

GRADMIX: MULTI-SOURCE TRANSFER ACROSS DOMAINS AND TASKS

Anonymous authors

Paper under double-blind review

ABSTRACT

The machine learning and computer vision community is witnessing an unprecedented rate of new tasks being proposed and addressed, thanks to the power of deep convolutional networks to find complex mappings from \mathcal{X} to \mathcal{Y} . The advent of each task often accompanies the release of a large-scale human-labeled dataset, for supervised training of the deep network. However, it is expensive and time-consuming to manually label sufficient amount of training data. Therefore, it is important to develop algorithms that can leverage off-the-shelf labeled dataset to learn useful knowledge for the target task. While previous works mostly focus on transfer learning from a single source, we study multi-source transfer across domains and tasks (MS-DTT), in a semi-supervised setting. We propose GradMix, a model-agnostic method applicable to any model trained with gradient-based learning rule. GradMix transfers knowledge via gradient descent, by weighting and mixing the gradients from all sources during training. Our method follows a meta-learning objective, by assigning layer-wise weights to the source gradients, such that the combined gradient follows the direction that can minimize the loss for a small set of samples from the target dataset. In addition, we propose to adaptively adjust the learning rate for each mini-batch based on its importance to the target task, and a pseudo-labeling method to leverage the unlabeled samples in the target domain. We perform experiments on two MS-DTT tasks: digit recognition and action recognition, and demonstrate the advantageous performance of the proposed method against multiple baselines.

1 INTRODUCTION

Deep convolutional networks (ConvNets) have significantly improved the state-of-the-art for visual recognition, by finding complex mappings from \mathcal{X} to \mathcal{Y} . Unfortunately, these impressive gains in performance come only when massive amounts of paired labeled data (x, y) s.t. $x \in \mathcal{X}, y \in \mathcal{Y}$ are available for supervised training. For many application domains, it is often prohibitive to manually label sufficient training data, due to the significant amount of human efforts involved. Hence, there is strong incentive to develop algorithms that can reduce the burden of manual labeling, typically by leveraging off-the-shelf labeled datasets from other related domains and tasks.

There has been a large amount of efforts in the research community to address adapting deep models across domains (Ganin & Lempitsky, 2015; Long et al., 2016; Tzeng et al., 2017), to transfer knowledge across tasks (Luo et al., 2017; He et al., 2017; Zamir et al., 2018), and to learn efficiently in a few shot manner (Finn et al., 2017; Ravi & Larochelle, 2017; Ren et al., 2018a). However, most works focus on a single-source and single-target scenario. Recently, some works (Xu et al., 2018; Mancini et al., 2018) propose deep approaches for multi-source domain adaptation, but they assume that the source and target domains have shared label space (task).

In many computer vision applications, there often exist multiple labeled datasets available from different domains and/or tasks related to the target application. Hence, it is important and practically valuable that we can transfer knowledge from as many source datasets as possible. In this work, we formalize this problem as multi-source domain and task transfer (MS-DTT). Given a set of labeled source dataset, $\mathcal{S} = \{S_1, S_2, \dots, S_k\}$, we aim to transfer knowledge to a sparsely labeled target dataset \mathcal{T} . Each source dataset S_i could come from a different domain compared to \mathcal{T} , or from a

different task, or different in both domain and task. We focus on a semi-supervised setting, where only few samples in \mathcal{T} have labels.

Most works achieve domain transfer by aligning the feature distribution of source domain and target domain (Long et al., 2015b; 2016; Ganin & Lempitsky, 2015; Tzeng et al., 2015; Mancini et al., 2018; Xu et al., 2018). However, this method could be suboptimal for MS-DTT. The reason is that in MS-DTT, the distribution of source data $p(x^{S_i}, y^{S_i})$ and target data $p(x^{\mathcal{T}}, y^{\mathcal{T}})$ could be significantly different in both input space and label space, thus simply aligning their input space may generate indiscriminative features for the target classes. In addition, feature alignment introduces additional layers and loss terms, which require careful design to perform well.

