

Learning What and Where to Transfer

Yunhun Jang^{*1,2} Hankook Lee^{*1} Sung Ju Hwang^{3,4,5} Jinwoo Shin^{1,4,5}

Abstract

As the application of deep learning has expanded to real-world problems with insufficient volume of training data, transfer learning recently has gained much attention as means of improving the performance in such small-data regime. However, when existing methods are applied between heterogeneous architectures and tasks, it becomes more important to manage their detailed configurations and often requires exhaustive tuning on them for the desired performance. To address the issue, we propose a novel transfer learning approach based on meta-learning that can automatically learn what knowledge to transfer from the source network to where in the target network. Given source and target networks, we propose an efficient training scheme to learn meta-networks that decide (a) which pairs of layers between the source and target networks should be matched for knowledge transfer and (b) which features and how much knowledge from each feature should be transferred. We validate our meta-transfer approach against recent transfer learning methods on various datasets and network architectures, on which our automated scheme significantly outperforms the prior baselines that find “what and where to transfer” in a hand-crafted manner.

1. Introduction

Learning deep neural networks (DNNs) requires large datasets, but it is expensive to collect a sufficient amount of labeled samples for each target task. A popular approach for handling such lack of data is transfer learning (Pan & Yang, 2010), whose goal is to transfer knowledge from a known source task to a new target task. The most widely used method for transfer learning is *pre-training with fine-*

^{*}Equal contribution ¹School of Electrical Engineering, KAIST, Korea ²OMNIOUS, Korea ³School of Computing, KAIST, Korea ⁴Graduate School of AI, KAIST, Korea ⁵AITRICS, Korea. Correspondence to: Jinwoo Shin <jinwoos@kaist.ac.kr>.

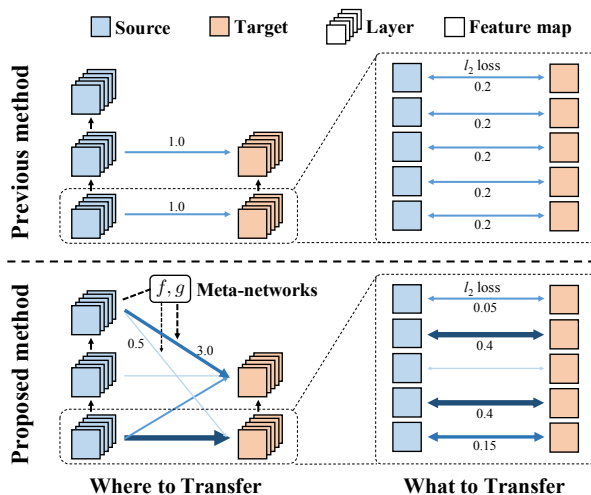


Figure 1. **Top:** Prior approaches. Knowledge transfer between two networks is done between *hand-crafted chosen pairs* of layers *without considering importance of channels*. **Bottom:** Our meta-transfer method. The meta-networks f, g automatically decide *amounts of knowledge transfer between layers* of the two networks and *importance of channels* when transfer. Line width indicates an amount of transfer in pairs of transferring layers and channels.

tuning (Razavian et al., 2014): first train a source DNN (e.g. ResNet (He et al., 2016)) with a large dataset (e.g. ImageNet (Deng et al., 2009)) and then, use the learned weights as an initialization to train a target DNN. Yet, fine-tuning definitely is not a panacea. If the source and target tasks are semantically distant, it may provide no benefit. Cui et al. (2018) suggest to sample from the source dataset depending on a target task for pre-training, but it is only possible when the source dataset is available. There is also no straightforward way to use fine-tuning, if the network architectures for the source and target tasks largely differ.

Several existing works can be applied to this challenging scenario of knowledge transfer between heterogeneous DNNs and tasks. Learning without forgetting (LwF) (Li & Hoiem, 2018) proposes to use knowledge distillation, suggested in Hinton et al. (2015), for transfer learning by introducing an additional output layer on a target model, and thus it can be applied to situations where the source and target tasks are different. FitNet (Romero et al., 2015) proposes a teacher-student training scheme for transferring the knowledge from a wider teacher network to a thinner student network, by

using teacher’s feature maps to guide the learning of the student. To guide the student network, FitNet uses ℓ_2 matching loss between the source and target features. Attention transfer (Zagoruyko & Komodakis, 2017) and Jacobian matching (Srinivas & Fleuret, 2018) suggest similar approaches to FitNet, but use attention maps generated from feature maps or Jacobians for transferring the source knowledge.

