathbf{x}_t) \right\} \\ \text{s. t. } \hat{\mathbf{y}}_t &\in \Delta^{K-1} \cup \{\mathbf{0}\}, \quad \forall t. \end{aligned} \quad (5)$$

For solving Step 1), there is a global optimizer for arbitrary $\hat{\mathbf{y}}_t = (\hat{y}_t^{(1)}, \dots, \hat{y}_t^{(K)})$ as¹⁰

$$\hat{y}_t^{(k)*} = \begin{cases} 1, & \text{if } k = \operatorname{argmax}_k \frac{p_{\mathbf{w}}(k|\mathbf{x}_t)}{\lambda_k} \quad \text{and} \quad p_{\mathbf{w}}(k|\mathbf{x}_t) > \lambda_k \\ 0, & \text{otherwise} \end{cases}. \quad (6)$$

Despite a long period of little progression, there have been some recent works^{32,35} using the sampler based on SGLD to train the large-scale EBMs on high-dimensional data, which is parameterized by deep neural networks.

Given that the goal of this study is to incorporate EBM training into the standard classification setting, the classification part is the same as Zou et al.,¹⁰ and we only change the regularizer. Therefore, we can follow¹³ to train the network with both CE loss and SGLD to ensure this distribution is being optimized with an unbiased objective. Similar to Du and Mordatch³⁵ we also adopt the contrastive divergence to estimate the expectation of derivative of the log-likelihood for a single example \mathbf{x} with respect to \mathbf{w} . Since it gives an order of magnitude savings in computation compared to seeding new chains at each iteration as in Nijkamp et al.³²

Step 2) Network retraining Fix $\hat{\mathbf{Y}}_T$ and minimize

$$-\sum_{s \in S} \sum_{k=1}^K y_s^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_s) - \sum_{t \in T} \sum_{k=1}^K \hat{y}_t^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_t) \quad (7)$$

w.r.t. \mathbf{w} . Carrying out steps 1) and 2) for one time is defined as one round in self-training.

For each class k , λ_k is determined by the confidence value selecting the most confident p portion of class k predictions in the entire target set.²⁴ If a sample's predication is relatively confident with $p(k^*|\mathbf{x}_i; \mathbf{w}) > \lambda_{k^*}$, it is selected and labeled as class $k^* = \arg \max_k \left\{ \frac{p(k|\mathbf{x}_i; \mathbf{w})}{\lambda_k} \right\}$. The less confident ones with $p(k^*|\mathbf{x}_i; \mathbf{w}) \leq \lambda_{k^*}$ are not selected. The same class-balanced λ_k strategy introduced in Zou et al.²⁴ is adopted for all self-training methods in this study.

3.4 | EBM loss for target domain

The essence of the EBM³¹ is a data-driven process that shapes the energy surface in such a way that the desired configurations get assigned low energies, while the incorrect ones are given high energies.

In this framework, the conventional supervised learning³⁷ for each \mathbf{x} in the training set, the energy of the pair $\{\mathbf{x}, \mathbf{y}\}$ takes low values when \mathbf{y} is the correct label and higher values for incorrect \mathbf{y} . Similarly, when modeling \mathbf{x} alone within an unsupervised learning setting, lower energy is attributed to the data manifold. Actually, the term contrastive sample is often used to refer to a data point causing an energy pull-up, such as the incorrect \mathbf{y} in supervised learning and points from low data density regions in the setting of unsupervised learning. Therefore, we can construct an efficient EBM loss for the target sample as

$$\mathcal{L}_{EBM} = -\log \sum_k \hat{y} \exp(f_{\mathbf{w}}(\mathbf{x})[k]) + \sum_k \hat{y}_t^{(k)} \log \lambda_k. \quad (8)$$

The soft pseudo-label \hat{y} used here is the original softmax prediction of the previous round classifier/segmentor without the one-hot processing. This is also essentially different from the soft-label used in Zou et al.¹⁰ which post-process the one-hot label to construct a smoothed one. Then, Equation (8) will be used to replace the objective in Equation (5). This configuration also dispenses with the validation of α .

This can be regarded as using \hat{y} as the weight to indicate which class's energy should be minimized. For the class with larger \hat{y} probability, the larger $f_{\mathbf{w}}(\mathbf{x})[k]$ will be emphasised. Therefore, the network will pron to produce the large $f_{\mathbf{w}}(\mathbf{x})[k]$ value at the pseudo-label's category.

With the soft pseudo-label setting, the confidence of each class prediction can also be encoded in the histogram distribution of \hat{y} . Although the additional information can enrich the training, the signal may be even noisier at the beginning. To avoid the model collapse, in practice, we propose a loss annealing scheme that starts with the EBM regularization and then smoothly changes to \mathcal{L}_{EBM} . The objective can be formulated as $\frac{1}{1+\beta} \{\mathcal{L}_{EBM} + \beta \mathcal{R}_{EBM}\}$ where β is a balance hyper-parameter and decrease β from 10 to 0 gradually in the training, that is, $\beta = \frac{10}{1+N^2}$ for $N \leq 5$, and $\beta = 0$ for $N > 5$, where N is the epoch number.

3.5 | Energy-based self-training as CEM

Amini and Gallinari³⁸ proves that semi-supervised logistic regression is a Classification Expectation-Maximization (CEM) to solve Classification Maximum Likelihood (CML). We generalize the statement for proposed energy-constrained self-training as follows.