In this work, we propose a generic and scalable method, namely GradMix, for semi-supervised MS-DTT. GradMix is a model-agnostic method, applicable to any model that uses gradient-based learning rule. Our method does not introduce extra layers or loss functions for feature alignment. Instead, we perform knowledge transfer via gradient descent, by weighting and mixing the gradients from all the source datasets during training. We follow a meta-learning paradigm and model the most basic assumption: *the combined gradient should minimize the loss for a set of unbiased samples from the target dataset*. We propose an online method to weight and mix the source gradients at each training iteration, such that the knowledge most useful for the target task is preserved through the gradient update. Our method can adaptively adjust the learning rate for each mini-batch based on its importance to the target task. In addition, we propose a pseudo-labeling method based on model ensemble to learn from the unlabeled data in target domain. We perform extensive experiments on two sets of MS-DTT task, including digit recognition and action recognition, and demonstrate the advantageous performance of the proposed method compared to multiple baselines. Our code is available at <https://www.url.com>.

2 RELATED WORK

Domain Adaptation. Domain adaptation seeks to learn from source domain a well-performing model on the target domain, by addressing the domain shift problem (Csurka, 2017). Most existing works focus on aligning the feature distribution of the source domain and target domain. Several works attempt to learn domain-invariant features by minimizing Maximum Mean Discrepancy (Long et al., 2015b; 2016; Sun & Saenko, 2016). Another class of method uses adversarial discriminative models, *i.e.* learn domain-agnostic representations by maximizing a domain confusion loss (Ganin & Lempitsky, 2015; Tzeng et al., 2015; Luo et al., 2017). Recently, multi-source domain adaptation with deep model has been studied. Mancini et al. (2018) use DA-layers (Carlucci et al., 2017; Li et al., 2017b) to minimize the distribution discrepancy of network activations. Xu et al. (2018) propose multi-way adversarial domain discriminator that minimizes the domain discrepancies between the target and each of the sources. However, both methods (Mancini et al., 2018; Xu et al., 2018) assume that the source and target domains have a shared label space.

Transfer Learning. Transfer learning extends domain adaptation into more general cases, where the source and target domain could be different in both input space and label space (Pan & Yang, 2010; Weiss et al., 2016). In computer vision, transfer learning has been widely studied to overcome the deficit of labeled data by adapting models trained for other tasks. With the advance of deep supervised learning, ConvNets trained on large datasets such as ImageNet (Russakovsky et al., 2015) have achieved state-of-the-art performance when transferred to other tasks (*e.g.* object detection (He et al., 2017), semantic segmentation (Long et al., 2015a), image captioning (Donahue et al., 2015), etc.) by simple fine-tuning. In this work, we focus on the setting where source and target domains have the same input space and different label spaces.

Meta-Learning. Meta-learning aims to utilize knowledge from past experiences to learn quickly on target tasks, from only a few annotated samples. Meta-learning generally seeks performing the learning at a level higher than where conventional learning occurs, *e.g.* learning the update rule of a learner (Ravi & Larochelle, 2017), or finding a good initialization point that can be easily fine-tuned (Finn et al., 2017). Recently Li et al. (2018) propose a meta-learning method to train models with good generalization ability to novel domains. Our method follows the meta-learning paradigm that uses validation loss as the meta-objective. Our method also resembles the example reweighting method by Ren et al. (2018b). However, they reweight samples in a batch for robust learning against noise, whereas we reweight source domain gradients layer-wise for transfer learning.

3 METHOD

3.1 PROBLEM FORMULATION

We first formally introduce the semi-supervised MS-DTT problem. We assume that there exists a set of k source domains $\mathcal{S} = \{S_1, S_2, \dots, S_k\}$, and a target domain \mathcal{T} . Each source domain S_i contains N^{S_i} images, $x^{S_i} \in \mathcal{X}^{S_i}$, with associated labels $y^{S_i} \in \mathcal{Y}^{S_i}$. Similarly, the target domain consists of $N^{\mathcal{T}}$ unlabeled images, $x^{\mathcal{T}} \in \mathcal{X}^{\mathcal{T}}$, as well as $M^{\mathcal{T}}$ labeled images, with associated labels $y^{\mathcal{T}} \in \mathcal{Y}^{\mathcal{T}}$. We assume that the target domain is only sparsely labeled, *i.e.* $M^{\mathcal{T}} \ll N^{\mathcal{T}}$. Our goal is to learn a strong target classifier that can predict labels $y^{\mathcal{T}}$ given $x^{\mathcal{T}}$.