Our motivation is that these methods, while allowing to transfer knowledge between heterogeneous source and target tasks/architectures, have no mechanism to identify which source information to transfer, between which layers of the networks. Some source information is more important than others, while some are irrelevant or even harmful depending on the task difference. For example, since network layers generate representations at different level of abstractions (Zeiler & Fergus, 2014), the information of lower layers might be more useful when the input domains of the tasks are similar, but the actual tasks are different (e.g., fine-grained image classification tasks). Furthermore, under heterogeneous network architectures, it is not straightforward to associate a layer from the source network with one from the target network. Yet, since there was no mechanism to learn what to transfer to where, existing approaches require a careful manual configuration of layer associations between the source and target networks depending on tasks, which cannot be optimal.

Contribution. To tackle this problem, we propose a novel transfer learning method based on the concept of *meta-learning* (Naik & Mammone, 1992; Thrun & Pratt, 2012) that learns *what* information to transfer to *where*, from source networks to target networks with heterogeneous architectures and tasks. Our goal is learning to learn transfer rules for performing knowledge transfer in an automatic manner, considering the difference in the architectures and tasks between source and target, without hand-crafted tuning of transfer configurations. Specifically, we learn meta-networks that generate the weights for each feature and between each pair of source and target layers, jointly with the target network. Thus, it can automatically learn to identify which source network knowledge is useful, and where it should transfer to (see Figure 1). We validate our method, learning to transfer what and where (L2T-ww), to multiple source and target task combinations between heterogeneous DNN architectures, and obtain significant improvements over existing transfer learning methods. Our contributions are as follows:

- We introduce meta-networks for transfer learning that automatically decide which feature maps (channels) of a source model are useful and relevant for learning a target task and which source layers should be transferred to which target layers.
- To learn the parameters of meta-networks, we propose

an efficient meta-learning scheme. Our main novelty is to evaluate the one-step adaptation performance (meta-objective) of a target model learned by minimizing the transfer objective only (as an inner-objective). This scheme significantly accelerates the inner-loop procedure, compared to the standard scheme.

- The proposed method achieves significant improvements over baseline transfer learning methods in our experiments. For example, in the ImageNet experiment, our meta-transfer learning method achieves 65.05% accuracy on CUB200, while the second best baseline obtains 58.90%. In particular, our method outperforms baselines with a large margin when the target task has an insufficient number of training samples and when transferring from multiple source models.

Organization. The rest of the paper is organized as follows. In Section 2, we describe our method for selective knowledge transfer, and training scheme for learning the proposed meta-networks. Section 3 shows our experimental results under various settings, and Section 4 states the conclusion.

2. Learning What and Where to Transfer

Our goal is to learn to transfer useful knowledge from the source network to the target network, without requiring manual layer association or feature selection. To this end, we propose a meta-learning method that learns what knowledge of the source network to transfer to which layer in the target network. In this paper, we primarily focus on transfer learning between convolutional neural networks, but our method is generic and is applicable to other types of deep neural networks as well.

In Section 2.1, we describe meta-networks that learn what to transfer (for selectively transfer only the useful channels/features to a target model), and where to transfer (for deciding a layer-matching configuration that encourages learning a target task). Section 2.2 presents how to train the proposed meta-networks jointly with the target network.