Proposition 1. *Energy-constrained self-training can be a regularized Classification Expectation-Maximization (CEM) for solving a generalized Regularized Classification Maximum Likelihood (RCML) problem and is convergent with gradient descent.*

The problem of CML was first proposed for clustering tasks in Celeux and Govaert,³⁹ and can be solved through CEM. Compared with traditional expectation maximization (EM) that has an Expectation (E) step and a Maximization (M) step, CEM has an additional Classification (C) step (between E and M step) that assigns a sample to the cluster with maximal posterior probability. In Amini and Gallinari,³⁸ the CML is generalized for discriminant semi-supervised learning with both labeled and unlabeled data that is defined as follows:

$$\log \mathcal{L}_C = \log \tilde{\mathcal{L}}_C + \sum_{i \in S, T} \log p(\mathbf{x}_i), \quad (9)$$

where

$$\log \tilde{\mathcal{L}}_C = \sum_{s \in S} \sum_{k=1}^K y_s^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_s) + \sum_{t \in T} \sum_{k=1}^K \hat{y}_t^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_t). \quad (10)$$

Note that $\hat{y}_t \in \{0, 1\}^K, \forall t$. $p_{\mathbf{w}}(k|\mathbf{x}_t)$ is a posterior probability modeled by probabilistic classifier such as logistic classifier, neural network, and so forth. \mathbf{w} is the learnable weight. As mentioned in Amini and Gallinari,³⁸ when using a discriminant classifier, we make no assumption about the data distribution $p(\mathbf{x}_t)$. Thus maximizing Equation (9) is equal to maximizing Equation (10). The CEM is used to solve Equation (10) via alternating E-step: estimating the posterior probability $p_{\mathbf{w}}(k|\mathbf{x}_t)$; C-step: assigning the pseudo-labels according to the maximal posterior probability $p_{\mathbf{w}}(k|\mathbf{x}_t)$; M-step: maximizing the log-likelihood in terms of \mathbf{w} .

For Proposition 1, we show the inherent connections between the energy-based self-training and RCML. We also show that the proposed alternating optimization is essentially a Classification Expectation-Maximization (CEM) toward maximizing RCML. The energy-based self-training minimization problem can be seen as the following maximization problem:

$$\begin{aligned} \max_{\mathbf{w}, \hat{\mathbf{Y}}_T} \log \tilde{\mathcal{L}}_C + EBM &= \sum_{s \in S} \sum_{k=1}^K y_s^{(k)} \log(p_{\mathbf{w}}(k|\mathbf{x}_s)) \\ &+ \sum_{t \in T} \left[\sum_{k=1}^K \hat{y}_t^{(k)} \log(p_{\mathbf{w}}(k|\mathbf{x}_t)) - \alpha E_{\mathbf{w}}(\mathbf{x}_t) \right] \\ \text{s.t. } \hat{\mathbf{y}}_t &\in \Delta^{(K-1)} \cup \{\mathbf{0}\}, \forall t \end{aligned} \quad (11)$$

where $\alpha > 0$. $\hat{\mathbf{y}}_t \in \Delta^{K-1} = \{(\hat{y}_t^{(1)}, \dots, \hat{y}_t^{(K)}) : \hat{y}_t^{(k)} \geq 0, \sum_{k=1}^K \hat{y}_t^{(k)} = 1\}, \forall t$. (Δ^{K-1} denotes a probability simplex) Compared with Equation (10), the above problem can be regarded as a generalized Classification Maximization Likelihood ($\log \tilde{\mathcal{L}}_C$) with a energy regularizer (E), relaxing the one-hot vector feasible set to the probability simplex. The alternating optimization can be regarded as a CEM for solving this RCML problem:

E-Step: Given the model weight \mathbf{w} , estimate the posterior probability $p_{\mathbf{w}}(\mathbf{x}_t), \forall t$.

C-Step: Fix \mathbf{w} and solve the following problem for $\hat{\mathbf{Y}}_T$.

$$\begin{aligned} \max_{\hat{\mathbf{Y}}_T} \sum_{t \in T} \left[\sum_{k=1}^K \hat{y}_t^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_t) - \alpha E_{\mathbf{w}}(\mathbf{x}_t) \right] \\ \text{s.t. } \hat{\mathbf{y}}_t \in \Delta^{(K-1)} \cup \{\mathbf{0}\}, \forall t. \end{aligned} \quad (12)$$

M-Step: Fix $\hat{\mathbf{Y}}_T$ and use gradient ascent to solve the following problem for \mathbf{w} .

$$\max_{\mathbf{w}} \sum_{s \in S} \sum_{k=1}^K y_s^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_s) + \sum_{t \in T} \left[\sum_{k=1}^K \hat{y}_t^{(k)} \log p_{\mathbf{w}}(k|\mathbf{x}_t) - \alpha E_{\mathbf{w}}(\mathbf{x}_t) \right]. \quad (13)$$

Now we have shown energy-based self-training can be regarded as a CEM.

In the next subsection, we also provide the convergence analysis of energy-constrained self-training for UDA.

3.6 | Convergence analysis for energy-constrained self-training

Considering that our energy-constrained self-training follows an alternating optimization scheme, which is different from the conventional single round updating, we further explain that the energy-constrained self-training is convergent. We emphasize the proposed energy regularizer is convex w.r.t. either prediction distributions or pseudo-label distributions. The problem can be solved by alternatively taking the two steps. The training is convergent only if the loss in both of the two steps are nonincreasing.

Specifically, we have the following analysis for these two steps:

Loss in step (1) is nonincreasing: We prove that the solver 6 is a global minimum of optimization problem 5. Specifically, the feasible set of the problem is a union of Δ^{K-1} and $\{\mathbf{0}\}$. Thus, two sub-problems share the same objective with either a probability simplex and $\mathbf{0}$ as a feasible set. For the proposed convex energy regularizer, the first sub-problem is a convex problem with a unique solution (see Table 1 of the main paper), while the latter one has a global optimum solver at $\mathbf{0}$. Therefore, 6 is the solver for the combined two subproblems, making it the global optimum. Thus step 1 is nonincreasing and can be minimized with the global solver.

Loss in step (2) is nonincreasing: Due to the convexity of proposed model regularizers, with a proper learning rate, the gradient descent can monotonically decrease the regularized objective function of 7.