Different from standard domain adaptation approaches that assume a shared label space between source and target domain ($\mathcal{Y}^{\mathcal{S}} = \mathcal{Y}^{\mathcal{T}}$), we study the problem of joint transfer across domains and tasks. Each source domain could have a partially overlapping label space with the target domain ($\mathcal{Y}^{S_i} \cap \mathcal{Y}^{\mathcal{T}} \subset \mathcal{Y}^{\mathcal{T}}$ and $\mathcal{Y}^{S_i} \cap \mathcal{Y}^{\mathcal{T}} \neq \emptyset$), or a non-overlapping label space ($\mathcal{Y}^{S_i} \cap \mathcal{Y}^{\mathcal{T}} = \emptyset$). However, we presume that at least one source domain should have the same label space as the target domain ($\exists S_i$ s.t. $\mathcal{Y}^{S_i} = \mathcal{Y}^{\mathcal{T}}$).

3.2 META-LEARNING OBJECTIVE

Let Θ denote the network parameters for our model. We consider a loss function $\mathcal{L}(x, y; \Theta) = f(\Theta)$ to minimize during training. For deep networks, SGD or its variants are commonly used to optimize the loss functions. At every step n of training, we forward a mini-batch of samples from each of the source domain $\{S_i\}_{i=1}^k$, and apply back-propagation to calculate the gradients w.r.t the parameters $\Theta_n, \nabla f_{s_i}(\Theta_n)$. The parameters are then adjusted according to the sum of the source gradients. For example, for vanilla SGD:

$$\Theta_{n+1} = \Theta_n - \alpha \sum_{i=1}^k \nabla f_{s_i}(\Theta_n), \quad (1)$$

where α is the learning rate.

In semi-supervised MS-DTT, we also have a small validation set \mathcal{V} that contains few labeled samples from the target domain. We want to learn a set of weights for the source gradients, $w = \{w_{s_i}\}_{i=1}^k$, such that when taking a gradient descent using their weighted combination $\sum_{i=1}^k w_{s_i} \nabla f_{s_i}(\Theta_n)$, the loss on the validation set is minimized:

$$\Theta^*(w) = \Theta_n - \alpha \sum_{i=1}^k w_{s_i} \nabla f_{s_i}(\Theta_n), \quad (2)$$

$$w^* = \arg \min_{w, w \geq 0} f_{\mathcal{V}}(\Theta^*(w)) \quad (3)$$

3.3 LAYER-WISE GRADIENT WEIGHTING

Calculating the optimal w^* requires two nested loops of optimization, which can be computationally expensive. Here we propose an approximation to the above objective. At each training iteration n , we do a forward-backward pass using the small validation set \mathcal{V} to calculate the gradient, $\nabla f_{\mathcal{V}}(\Theta_n)$. We take a first-order approximation and assume that adjusting Θ_n in the direction of $\nabla f_{\mathcal{V}}(\Theta_n)$ can minimize $f_{\mathcal{V}}(\Theta_n)$. Therefore, we find the optimal w^* by maximizing the cosine similarity between the combined source gradient and the validation gradient:

$$w^* = \arg \max_{w, w \geq 0} \left\langle \sum_{i=1}^k w_{s_i} \nabla f_{s_i}(\Theta_n), \nabla f_{\mathcal{V}}(\Theta_n) \right\rangle, \quad (4)$$

where $\langle a, b \rangle$ denotes the cosine similarity between vector a and b . This method is a cheap estimation for the meta-objective, which can also prevent the model from over-fitting to \mathcal{V} .

Instead of using a global weight value for each source gradient, we propose a layer-wise gradient weighting, where the gradient for each layer of the network are weighted separately. This enables a finer level of gradient combination. Specifically, in our MS-DTT setting, the source domains

and the target domain share the same parameters up to the last fully-connected (fc) layer, which is task-specific. Therefore, for each layer l with parameter θ^l , and for each source domain S_i , we have a corresponding weight $w_{s_i}^l$. We can then write Equation 4 as:

$$w^* = \arg \max_{w, w \geq 0} \sum_{l=1}^{L-1} \langle \sum_{i=1}^k w_{s_i}^l \nabla f_{s_i}(\theta_n^l), \nabla f_V(\theta_n^l) \rangle, \quad (5)$$

where L is the total number of layers for the ConvNet. We constrain $w_{s_i}^l \geq 0$ for all i and l , since negative gradient update can usually result in unstable behavior. To efficiently solve the above constrained non-linear optimization problem, we utilize a sequential quadratic programming method, SLSQP, implemented in NLOpt (Johnson).