2.1. Weighted Feature Matching

If a convolutional neural network is well-trained on a task, then its intermediate feature spaces should have useful knowledge for the task. Thus, mimicking the well-trained features might be helpful for training another network. To formalize the loss forcing this effect, let x be an input, and y be the corresponding (ground-truth) output. For image classification tasks, $\{x\}$ and $\{y\}$ are images and their class labels. Let $S^m(x)$ be intermediate feature maps of the m^{th} layer of the pre-trained source network S . Our goal is then to train another target network T_θ with parameter θ utilizing the knowledge of S . Let $T_\theta^n(x)$ be intermediate

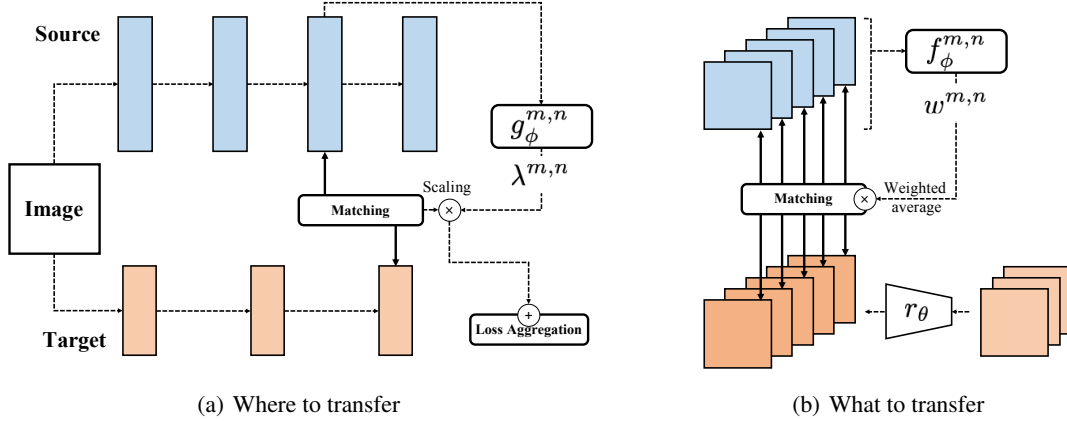


Figure 2. Our meta-transfer learning method for selective knowledge transfer. The meta-transfer networks are parameterized by ϕ and are learned via meta-learning. The dashed lines indicate flows of tensors such as feature maps, and solid lines denote ℓ_2 feature matching. (a) $g_\phi^{m,n}$ outputs weights of matching pairs $\lambda^{m,n}$ between the m^{th} and n^{th} layers of the source and target models, respectively, and (b) $f_\phi^{m,n}$ outputs weights for each channel.

feature maps of the n^{th} layer of the target network. Then, we minimize the following ℓ_2 objective, similar to that used in FitNet (Romero et al., 2015), to transfer the knowledge from $S^m(x)$ to $T_\theta^n(x)$:

$$\|r_\theta(T_\theta^n(x)) - S^m(x)\|_2^2$$

where r_θ is a linear transformation parameterized by θ such as a pointwise convolution. We refer to this method as *feature matching*. Here, the parameter θ consists of both the parameter for linear-transformation r_θ and non-linear neural network T_θ , where the former is only necessary in training the latter and is not required at testing time.

What to transfer. In general transfer learning settings, the target model is trained for a task that is different from that of the source model. In this case, not all the intermediate features of the source model may be useful to learn the target task. Thus, to give more attention on the useful channels, we consider a weighted feature matching loss that can emphasize the channels according to their utility on the target task:

$$\begin{aligned} \mathcal{L}_{\text{wfm}}^{m,n}(\theta|x, w^{m,n}) \\ = \frac{1}{HW} \sum_c w_c^{m,n} \sum_{i,j} (r_\theta(T_\theta^n(x))_{c,i,j} - S^m(x)_{c,i,j})^2, \end{aligned}$$

where $H \times W$ is the spatial size of $S^m(x)$ and $r_\theta(T_\theta^n(x))$, the inner-summation is over $i \in \{1, 2, \dots, H\}$ and $j \in \{1, 2, \dots, W\}$, and $w_c^{m,n}$ is the non-negative weight of channel c with $\sum_c w_c^{m,n} = 1$. Since the important channels to transfer can vary for each input image, we set channel weights as a function, $w^{m,n} = [w_c^{m,n}] = f_\phi^{m,n}(S^m(x))$, by taking the softmax output of a small meta-network which takes features of source models as an input. We let ϕ denote the parameters of meta-networks throughout this paper.