TABLE 1 Experimental results for VisDA17-val setting

Method	Aero	Bike	Bus	Car	Horse	Knife	Motor	Person	Plant	Skateboard	Train	Truck	Mean
Source-Res101 ⁴⁰	55.1	53.3	61.9	59.1	80.6	17.9	79.7	31.2	81.0	26.5	73.5	8.5	52.4
MMD ¹⁹	87.1	63.0	76.5	42.0	90.3	42.9	85.9	53.1	49.7	36.3	85.8	20.7	61.1
DANN ⁴¹	81.9	77.7	82.8	44.3	81.2	29.5	65.1	28.6	51.9	54.6	82.8	7.8	57.4
ENT ⁴²	80.3	75.5	75.8	48.3	77.9	27.3	69.7	40.2	46.5	46.6	79.3	16.0	57.0
MCD ⁴³	87.0	60.9	83.7	64.0	88.9	79.6	84.7	76.9	88.6	40.3	83.0	25.8	71.9
ADR ⁴⁴	87.8	79.5	83.7	65.3	92.3	61.8	88.9	73.2	87.8	60.0	85.5	32.3	74.8
DEV ⁴⁵	81.8	53.5	82.9	71.6	89.2	72.0	89.4	75.7	97.0	55.5	71.2	29.2	72.4
TDDA ⁴⁰	88.2	78.5	79.7	71.1	90.0	81.6	84.9	72.3	92.0	52.6	82.9	18.4	74.03
PANDA ⁴⁶	90.9	50.5	72.3	82.7	88.3	88.3	90.3	79.8	89.7	79.2	88.1	39.4	78.3
DMRL ⁴⁷	-	-	-	-	-	-	-	-	-	-	-	-	75.5
CBST ²⁴	87.1 ± 1.2	79.5 ± 2.3	58.3 ± 2.6	50.4 ± 3.9	82.8 ± 2.1	73.7 ± 7.2	80.9 ± 2.6	71.8 ± 3.1	81.6 ± 3.2	88.4 ± 3.3	75.2 ± 1.2	68.4 ± 3.4	74.8 ± 0.5
CBST + R _{EEM}	87.9 ± 1.6	79.6 ± 1.5	68.5 ± 1.3	68.6 ± 1.9	83.2 ± 1.2	78.4 ± 1.9	83.5 ± 1.5	72.2 ± 1.5	82.2 ± 1.6	84.3 ± 1.5	80.9 ± 1.4	67.5 ± 1.3	77.0 ± 0.6
CBST + L _{EEM}	87.8 ± 1.2	80.4 ± 1.9	69.3 ± 1.7	67.5 ± 1.6	83.2 ± 1.0	76.8 ± 1.4	81.6 ± 1.8	72.3 ± 1.5	81.6 ± 1.8	84.5 ± 1.3	81.7 ± 1.1	67.0 ± 1.4	76.8 ± 0.5
CRST ¹⁰	89.2 ± 1.6	79.6 ± 4.6	64.2 ± 4.0	57.8 ± 3.4	87.8 ± 1.9	79.6 ± 8.5	85.6 ± 2.6	75.9 ± 4.2	86.5 ± 2.2	85.1 ± 2.4	77.7 ± 2.2	68.5 ± 0.9	-
CRST + R _{EEM}	90.3 ± 1.5	82.6 ± 1.2	72.4 ± 1.5	71.7 ± 1.8	87.6 ± 1.8	80.5 ± 1.9	85.4 ± 1.5	80.8 ± 1.5	87.1 ± 1.6	89.9 ± 1.5	83.6 ± 1.6	71.5 ± 1.3	80.2 ± 0.5
CRST + L _{EEM}	91.9 ± 1.7	81.8 ± 1.8	72.8 ± 1.2	71.5 ± 1.3	88.4 ± 1.6	81.3 ± 1.3	89.3 ± 1.1	78.8 ± 1.2	87.2 ± 1.7	88.0 ± 1.5	85.3 ± 1.0	71.2 ± 1.6	80.4 ± 0.7
SimNet ⁴⁸	94.3	82.3	73.5	47.2	87.9	49.2	75.1	79.7	85.3	68.5	81.1	50.3	72.9
GTA ⁴⁹	-	-	-	-	-	-	-	-	-	-	-	-	77.1
CRST + R _{EEM} ^a	93.2 ± 1.3	85.8 ± 1.2	73.7 ± 1.0	74.3 ± 1.5	89.5 ± 0.6	87.6 ± 1.6	88.2 ± 1.6	82.2 ± 1.6	90.9 ± 1.3	91.6 ± 1.8	85.1 ± 1.4	79.9 ± 1.4	82.8 ± 0.5
CRST + L _{EEM} ^a	93.0 ± 1.6	84.4 ± 1.6	74.4 ± 1.0	73.5 ± 1.4	91.2 ± 1.5	82.4 ± 1.5	88.8 ± 1.6	81.8 ± 1.2	89.8 ± 1.2	90.8 ± 1.4	86.5 ± 1.2	77.1 ± 0.9	82.5 ± 0.5

Note: We use ResNet101 as backbone except SimNet and GTA. The best performance with the same backbone are bold.

^aResNet152 backbone.

Since the proposed energy regularization are lower-bounded by 0, it's easy to see that energy regularized self-training loss is lower-bounded by $\sum_{t \in T} \sum_{k=1, \dots, K} \log \lambda_k$. Thus energy-regularized self-training is convergent (Table 2).

4 | EXPERIMENTS

This section provides comprehensive evaluations of our proposed EBM-based regularization and loss on several UDA benchmarks, including image classification and semantic segmentation. We implement our methods with the PyTorch platform⁵⁴ on a V100GPU.

4.1 | Domain adaptation for image classification

We test on two challenging benchmarks: VisDA17¹¹ and Office-31.⁵⁵

The VisDA17¹¹ data set involves 12-class for UDA classification problem. Following the standard protocol as,^{10,49} the source domain utilizes the training set with 152, 409 synthetic 2D images, while the target domain uses the validation set with 55, 400 real images from the COCO data set.

For a fair comparison, we use the same backbone network for VisDA17 as the compared works, for example, ResNet-101 and ResNet-152. Moreover, the model was also pre-trained with the ImageNet data set as the previous works and fine-tuned in source domain by SGD.¹⁰ We use a fixed learning rate 1×10^{-3} , weight decay 5×10^{-4} , momentum 0.9 and batch size 32. From Figure 3, we can see that the performance is not sensitive to α when $\alpha \in [0.8, 1.1]$. Therefore, we simply choose $\alpha = 1$ consistently for all settings (Figures 2 and 3, Table 3).