In practice, we normalize the weights for each layer across all source domains so that they sum up to one:

$$\tilde{w}_{s_i}^l = \frac{w_{s_i}^l}{\sum_{i=1}^k w_{s_i}^l} \quad (6)$$

3.4 ADAPTIVE LEARNING RATE

Intuitively, certain mini-batches from the source domains contain more useful knowledge that can be transferred to the target domain, whereas some mini-batches contain less. Therefore, we want to adaptively adjust our training to pay more attention to the important mini-batches. To this end, we measure the importance score ρ of a mini-batch using the cosine similarity between the optimally combined gradient and the validation gradient:

$$\rho = \sum_{l=1}^{L-1} \langle \sum_{i=1}^k \tilde{w}_{s_i}^l \nabla f_{s_i}(\theta_n^l), \nabla f_V(\theta_n^l) \rangle \quad (7)$$

Based on ρ , we calculate a scaling term η bounded between 0 and 1:

$$\eta = \frac{1}{1 + e^{-\beta\rho - \gamma}}, \quad (8)$$

where β controls the rate of saturation for η , and γ defines the value of $\beta\rho$ where $\eta = 0.5$. We determine the value of β and γ empirically through experiments.

Finally, we multiply η to the learning rate α , and perform SGD to update the parameters:

$$\theta_{n+1}^l = \theta_n^l - \eta\alpha \sum_{i=1}^k \tilde{w}_{s_i}^l \nabla f_{s_i}(\theta_n^l), \text{ for } l = 1, 2, \dots, L - 1 \quad (9)$$

3.5 PSEUDO-LABEL WITH ENSEMBLES

In our semi-supervised MS-DTT setting, there also exists a large set of unlabeled images in the target domain, denoted as $\mathcal{U} = \{(x_n^T)\}_{n=1}^{N^T}$. We want to learn target-discriminative knowledge from \mathcal{U} . To achieve this, we propose a method to calculate pseudo-labels \hat{y}_n^T for the unlabeled images, and construct a pseudo-labeled dataset $S_u = \{(x_n^T, \hat{y}_n^T)\}_{n=1}^{N^P}$. Then we leverage S_u using the same gradient mixing method as described above. Specifically, we consider a loss $\mathcal{L}_u(x, \hat{y}; \Theta)$ to minimize during training, where $(x, \hat{y}) \in S_u$. At each training iteration n , we sample a mini-batch from S_u , calculate the gradient $\nabla f_{s_u}(\Theta_n)$, and combine it with the source gradients $\{\nabla f_{s_i}(\Theta_n)\}_{i=1}^k$ using the proposed layer-wise weighting method.

In order to acquire the pseudo-labels, we perform a first step to train a model using the source domain datasets following the proposed gradient mixing method, and use the learned model to label \mathcal{U} . However, the learned model would inevitably create some false pseudo-labels. Previous studies found that ensemble of models helps to produce more reliable pseudo-labels (Saito et al., 2017; Laine & Aila, 2017; Tarvainen & Valpola, 2017). Therefore, in our first step, we train multiple models with different combination of β and γ in Equation 8. Then we pick the top R models with the best accuracies on the hyper-validation set ($R = 3$ in our experiments), and use their ensemble to create



Figure 1: An illustration of the two experimental settings (digit recognition and action recognition) for multi-source domain and task transfer (MS-DTT). Our method effectively transfers knowledge from multiple sources to the target task.

pseudo-labels. The difference in hyper-parameters during training ensures that different models learn significantly different sets of weight, hence the ensemble of their prediction is less biased.

Here we propose two approaches to create pseudo-labels, namely hard label and soft label:

Hard label. In this approach, we assume that the pseudo-label is more likely to be correct if all the models can reach an agreement with high confidence. We assign a pseudo-label $\hat{y} = C$ to an image $x \in \mathcal{U}$, where C is a class number, if the two following conditions are satisfied. First, all of the R models should predict C as the class with the maximum probability. Second, for all models, the probability for C should exceed certain threshold, which is set as 0.8 in our experiments. If these two conditions are satisfied, we will add (x, \hat{y}) into S_u . During training, the loss $\mathcal{L}_u(x, \hat{y}; \Theta)$ is the standard cross entropy loss.