Where to transfer. When transferring knowledge from a source model to a target model, deciding pairs (m, n) of layers in the source and target model is crucial to its effectiveness. Previous approaches (Romero et al., 2015; Zagoruyko & Komodakis, 2017) select the pairs manually based on prior knowledge of architectures or semantic similarities between tasks. For example, attention transfer (Zagoruyko & Komodakis, 2017) matches the last feature maps of each group of residual blocks in ResNet (He et al., 2016). However, finding the optimal layer association is not a trivial problem and requires exhaustive tuning based on trial-and-error, given models with different numbers of layers or heterogeneous architectures, e.g., between ResNet (He et al., 2016) and VGG (Simonyan & Zisserman, 2015). Hence, we introduce a learnable parameter $\lambda^{m,n} \geq 0$ for each pair (m, n) which can decide the amount of transfer between the m^{th} and n^{th} layers of source and target models, respectively. We also set $\lambda^{m,n} = g_\phi^{m,n}(S^m(x))$ for each pair (m, n) as an output of a meta-network $g_\phi^{m,n}$ that automatically decides important pairs of layers for learning the target task.

The combined transfer loss given the weights of channels w and weights of matching pairs λ is

$$\mathcal{L}_{\text{wfm}}(\theta|x, \phi) = \sum_{(m,n) \in \mathcal{C}} \lambda^{m,n} \mathcal{L}_{\text{wfm}}^{m,n}(\theta|x, w^{m,n}),$$

where \mathcal{C} be a set of candidate pairs. Our final loss $\mathcal{L}_{\text{total}}$ to train a target model then is given as:

$$\mathcal{L}_{\text{total}}(\theta|x, y, \phi) = \mathcal{L}_{\text{org}}(\theta|x, y) + \beta \mathcal{L}_{\text{wfm}}(\theta|x, \phi).$$

where \mathcal{L}_{org} is the original loss (e.g., cross entropy) and $\beta > 0$ is a hyper-parameter. We note that $w^{m,n}$ and $\lambda^{m,n}$ decide what and where to transfer, respectively. We provide an illustration of our transfer learning scheme in Figure 2.

2.2. Training Meta-Networks and Target Model

Our goal is to achieve high performance on the target task when the target model is learned using the training objective $\mathcal{L}_{\text{total}}(\cdot|x, y, \phi)$. To maximize the performance, the feature matching term $\mathcal{L}_{\text{wfm}}(\cdot|x, \phi)$ should encourage learning of useful features for the target task, e.g., predicting labels. To measure and increase usefulness of the feature matching decided by meta-networks parameterized by ϕ , a standard approach is to use the following bilevel scheme (Colson et al., 2007) to train ϕ , e.g., see (Finn et al., 2017; Franceschi et al., 2018):

1. Update θ to minimize $\mathcal{L}_{\text{total}}(\theta|x, y, \phi)$ for T times.
2. Measure $\mathcal{L}_{\text{org}}(\theta|x, y)$ and update ϕ to minimize it.

In the above, the actual objective $\mathcal{L}_{\text{total}}$ for learning the target model is used in the inner-loop, and the original loss \mathcal{L}_{org} is used as a meta-objective to measure the effectiveness of $\mathcal{L}_{\text{total}}$ for learning the target model to perform well.

However, since our meta-networks affect the learning procedure of the target model weakly through the regularization term \mathcal{L}_{wfm} , their influence on \mathcal{L}_{org} can be very marginal, unless one uses a very large number of inner-loop iterations T . Consequently, it causes difficulties on updating ϕ using gradient $\nabla_{\phi}\mathcal{L}_{\text{org}}$. To tackle this challenge, we propose the following alternative scheme:

1. Update θ to minimize $\mathcal{L}_{\text{wfm}}(\theta|x, \phi)$ for T times.
2. Update θ to minimize $\mathcal{L}_{\text{org}}(\theta|x, y)$ once.
3. Measure $\mathcal{L}_{\text{org}}(\theta|x, y)$ and update ϕ to minimize it.