TABLE 2 Experimental results for Office-31

Method	A → W	D → W	W → D	A → D	D → A	W → A	Mean accuracy
ResNet-50 ⁵⁰	68.4 ± 0.2	96.7 ± 0.1	99.3 ± 0.1	68.9 ± 0.2	62.5 ± 0.3	60.7 ± 0.3	76.1
DAN [ICML2015] ¹⁹	80.5 ± 0.4	97.1 ± 0.2	99.6 ± 0.1	78.6 ± 0.2	63.6 ± 0.3	62.8 ± 0.2	80.4
RTN [NeurIPS2016] ⁵¹	84.5 ± 0.2	96.8 ± 0.1	99.4 ± 0.1	77.5 ± 0.3	66.2 ± 0.2	64.8 ± 0.3	81.6
DANN [JMLR2016] ⁴¹	82.0 ± 0.4	96.9 ± 0.2	99.1 ± 0.1	79.7 ± 0.4	68.2 ± 0.4	67.4 ± 0.5	82.2
ADDA [CVPR2017] ⁵²	86.2 ± 0.5	96.2 ± 0.3	98.4 ± 0.3	77.8 ± 0.3	69.5 ± 0.4	68.9 ± 0.5	82.9
JAN [JMLR2017] ⁵³	85.4 ± 0.3	97.4 ± 0.2	99.8 ± 0.2	84.7 ± 0.3	68.6 ± 0.3	70.0 ± 0.4	84.3
GTA [CVPR2018] ⁴⁹	89.5 ± 0.5	97.9 ± 0.3	99.8 ± 0.4	87.7 ± 0.5	72.8 ± 0.3	71.4 ± 0.4	86.5
DMRL [ECCV2020] ⁴⁷	90.8 ± 0.3	99.0 ± 0.2	100.0 ± 0.0	93.4 ± 0.5	73.0 ± 0.3	71.2 ± 0.3	87.9
CBST [ECCV2018] ²⁴	87.8 ± 0.8	98.5 ± 0.1	100 ± 0.0	86.5 ± 1.0	71.2 ± 0.4	70.9 ± 0.7	85.8
CBST + R_{EBM}	89.4 ± 0.2	99.6 ± 0.1	100 ± 0.0	88.7 ± 0.9	73.8 ± 0.2	73.1 ± 0.2	87.4
CBST + \mathcal{L}_{EBM}	89.0 ± 0.4	99.6 ± 0.1	100 ± 0.0	88.4 ± 0.8	74.7 ± 0.2	73.0 ± 0.4	87.4
CRST [ICCV2019] ¹⁰	89.4 ± 0.7	98.9 ± 0.4	100 ± 0.0	88.7 ± 0.8	72.6 ± 0.7	70.9 ± 0.5	86.8
CRST + R_{EBM}	90.6 ± 0.4	99.7 ± 0.1	100 ± 0.0	91.0 ± 0.8	74.0 ± 0.6	74.0 ± 0.3	88.6
CRST + \mathcal{L}_{EBM}	90.4 ± 0.9	99.7 ± 0.1	100 ± 0.0	90.0 ± 0.9	73.7 ± 0.8	73.9 ± 0.4	88.3

Note: The best performance with the same backbone are bold.

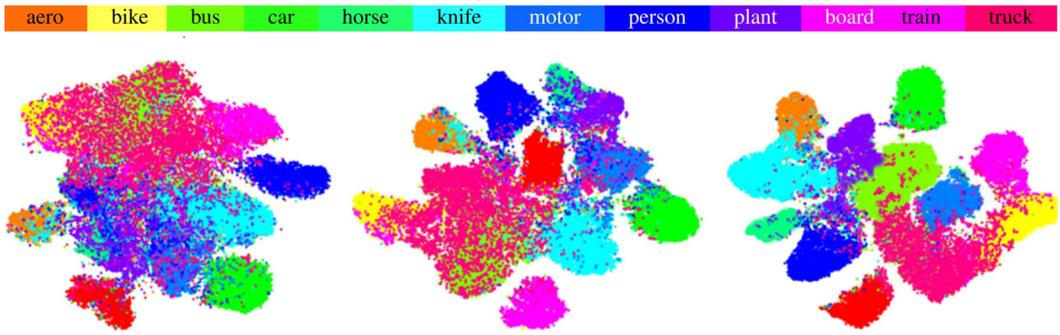


FIGURE 2 t-SNE visualization of learned VisDA17 features with source model, class-balanced self-training (CBST)²⁴ and CBST + R_{EBM} , respectively (from left to right) [Color figure can be viewed at wileyonlinelibrary.com]

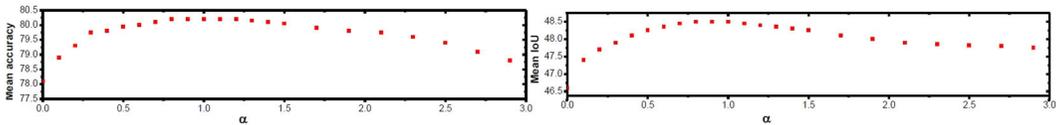


FIGURE 3 Sensitive analysis of hyper-parameter α in VisDA17 (top) and CTA52Sityscapes (bottom) with CRST + R_{EBM} [Color figure can be viewed at wileyonlinelibrary.com]

We present the results on VisDA 17 in Table 1 regarding per-class accuracy and mean accuracy. For each proposed approach, we independently run five times, and report the average accuracy and the corresponding standard deviation (\pm SD).

Among the compared methods, Class-balanced self-training (CBST) and Conservative regularized self-training (CRST)¹⁰ are the state-of-the-art self-training UDA methods. Our regularization can be simply added on CBST/CRST. We denote the CRST with EBM regularization with CRST + R_{EBM} .

Contributed by the additional EBMs objective, our proposed CBST + R_{EBM} and CRST + R_{EBM} can outperform CBST and CRST significantly. More appealingly, it can outperform the recent UDA methods other than self-training by a large margin. More appealing, the adversarial training can be simply added on to further boost the performance of self-training.¹⁰ Our proposed energy-based constraint is able to adaptively regularize the training w.r.t. the input \mathbf{x} and the present network parameters, which is more flexible than the manually pre-defined label smoothing.¹⁰

We can see that CRST + R_{EBM} achieves a better performance than the recent adversarial training,⁴⁰ dropout⁶³ and moment matching methods, which revokes the potential of self-training in UDA.

To demonstrate the generality of our EBM-based regularizer and loss for different backbones, we also tested with the more powerful backbones that have also been applied and shown better results.^{48,49} The EBM objective with the ResNet152 backbone outperforms the other state-of-the-arts.^{48,49} It also demonstrates the flexibility of R_{EBM} for different backbones.

Figure 2 visualizes the feature representations of VisDA-2017 classification. From it, we can clearly see an improved separation of target features with the EBM objective, resulting in better classification performance.