Soft label. Let p_r denote the output from the r -th model’s softmax layer for an input x , which represents the probability over classes. We calculate the average of p_r across all of the R pre-trained models as the soft pseudo-label for x : $\hat{y} = \frac{1}{R} \sum_{r=1}^R p_r$. Every unlabeled image $x \in \mathcal{U}$ will be assigned a soft label and added to S_u . During training, let p_Θ be the output probability from the model, we want to minimize the KL-divergence between p_Θ and the soft pseudo-label for all pairs $(x, \hat{y}) \in S_u$. Therefore, the loss is $\mathcal{L}_u(x, \hat{y}; \Theta) = D_{KL}(p_\Theta, \hat{y})$.

For both hard and soft label approach, after getting the pseudo-labels, we train a model from scratch using all available datasets $\{S_i\}_{i=1}^k, S_u$ and \mathcal{V} . Since the proposed gradient mixing method relies on \mathcal{V} to estimate the model’s performance on the target domain, we enlarge the size of \mathcal{V} to 100 samples per class, by adding hard-labeled images from S_u using the method described above. The enlarged \mathcal{V} can represent the target domain with less bias, which helps to calculate better weights on the source gradients, such that the model’s performance on the target domain is maximized.

4 EXPERIMENT

4.1 EXPERIMENTAL SETUP

Datasets. In our experiment we perform MS-DTT across two different groups of data settings, as shown in Figure 1. First we do transfer learning across different digit domains using MNIST (LeCun et al., 1998) and Street View House Numbers (SVHN) (Netzer et al., 2011). MNIST is the popular benchmark for handwritten digit recognition, which contains a training set of 60,000 examples, and a test set of 10,000 examples. SVHN is a real-world dataset consisting of images with colored background and blurred digits. It has 73,257 examples for training and 26,032 examples for testing.

For our second setup, we study MS-DTT from human activity images in MPII dataset (Andriluka et al., 2014) and human action images from the Web (BU101 dataset) (Ma et al., 2017), to video action recognition using UCF101 (Soomro et al., 2012) dataset. MPII dataset consists of 28,821 images covering 410 human activities including home activities, religious activities, occupation, etc. UCF101 is a benchmark action recognition dataset collected from YouTube. It consists of 13,320 videos from 101 action categories, captured under various lighting conditions with camera motion

Table 1: Classification accuracy (%) of the baselines and our method on the test split of MNIST 5-9. We report the mean and the standard error of each method across 10 runs with different randomly sampled \mathcal{V} .

Method	Datasets	k=2	k=3	k=4	k=5
Target only	\mathcal{V}	71.35±1.85	77.15±1.36	81.43±1.41	84.83±1.10
Source only	S_1, S_2	82.39	82.39	82.39	82.39
Fine-tune	S_1, S_2, \mathcal{V}	89.94±0.35	89.86±0.46	90.89±0.48	91.96±0.39
GradMix w/o AdaLR	S_1, S_2, \mathcal{V}	90.10±0.37	90.22±0.62	92.14±0.43	92.92±0.29
GradMix	S_1, S_2, \mathcal{V}	91.17±0.37	91.45±0.52	92.14±0.40	93.06±0.46
MDDA (Mancini et al., 2018)	$S_1, S_2, \mathcal{V}, \mathcal{U}$	90.23±0.40	90.28±0.50	91.45±0.37	91.85±0.31
DCTN (Xu et al., 2018)	$S_1, S_2, \mathcal{V}, \mathcal{U}$	91.81±0.26	92.34±0.28	92.42±0.39	92.97±0.37
GradMix w/ soft label	$S_1, S_2, \mathcal{V}, \mathcal{U}$	94.62±0.18	95.03±0.30	95.26±0.17	95.74±0.21
GradMix w/ hard label	$S_1, S_2, \mathcal{V}, \mathcal{U}$	96.02±0.24	96.24±0.33	96.63±0.17	96.84±0.20

and occlusion. We take the first split of UCF101 for our experiment. BU101 contains 23,800 images collected from the Web, with the same action categories as UCF101. It contains professional photos, commercial photos, and artistic photos, which can differ significantly from video frames.

Network and implementation details. For our first setting, we use the same ConvNet architecture as Luo et al. (2017), which has 4 Conv layers and 2 fc layers. We randomly initialize the weights, and train the network using SGD with learning rate $\alpha = 0.05$, and a momentum of 0.9. For fine-tuning we reduce the learning rate to 0.005. For our second setting, we use the ResNet-18 (He et al., 2016) architecture. We initialize the network with ImageNet pre-trained weights, which is important for all baseline methods to perform well. The learning rate is 0.001 for training and $5e-5$ for fine-tuning.