In the first stage, given the current parameter $\theta_0 = \theta$, we update the target model for T times via gradient-based algorithms for minimizing \mathcal{L}_{wfm} . Namely, the resulting parameter θ_T is learned only using the knowledge of the source model. Since transfer is done by the form of feature matching, it is feasible to train useful features for the target task by selectively mimic the source features. More importantly, it increases the influence of the regularization term \mathcal{L}_{wfm} on the learning procedure of the target model in the inner-loop, since the target features are solely trained by the source knowledge (without target labels). The second stage is an one-step adaptation θ_{T+1} from θ_T toward the target label. Then, in the third stage, the task-specific objective $\mathcal{L}_{\text{org}}(\theta_{T+1})$ can measure how quickly the target model has adapted (via only one step from θ_T) to the target task, under the sample used in the first and second stage. Finally, the meta-parameter ϕ can be trained by minimizing $\mathcal{L}_{\text{org}}(\theta_{T+1})$. The above 3-stage scheme encourages significantly faster training of ϕ , compared the standard 2-stage one. This is because the former measures the effect of the regularization

Algorithm 1 Learning of θ with meta-parameters ϕ

Input: Dataset $\mathcal{D}_{\text{train}} = \{(x_i, y_i)\}$, learning rate α
repeat
 Sample a batch $\mathcal{B} \subset \mathcal{D}_{\text{train}}$ with $|\mathcal{B}| = B$
 Update θ to minimize $\frac{1}{B} \sum_{(x,y) \in \mathcal{B}} \mathcal{L}_{\text{total}}(\theta|x, y, \phi)$
 Initialize $\theta_0 \leftarrow \theta$
 for $t = 0$ **to** $T - 1$ **do**
 $\theta_{t+1} \leftarrow \theta_t - \alpha \nabla_{\theta} \frac{1}{B} \sum_{(x,y) \in \mathcal{B}} \mathcal{L}_{\text{wfm}}(\theta_t|x, \phi)$
 end for
 $\theta_{T+1} \leftarrow \theta_T - \alpha \nabla_{\theta} \frac{1}{B} \sum_{(x,y) \in \mathcal{B}} \mathcal{L}_{\text{org}}(\theta_T|x, y)$
 Update ϕ using $\nabla_{\phi} \frac{1}{B} \sum_{(x,y) \in \mathcal{B}} \mathcal{L}_{\text{org}}(\theta_{T+1}|x, y)$
until done

term \mathcal{L}_{wfm} more directly to the original \mathcal{L}_{org} , and allows to choose a small T to update ϕ meaningfully (we choose $T = 2$ in our experiments).

In the case of using the vanilla gradient descent algorithm for updates, the 3-stage training scheme to learn meta-parameters ϕ can be formally written as the following optimization task:

$$\begin{aligned} & \underset{\phi}{\text{minimize}} && \mathcal{L}_{\text{org}}(\theta_{T+1}|x, y) \\ & \text{subject to} && \theta_{T+1} = \theta_T - \alpha \nabla_{\theta} \mathcal{L}_{\text{org}}(\theta_T|x, y), \\ & && \theta_{t+1} = \theta_t - \alpha \nabla_{\theta} \mathcal{L}_{\text{wfm}}(\theta_t|x, \phi), \\ & && t = 0, \dots, T - 1, \end{aligned}$$

where $\alpha > 0$ is a learning rate. To solve the above optimization problem, we use Reverse-HG (Franceschi et al., 2017) that can compute $\nabla_{\phi}\mathcal{L}_{\text{org}}(\theta_{T+1}|x, y)$ efficiently using Hessian-vector products.

To train the target model jointly with meta-networks, we alternatively update the target model parameters θ and the meta-network parameters ϕ . We first update the target model for a single step with objective $\mathcal{L}_{\text{total}}(\theta|x, y, \phi)$. Then, given current target model parameters, we update the meta-networks parameters ϕ using the 3-stage bilevel training scheme described above. This eliminates an additional meta-training phase for learning ϕ . The proposed training scheme is formally outlined in Algorithm 1.

3. Experiments

We validate our meta-transfer learning method that learns what and where to transfer, between heterogeneous network architectures and tasks.