Office-31 is another standard UDA data set containing images belonging to 31 classes from three different domains, that is, Amazon (A), Webcam (W), and DSLR (D), each containing 2817, 795, and

TABLE 3 Experimental results for GTA5 to Cityscapes

Method	Base Net	Road	SW	Build	Wall	Fence	Pole	TL	TS	Veg.	Terrain	Sky	PR	Rider	Car	Truck	Bus	Train	Motor	Bike	mIoU
Source	DRN26	42.7	26.3	51.7	5.5	6.8	13.8	23.6	6.9	75.5	11.5	36.8	49.3	0.9	46.7	3.4	5.0	0.0	5.0	1.4	21.7
CyCADA ⁵⁶		79.1	33.1	77.9	23.4	17.3	32.1	33.3	31.8	81.5	26.7	69.0	62.8	14.7	74.5	20.9	25.6	6.9	18.8	20.4	39.5
Source	DRN105	36.4	14.2	67.4	16.4	12.0	20.1	8.7	0.7	69.8	13.3	56.9	37.0	0.4	53.6	10.6	3.2	0.2	0.9	0.0	22.2
MCD ⁴³		90.3	31.0	78.5	19.7	17.3	28.6	30.9	16.1	83.7	30.0	69.1	58.5	19.6	81.5	23.8	30.0	5.7	25.7	14.3	39.7
Source	PSPNet	69.9	22.3	75.6	15.8	20.1	18.8	28.2	17.1	75.6	8.00	73.5	55.0	2.9	66.9	34.4	30.8	0.0	18.4	0.0	33.3
DCAN ⁵⁷		85.0	30.8	81.3	25.8	21.2	22.2	25.4	26.6	83.4	36.7	76.2	58.9	24.9	80.7	29.5	42.9	2.50	26.9	11.6	41.7
Source	DeepLabv2	75.8	16.8	77.2	12.5	21.0	25.5	30.1	20.1	81.3	24.6	70.3	53.8	26.4	49.9	17.2	25.9	6.5	25.3	36.0	36.6
AdaptSegNet ⁵⁸		86.5	36.0	79.9	23.4	23.3	23.9	35.2	14.8	83.4	33.3	75.6	58.5	27.6	73.7	32.5	35.4	3.9	30.1	28.1	42.4
AdvEnt ⁵⁹		89.4	33.1	81.0	26.6	26.8	27.2	33.5	24.7	83.9	36.7	78.8	58.7	30.5	84.8	38.5	44.5	1.7	31.6	32.4	45.5
Source	DeepLabv2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	29.2
FCAN ⁶⁰		-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	46.6
Source	DeepLabv2	75.8	16.8	77.2	12.5	21.0	25.5	30.1	20.1	81.3	24.6	70.3	53.8	26.4	49.9	17.2	25.9	6.5	25.3	36.0	36.6
DPR ⁶¹		92.3	51.9	82.1	29.2	25.1	24.5	33.8	33.0	82.4	32.8	82.2	58.6	27.2	84.3	33.4	46.3	2.2	29.5	32.3	46.5
Source	DeepLabv2	73.8	16.0	66.3	12.8	22.3	29.0	30.3	10.2	77.7	19.0	50.8	55.2	20.4	73.6	28.3	25.6	0.1	27.5	12.1	34.2
PyCDA ⁶²		90.5	36.3	84.4	32.4	28.7	34.6	36.4	31.5	86.8	37.9	78.5	62.3	21.5	85.6	27.9	34.8	18.0	22.9	49.3	47.4
Source	DeepLabv2	71.3	19.2	69.1	18.4	10.0	35.7	27.3	6.8	79.6	24.8	72.1	57.6	19.5	55.5	15.5	15.1	11.7	21.1	12.0	33.8
CBST ²⁴		89.9	55.0	79.9	29.5	20.6	37.8	32.9	13.9	84.0	31.2	75.5	60.2	27.1	81.8	29.7	40.5	7.62	28.7	41.4	45.6
CBST+ R_{EBM}		91.1	53.9	80.6	31.6	21.0	40.4	35.0	19.8	86.6	35.9	76.4	63.3	31.4	83.0	22.5	38.6	24.2	32.2	39.4	47.8
CBST+ \mathcal{L}_{EBM}		92.7	56.2	79.6	27.1	21.7	39.3	35.5	24.2	85.3	38.0	77.9	63.6	29.1	84.8	25.3	39.6	17.8	34.8	43.4	48.2
Source	DeepLabv2	71.3	19.2	69.1	18.4	10.0	35.7	27.3	6.8	79.6	24.8	72.1	57.6	19.5	55.5	15.5	15.1	11.7	21.1	12.0	33.8
CRST ¹⁰		89.0	51.2	79.4	31.7	19.1	38.5	34.1	20.4	84.7	35.4	76.8	61.3	30.2	80.7	27.4	39.4	10.2	32.2	43.3	46.6
CRST+ R_{EBM}		92.5	56.6	80.9	26.2	20.5	40.2	35.3	24.4	86.6	37.3	77.5	63.4	30.5	81.3	28.8	39.2	20.6	33.5	41.3	48.5
CRST+ \mathcal{L}_{EBM}		92.7	55.1	83.6	34.5	20.4	39.4	36.5	19.6	85.7	35.7	78.3	64.2	30.1	84.5	32.0	40.4	28.0	33.9	44.2	49.0

Note: The best performance with the same backbone are bold.