4.2 SVHN 5-9 + MNIST 0-4 \rightarrow MNIST 5-9

Experimental setting. In this experiment, we define four sets of training data: (1) labeled images of digits 5-9 from the training split of SVHN dataset as the first source S_1 , (2) labeled images of digits 0-4 from the training split of MNIST dataset as the second source S_2 , (3) few labeled images of digits 5-9 from the training split of MNIST dataset as the validation set \mathcal{V} , (4) unlabeled images from the rest of the training split of MNIST 5-9 as \mathcal{U} . We subsample k examples from each class of MNIST 5-9 to construct the unbiased validation set \mathcal{V} . We experiment with $k = 2, 3, 4, 5$, which corresponds to 10, 15, 20, 25 labeled examples. Since \mathcal{V} is randomly sampled, we repeat our experiment 10 times with different \mathcal{V} . In order to monitor training progress and tune hyper-parameters (e.g. α, β, γ), we split out another 1000 labeled samples from MNIST 5-9 as the hyper-validation set. The hyper-validation set is the traditional validation set, which is fixed across the 10 runs.

Baselines. We compare the proposed method to multiple baseline methods: (1) *Target only*: the model is trained using \mathcal{V} . (2) *Source only*: the model is trained using S_1 and S_2 without gradient reweighting. (3) *Fine-tune*: the *Source only* model is fine-tuned using \mathcal{V} . (4) MDDA (Mancini et al., 2018): Multi-domain domain alignment layers that shift the network activations for each domain using a parameterized transformation equivalent to batch normalization. (5) DCTN (Xu et al., 2018): Deep Cocktail Network, which uses multi-way adversarial adaptation to align the distribution of multiple source domains and the target domain.

We also evaluate different variants of our model with and without certain component to show its effect: (6) *GradMix w/o AdaLR*: the method in Section 3.3 without the adaptive learning rate (Section 3.4). (7) *GradMix*: the proposed method that uses S_1, S_2 and \mathcal{V} during training. (8) *GradMix w/ hard label*: using the hard label approach to create pseudo-labels for \mathcal{U} , and train a model with all available datasets. (9) *GradMix w/ soft label*: using the soft label approach to create pseudo-labels for \mathcal{U} , and train a model with all available datasets.

Results. Table 1 shows the results for methods described above. We report the mean and standard error of classification accuracy across 10 runs with randomly sampled \mathcal{V} . Methods in the upper part of the table do not use the unlabeled target domain data \mathcal{U} . Among these methods, the proposed GradMix has the best performance. If we remove the adaptive learning rate, the accuracy would decrease. As expected, the performance improves as k increases, which indicates more samples in




Figure 2: Loss on the hyper-validation set as training proceeds. We define 1 epoch as training for 100 mini-batches (gradient descents).

Table 2: Results of GradMix using different β and γ when $k = 3$. Numbers indicate the test accuracy (%) on MNIST 5-9 (averaged across 10 runs). The ensemble of the top three models is used to create pseudo-labels.

	$\gamma = 0$	$\gamma = 0.1$	$\gamma = 0.2$	$\gamma = 0.3$	$\gamma = 0.4$	$\gamma = 0.5$	$\gamma = 0.6$	$\gamma = 0.7$	$\gamma = 0.8$
$\beta = 5$	90.92	90.96	90.95	90.58	90.75	90.75	90.51	90.63	91.12
$\beta = 6$	90.41	90.75	89.95	90.79	90.59	89.95	90.58	90.63	90.56
$\beta = 7$	89.76	90.44	90.42	90.94	90.28	90.40	90.52	90.70	90.66
$\beta = 8$	90.05	90.89	90.93	90.57	90.77	90.69	89.99	90.58	90.71
$\beta = 9$	90.32	90.70	90.48	90.94	90.47	90.92	90.20	90.23	90.86
$\beta = 10$	90.52	90.03	89.67	90.01	89.84	90.51	91.45	90.58	90.70

Table 3: Results of GradMix w/ hard label using different number of pre-trained models for ensemble.