TABLE 4 Experimental results of SYNTHIA Cityscapes

Method	Base Net	Road	SW	Build	Wall*	Fence*	Pole*	TL	TS	Veg.	Sky	PR	Rider	Car	Bus	Motor	Bike	mIoU	mIoU*
Source	DRN105	14.9	11.4	58.7	1.9	0.0	24.1	1.2	6.0	68.8	76.0	54.3	7.1	34.2	15.0	0.8	0.0	23.4	26.8
MCD ⁴³		84.8	43.6	79.0	3.9	0.2	29.1	7.2	5.5	83.8	83.1	51.0	11.7	79.9	27.2	6.2	0.0	37.3	43.5
Source	PSPNet	56.0	24.6	76.5	5.0	0.2	19.0	5.7	7.8	77.5	78.9	44.7	7.70	35.3	7.9	1.5	24.0	29.5	34.5
DCAN ⁶⁴		82.8	36.4	75.7	5.1	0.1	25.8	8.0	18.7	74.7	76.9	51.1	15.9	77.7	24.8	4.11	37.3	38.4	44.9
Source only	DeepLabv2	55.6	23.8	74.6	—	—	—	6.1	12.1	74.8	79.0	55.3	19.1	39.6	23.3	13.7	25.0	—	38.6
AdaptSegNet ⁵⁸		84.3	42.7	77.5	—	—	—	4.7	7.0	77.9	82.5	54.3	21.0	72.3	32.2	18.9	32.3	—	46.7
Source only	ResNet-38	32.6	21.5	46.5	4.8	0.1	26.5	14.8	13.1	70.8	60.3	56.6	3.5	74.1	20.4	8.9	13.1	29.2	33.6
CBST ²⁴		53.6	23.7	75.0	12.5	0.3	36.4	23.5	26.3	84.8	74.7	67.2	17.5	84.5	28.4	15.2	55.8	42.5	48.4
CBST + R_{EBM}		77.8	37.8	75.1	17.0	10.8	36.3	20.9	32.8	78.2	79.4	55.7	21.9	82.2	13.0	20.6	45.3	43.6	49.6
CBST + \mathcal{L}_{EBM}		65.7	26.4	78.1	19.0	3.2	31.5	21.9	32.5	81.4	78.4	56.8	29.8	82.3	29.7	21.0	45.4	42.8	49.8
Source only	DeepLabv2	64.3	21.3	73.1	2.4	1.1	31.4	7.0	27.7	63.1	67.6	42.2	19.9	73.1	15.3	10.5	38.9	34.9	40.3
CBST ²⁴		68.2	32.0	76.6	15.1	2.5	29.2	20.7	33.8	73.2	78.1	60.0	25.0	83.5	15.1	20.2	40.2	42.1	48.8
CRST ¹⁰		65.4	27.7	80.1	18.3	0.7	38.1	22.4	30.8	83.2	80.0	60.6	27.5	82.6	33.5	19.0	42.8	44.6	50.4
CRST + R_{EBM}		67.6	24.9	78.9	18.0	9.2	36.9	26.9	35.5	82.4	79.4	59.5	29.8	85.0	29.9	22.0	47.4	45.7	51.7
CRST + \mathcal{L}_{EBM}		64.6	30.3	78.7	16.8	5.8	39.0	25.2	31.8	85.6	80.7	64.5	29.3	83.5	35.3	20.1	44.5	45.6	51.8

Note: The best performance with the same backbone are bold.

498 images, respectively. We follow the standard protocol in Sankaranarayanan and Colleagues,^{49,55} and evaluate on six transfer tasks $A \rightarrow W$, $D \rightarrow W$, $W \rightarrow D$, $A \rightarrow D$, $D \rightarrow A$, and $W \rightarrow A$.

We compare the performances of different methods on Office-31 with the same basenet ResNet-50 in Table 4. CBST or CRST + R_{EBM} achieves the best performance that outperforms their baselines CBST and CRST.

Actually, our R_{EBM} can be orthogonal with the recent advanced self-training UDA methods. CBST²⁴ is a pioneer of vanilla adversarial UDA, and its performance on VisDA17 is reported on Zou et al.¹⁰ We note that the label smoothing or entropy regularization proposed in CRST¹⁰ or the vanilla self-training UDA²⁴ can be simply add-on our R_{EBM} to further improve the performance. The proposed CRST + R_{EBM} obtains better or competitive performances in all of the settings. Especially, it is able to outperform the sophisticated generative pixel-level domain adaptation method GTA in both architecture and objectives.

4.2 | UDA for semantic segmentation

The typical deep semantic segmentation is essentially making the pixel-wise classification. We consider the challenging segmentation adaptation settings in the synthetic-to-real scenario:

- 1) GTA5⁶⁵ to Cityscapes.³ There are 19 shared classes between these two datasets. GTA5 data set has 24,966 annotated images with the size of 1052×1914 , which are rendered by the GTA5 game engine.
- 2) SYNTHIA⁶⁶ to Cityscapes.³ There are 16 shared classes between these two data sets. We use the SYNTHIA-RAND-CITYSCAPES set, which consists of 9400 labeled images. Following the standard protocols in Hoffman and Colleagues,^{56,58} we use the full set of SYNTHIA/GTA5, and propose to adapt the segmentation model to the Cityscapes training set with 2975 images. In testing, we evaluate on the Cityscapes validation set with 500 images.

To make a fair comparison with the other methods, we use the ResNet101 as the backbone network as Zou and Colleagues.^{10,24} Noticing that in PSPNET,⁶⁷ Wide ResNet38 is a stronger

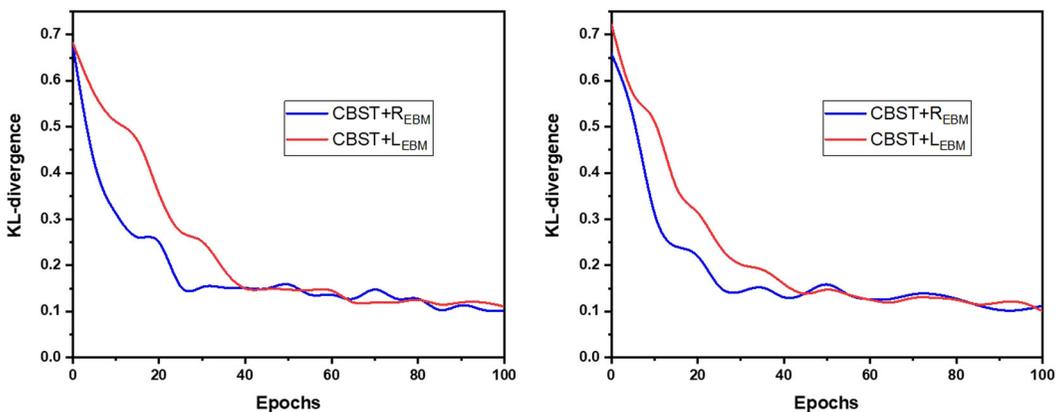


FIGURE 4 The KL-divergence between the predicted $p_t(y)$ and the ground of truth $p_t(y)$ in the testing stage. Left: VisDA17 task with CBST backbone, right: GTA5 to Cityscapes task with CRST backbone. The small value indicates the more accurate estimation [Color figure can be viewed at wileyonlinelibrary.com]

basenet than ResNet101. The basenet is pre-trained in ImageNet and fine-tuned in source domain by SGD with learning rate 2.5×10^{-4} , weight decay 5×10^{-4} , momentum 0.9, batch size 2, patch size 512×1024 and data augmentation of multi-scale training (0.5 – 1.5) and horizontal flipping. All results on this data set in the main paper are unified to report the mIoU of the models at the end of the 6th epoch.

We compare CBST/CRST + R_{EBM} with the other methods in Table 3. Based on the previous self-training UDA methods CBST or CRST, our additional EBM objective outperforms CBST or CRST by about 2% w.r.t. the mean IoUs.

We achieve a new state-of-the-art, even compared with the generative pixel-level domain adaptation method GTA,⁶⁰ which is a relatively complex algorithm in both architecture and objectives. As shown in Figure 4, R_{EBM} converges stably along the iterations. The performance on SYNTHIA to Cityscapes is also consistent with the GTA5 to Cityscapes as shown in Table 4.

5 | CONCLUSIONS

In this paper, we introduced the EBM's objective into the self-training UDA. Considering the lack of the target label, we resort to the pseudo-label, which is usually noisy. The EBM optimization objective provides an additional signal that is independent of the pseudo-labels. This approach can be more promising in the UDA setting than in the supervised learning setting as it can be added as either a regularization term or combined with a confidence-aware soft pseudo-label in an EBM loss. In addition, our solution is orthogonal with the recent advances in self-training, and can be added in a plug-and-play manner without large computation and change in network structure. Extensive experiments on both UDA classification and semantic segmentation demonstrated the effectiveness and generality of our approach. More advanced EBM training and self-training methods such as Han et al.³⁰ can also be adapted to further improve the performance, which is subject to our future work.

DATA AVAILABILITY STATEMENT

This study does not involve new data collection. All of the data sets used in this study are public available.

ORCID

Xiaofeng Liu  <http://orcid.org/0000-0002-4514-2016>

ENDNOTE

¹ For example, class-wise adversarial learning, or discriminators taking network predictions as input.

REFERENCES

1. Wang Z, Du B, Guo Y. Domain adaptation with neural embedding matching. *IEEE Trans Neural Netw Learn Syst.* 2019;31(7):2387-2397.
2. Li S, Liu CH, Su L, et al. Discriminative transfer feature and label consistency for cross-domain image classification. *IEEE Trans Neural Netw Learn Syst.* 2020;31(11):4842-4856.
3. Cordts M, Omran M, Ramos S, et al. The cityscapes dataset for semantic urban scene understanding. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016:3213-3223.
4. Yang B, Ma AJ, Yuen PC. Learning domain-shared group-sparse representation for unsupervised domain adaptation. *Pattern Recogn.* 2018;81:615-632.

5. Yang B, Yuen PC. Learning adaptive geometry for unsupervised domain adaptation. *Pattern Recogn.* 2021;110:107638.
6. Liu X, Liu X, Hu B, et al. Subtype-aware unsupervised domain adaptation for medical diagnosis. AAAI.
7. Li R, Cao W, Jiao Q, Wu S, Wong HS. Simplified unsupervised image translation for semantic segmentation adaptation. *Pattern Recogn.* 2020;105:107343.
8. Jia J, Zhai L, Ren W, Wang L, Ren Y, Zhang L. Transferable heterogeneous feature subspace learning for JPEG mismatched steganalysis. *Pattern Recogn.* 2020;100:107105.
9. Lee DH. Pseudo-label: the simple and efficient semi-supervised learning method for deep neural networks. ICML Workshop on Challenges in Representation Learning; 2013.
10. Zou Y, Yu Z, Liu X, Kumar B, Wang J. Confidence regularized self-training. ICCV.
11. Peng X, Usman B, Kaushik N, et al. VisDA: A synthetic-to-real benchmark for visual domain adaptation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops; 2018:2021-2026.
12. Bagherinezhad H, Horton M, Rastegari M, Farhadi A. Label refinery: Improving imagenet classification through label progression. *arXiv preprint arXiv:1805.02641*.
13. Grathwohl W, Wang K, Jacobsen J, Duvenaud D, Norouzi M, Swersky K. Your classifier is secretly an energy based model and you should treat it like one. *arXiv*.
14. Liu X, Hu B, Liu X, Lu J, You J, Kong L. Energy-constrained self-training for unsupervised domain adaptation. 25th International Conference on Pattern Recognition (ICPR 2020).
15. Li J, Zhang B, Lu G, et al. Relaxed asymmetric deep hashing learning: Point-to-angle matching. *IEEE Trans Neural Netw Learn Syst.* 2019;31(11):4791-4805.
16. Che T, Liu X, Li S, et al. Deep verifier networks: Verification of deep discriminative models with deep generative models. *ArXiv*; 2019.
17. Liu X, Ji W, You J, Fakhri GE, Woo J. Severity-aware semantic segmentation with reinforced Wasserstein training. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2020: 12566-12575.
18. Glorot X, Bordes A, Bengio Y. Domain adaptation for large-scale sentiment classification: a deep learning approach. ICML; 2011:513-520.
19. Long M, Cao Y, Wang J, Jordan MI. Learning transferable features with deep adaptation networks. ICML; 2015.
20. Pinheiro PO. Unsupervised domain adaptation with similarity learning. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2018:8004-8013. doi:10.1109/CVPR.2018.00835
21. Tzeng E, Hoffman J, Darrell T, Saenko K. Simultaneous deep transfer across domains and tasks. 2015 IEEE International Conference on Computer Vision (ICCV); 2015:4068-4076. doi:10.1109/ICCV.2015.463
22. Kouw WM. An introduction to domain adaptation and transfer learning. *arXiv:1812.11806*.
23. He G, Liu X, Fan F, You J. Classification-aware semi-supervised domain adaptation. The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops; 2020.
24. Zou Y, Yu Z, VijayaKumar BVK, Wang J. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. ECCV; 2018.
25. Chen C, Xie W, Huang W, et al. Progressive feature alignment for unsupervised domain adaptation. CVPR; 2019.
26. Kang G, Jiang L, Yang Y, Hauptmann AG. Contrastive adaptation network for unsupervised domain adaptation. CVPR; 2019.
27. Gu X, Sun J, Xu Z. Spherical space domain adaptation with robust pseudo-label loss. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2020:9101-9110.
28. Triguero I, García S, Herrera F. Self-labeled techniques for semi-supervised learning: taxonomy, software and empirical study. *Knowl Inform Syst.* 2015;42(2):245-284.
29. Busto PP, Iqbal A, Gall J. Open set domain adaptation for image and action recognition. *IEEE Trans Pattern Anal Mach Int.* 2018;42(2):413-429.
30. Han L, Zou Y, Gao R, Wang L, Metaxas D. Unsupervised domain adaptation via calibrating uncertainties. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops; 2019: 99-102.
31. LeCun Y, Chopra S, Hadsell R, Ranzato M, Huang F. A tutorial on energy-based learning. In *Predicting Structured Data*. Vol 1.