Num. of Models	k=2	k=3	k=4	k=5
$R = 1$	95.80±0.28	95.92±0.28	96.29±0.20	96.46±0.25
$R = 2$	95.30±0.23	95.80±0.18	96.18±0.22	96.57±0.26
$R = 3$	96.02±0.24	96.24±0.33	96.63±0.17	96.84±0.20

\mathcal{V} can help the GradMix method to better combine the gradients during training. The lower part of the table shows methods that use all available datasets including S_1 , S_2 , \mathcal{V} and \mathcal{U} . The proposed GradMix without \mathcal{U} can achieve comparable performance with state-of-the-art baselines that use \mathcal{U} (MDDA and DCTN). Using pseudo-label with model ensemble significantly improves performance compared to baseline methods. Comparing soft label to hard label, the hard label approach achieves better performance. More detailed results about model ensemble for pseudo-labeling is shown later in the ablation study.

Ablation Study. In this section, we perform multiple ablation experiments to demonstrate the effectiveness of our method and the effect of different hyper-parameters. First, Figure 2 shows two examples of the hyper-validation loss as training proceeds. We show the loss for the baseline *Source only* method and the proposed GradMix, where we perform hyper-validation every 100 mini-batches (gradient descents). In both examples with different k , GradMix achieves a quicker and steadier decrease in the hyper-validation loss.

In Table 2, we show the results using GradMix with different combination of β and γ when $k = 3$. We perform a grid search with $\beta = [5, 6, \dots, 10]$ and $\gamma = [0, 0.1, \dots, 0.8]$. The accuracy is the highest for $\beta = 10$ and $\gamma = 0.6$. The top three models are selected for ensemble to create pseudo-labels for the unlabeled set \mathcal{U} .

In addition, we perform experiments with different number of models used for ensemble. Table 3 shows the results for $R = 1, 2, 3$ across all values of k . $R = 1$ and $R = 2$ have comparable performance, whereas $R = 3$ performs better. This indicates that using the top three models for ensemble can create more reliable pseudo-labels.

Table 4: Classification accuracy (%) of the baselines and our method on the test split of UCF101. We report the mean accuracy of each method across two runs with different randomly sampled \mathcal{V} .

Method	Datasets	per-frame			per-video		
		k=3	k=5	k=10	k=3	k=5	k=10
Target only	\mathcal{V}	42.58	53.31	63.05	43.74	55.50	64.74
Source only	S_1, S_2	41.96	41.96	41.96	43.46	43.46	43.46
Fine-tune	S_1, S_2, \mathcal{V}	55.86	60.55	66.77	58.57	66.01	70.21
EnergyNet (Li et al., 2017a)	S_1, S_2, \mathcal{V}	55.93	60.82	66.73	58.70	66.23	70.25
GradMix	S_1, S_2, \mathcal{V}	56.25	61.73	67.30	59.41	66.27	71.49
MDDA (Mancini et al., 2018)	$S_1, S_2, \mathcal{V}, \mathcal{U}$	56.65	61.58	67.65	60.00	65.14	71.54
GradMix w/ hard label	$S_1, S_2, \mathcal{V}, \mathcal{U}$	68.92	68.76	69.25	72.58	72.34	73.48

4.3 MPII + BU101 \rightarrow UCF101

Experimental setting. In the action recognition experiment, we have four sets of training data similar to the digit recognition experiment, which include (1) S_1 : labeled images from the training split of MPII, (2) S_2 : labeled images from the training split of BU101, (3) \mathcal{V} : k labeled video clips per class randomly sampled from the training split of UCF101, (4) \mathcal{U} : unlabeled images from the rest of the training split of UCF101. We experiment with $k = 3, 5, 10$ which corresponds to 303, 505, 1010 video clips. Each experiment is run two times with different \mathcal{V} . We report the mean accuracy across the two runs for both per-frame classification and per-video classification. Per-frame classification is the same as doing individual image classification for every frame in the video, and per-video classification is done by averaging the softmax score for all the frames in a video as the video’s score.

Baselines. We compare our method with multiple baselines described in Section 4.2, including (1) *Target only*, (2) *Source only*, (3) *Fine-tune*. In addition, we evaluate another baseline for knowledge transfer in action recognition, namely (4) EnergyNet (Li et al., 2017a): The ConvNet (ResNet-18) is first trained on MPII and BU101, then knowledge is transferred to UCF101 through spatial attention maps using a Siamese Energy Network.