32. Nijkamp E, Zhu SC, Wu YN. On learning non-convergent short-run MCMC toward energy-based model. *arXiv*.
33. Welling M, Teh YW. Bayesian learning via stochastic gradient Langevin dynamics. ICML; 2011.
34. Xie J, Lu Y, Zhu SC, Wu Y. A theory of generative convnet. ICML; 2016:2635-2644.
35. Du Y, Mordatch I. Implicit generation and generalization in energy-based models. *arXiv*.
36. Song Y, Ou Z. Learning neural random fields with inclusive auxiliary generators. *arXiv*.
37. LeCun Y, Huang FJ. Loss functions for discriminative training of energy-based models. *AIStats*; Vol 6; 2005: 34.
38. Amini MR, Gallinari P. Semi-supervised logistic regression. ECAI; 2002.
39. Celeux G, Govaert G. A classification EM algorithm for clustering and two stochastic versions. *Comput Stat Data Anal*. 1992;14(3):315-332.
40. Gholami B, Sahu P, Kim M, Pavlovic V. Task-discriminative domain alignment for unsupervised domain adaptation. ICCV; 2019.
41. Ganin Y, Ustinova E, Ajakan H, et al. Domain-adversarial training of neural networks. *J Mach Learn Res*. 2016;17(1):2096-2030.
42. Grandvalet Y, Bengio Y. Semi-supervised learning by entropy minimization. NIPS; 2005.
43. Saito K, Watanabe K, Ushiku Y, Harada T. Maximum classifier discrepancy for unsupervised domain adaptation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2018: 3723-3732.
44. Saito K, Ushiku Y, Harada T, Saenko K. Adversarial dropout regularization. ICLR; 2018.
45. You K, Wang X, Long M, Jordan M. Towards accurate model selection in deep unsupervised domain adaptation. ICML; 2019.
46. Hu D, Liang J, Hou Q, et al. PANDA: Prototypical unsupervised domain adaptation. ECCV.
47. Wu Y, Inkpen D, El-Roby A. Dual mixup regularized learning for adversarial domain adaptation. ECCV.
48. Pinheiro PO. Unsupervised domain adaptation with similarity learning. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2018:8004-8013.
49. Sankaranarayanan S, Balaji Y, Castillo CD, Chellappa R. Generate to adapt: Aligning domains using generative adversarial networks. CVPR; 2018.
50. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2016:770-778.
51. Long M, Zhu H, Wang J, Jordan MI. Unsupervised domain adaptation with residual transfer networks Advances in Neural Information Processing Systems. 2016:136-144.
52. Tzeng E, Hoffman J, Saenko K, Darrell T. Adversarial discriminative domain adaptation. CVPR; 2017.
53. Long M, Zhu H, Wang J, Jordan MI. Deep transfer learning with joint adaptation networks. Proceedings of the 34th International Conference on Machine Learning-Volume 70. JMLR. org; 2017: 2208-2217.
54. Paszke A, Gross S, Chintala S, et al. Automatic differentiation in PyTorch.
55. Saenko K, Kulis B, Fritz M, Darrell T. Adapting visual category models to new domains. European Conference on Computer Vision. Springer; 2010:213-226.
56. Hoffman J, Tzeng E, Park T, et al. CyCADA: Cycle-consistent adversarial domain adaptation. ICML; 2018.
57. Wu Z, Han X, Lin YL, et al. DCAN: Dual channel-wise alignment networks for unsupervised scene adaptation. The European Conference on Computer Vision (ECCV); 2018.
58. Tsai YH, Hung WC, Schuler S, Sohn K, Yang MH, Chandraker M. Learning to adapt structured output space for semantic segmentation. CVPR; 2018.
59. Vu TH, Jain H, Bucher M, Cord M, Pérez P. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. CVPR; 2019.
60. Zhang Y, Qiu Z, Yao T, Liu D, Mei T. Fully convolutional adaptation networks for semantic segmentation. CVPR; 2018.
61. Tsai YH, Sohn K, Schuler S, Chandraker M. Domain adaptation for structured output via discriminative representations. ICCV.
62. Lian Q, Lv F, Duan L, Gong B. Constructing self-motivated pyramid curriculums for cross-domain semantic segmentation: a non-adversarial approach. ICCV.
63. Saito K, Ushiku Y, Harada T, Saenko K. Adversarial Dropout Regularization. *arXiv preprint arXiv:1711.01575*.

64. Wu Z, Han X, Lin YL, et al. DCAN: Dual channel-wise alignment networks for unsupervised scene adaptation. Proceedings of the European Conference on Computer Vision (ECCV); 2018:518-534.
65. Richter SR, Vineet V, Roth S, Koltun V. Playing for data: ground truth from computer games. ECCV; 2016.
66. Ros G, Sellart L, Materzynska J, Vazquez D, Lopez AM. The SYNTHIA dataset: a large collection of synthetic images for semantic segmentation of urban scenes. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2016:3234-3243.
67. Zhao H, Shi J, Qi X, Wang X, Jia J. Pyramid scene parsing network. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2017:2881-2890.

How to cite this article: Kong L, Hu B, Liu X, Lu J, You J, Liu X. Constraining pseudo-label in self-training unsupervised domain adaptation with energy-based model. *Int J Intell Syst.* 2022;37:8092-8112. doi:10.1002/int.22930