Results. Table 4 shows the results for action recognition. *Target only* has better performance compared to *Source only* even for $k = 3$, which indicates a strong distribution shift between source data and target data for actions in the wild. For all values of k , the proposed GradMix outperforms baseline methods that use S_1, S_2 and \mathcal{V} for training in both per-frame and per-video accuracy. GradMix also has comparable performance with MDDA that uses the unlabeled dataset \mathcal{U} . The proposed pseudo-label method achieves significant gain in accuracy by assigning hard labels to \mathcal{U} and learn target-discriminative knowledge from the pseudo-labeled dataset.

5 CONCLUSION

In this work, we propose GradMix, a method for semi-supervised MS-DTT: multi-source domain and task transfer. GradMix assigns layer-wise weights to the gradients calculated from each source objective, in a way such that the combined gradient can optimize the target objective, measured by the loss on a small validation set. GradMix can adaptively adjust the learning rate for each mini-batch based on its importance to the target task. In addition, we assign pseudo-labels to the unlabeled samples using model ensembles, and consider the pseudo-labeled dataset as a source during training. We validate the effectiveness our method with extensive experiments on two MS-DTT settings, namely digit recognition and action recognition. GradMix is a generic framework applicable to any models trained with gradient descent. For future work, we intend to extend GradMix to other problems where labeled data for the target task is expensive to acquire, such as image captioning.

REFERENCES

Mykhaylo Andriluka, Leonid Pishchulin, Peter V. Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, pp. 3686–3693, 2014.

- Fabio Maria Carlucci, Lorenzo Porzi, Barbara Caputo, Elisa Ricci, and Samuel Rota Bulò. Autodial: Automatic domain alignment layers. In *ICCV*, pp. 5077–5085, 2017.
- Gabriela Csurka. A comprehensive survey on domain adaptation for visual applications. In *Domain Adaptation in Computer Vision Applications*, pp. 1–35. Springer, 2017.
- Jeff Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Trevor Darrell, and Kate Saenko. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, pp. 2625–2634, 2015.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pp. 1126–1135, 2017.
- Yaroslav Ganin and Victor S. Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pp. 1180–1189, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pp. 770–778, 2016.
- Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *ICCV*, pp. 2980–2988, 2017.
- Steven G. Johnson. The NLOpt nonlinear-optimization package. URL <http://ab-initio.mit.edu/nlopt>.
- Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018.
- Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S. Kankanhalli. Attention transfer from web images for video recognition. In *ACM Multimedia*, pp. 1–9, 2017a.
- Yanhao Li, Naiyan Wang, Jianping Shi, Jiaying Liu, and Xiaodi Hou. Revisiting batch normalization for practical domain adaptation. In *ICLR*, 2017b.
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pp. 3431–3440, 2015a.
- Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation networks. In *ICML*, pp. 97–105, 2015b.
- Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. Unsupervised domain adaptation with residual transfer networks. In *NIPS*, pp. 136–144, 2016.
- Zelun Luo, Yuliang Zou, Judy Hoffman, and Fei-Fei Li. Label efficient learning of transferable representations across domains and tasks. In *NIPS*, pp. 164–176, 2017.
- Shugao Ma, Sarah Adel Bargal, Jianming Zhang, Leonid Sigal, and Stan Sclaroff. Do less and achieve more: Training cnns for action recognition utilizing action images from the web. *Pattern Recognition*, 68:334–345, 2017.
- Massimiliano Mancini, Lorenzo Porzi, Samuel Rota Bulò, Barbara Caputo, and Elisa Ricci. Boosting domain adaptation by discovering latent domains. In *CVPR*, pp. 3771–3780, 2018.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, 2011.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.*, 22(10):1345–1359, 2010.

- Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. In *ICLR*, 2018a.
- Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *ICML*, pp. 4331–4340, 2018b.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *ICML*, pp. 2988–2997, 2017.
- Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- Baochen Sun and Kate Saenko. Deep CORAL: correlation alignment for deep domain adaptation. In *ECCV Workshops*, pp. 443–450, 2016.
- Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, pp. 1195–1204, 2017.
- Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *ICCV*, pp. 4068–4076, 2015.
- Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pp. 2962–2971, 2017.
- Karl R. Weiss, Taghi M. Khoshgoftaar, and Dingding Wang. A survey of transfer learning. *Journal of Big Data*, 3:9, 2016.
- Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *CVPR*, pp. 3964–3973, 2018.
- Amir R. Zamir, Alexander Sax, William Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, pp. 3712–3722, 2018.