3.1. Setups

Network architectures and tasks for source and target. To evaluate various transfer learning methods including ours, we perform experiments on two scales of image classification tasks, 32×32 and 224×224 . For 32×32 scale,

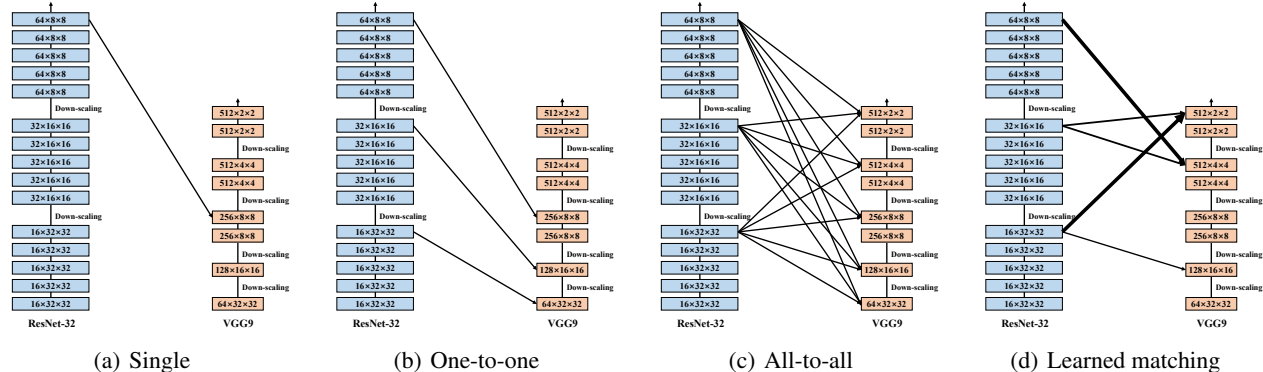


Figure 3. (a)-(c) Matching configurations \mathcal{C} between ResNet32 (left) and VGG9 (right). (d) The amount $\lambda^{m,n}$ of transfer between layers after learning. Line widths indicates the transfer amount. We omit the lines when $\lambda^{m,n}$ is less than 0.1.

we use the TinyImageNet¹ dataset as a source task, and CIFAR-10, CIFAR-100 (Krizhevsky & Hinton, 2009), and STL-10 (Coates et al., 2011) datasets as target tasks. We train 32-layer ResNet (He et al., 2016) and 9-layer VGG (Simonyan & Zisserman, 2015) on the source and target tasks, respectively. For 224×224 scale, the ImageNet (Deng et al., 2009) dataset is used as a source dataset, and Caltech-UCSD Bird 200 (Wah et al., 2011), MIT Indoor Scene Recognition (Quattoni & Torralba, 2009), Stanford 40 Actions (Yao et al., 2011) and Stanford Dogs (Khosla et al., 2011) datasets as target tasks. For these datasets, we use 34-layer and 18-layer ResNet as a source and target model, respectively, unless otherwise stated.

Meta-network architecture. For all experiments, we construct the meta-networks as 1-layer fully-connected networks for each pair $(m, n) \in \mathcal{C}$ where \mathcal{C} is the set of candidates of pairs, or *matching configuration* (see Figure 3). It takes the globally average pooled features of the m^{th} layer of the source network as an input, and outputs $w_c^{m,n}$ and $\lambda^{m,n}$. As for the channel assignments w , we use the softmax activation to generate them while satisfying $\sum_c w_c^{m,n} = 1$, and for transfer amount λ between layers, we commonly use ReLU6 (Krizhevsky & Hinton, 2010), $\max(0, \min(6, x))$ to ensure non-negativeness of λ and to prevent $\lambda^{m,n}$ from becoming too large.

Compared schemes for transfer learning. We compare our methods with the following prior methods and their combinations: learning without forgetting (LwF) (Li & Hoiem, 2018), attention transfer (AT) (Zagoruyko & Komodakis, 2017) and unweighted feature matching (FM) (Romero et al., 2015).² Here, AT and FM transfer knowledge on

feature-level as like ours by matching attention maps or feature maps between source and target layers, respectively. The feature-level transfer methods generally choose layers just before down-scaling, e.g., the last layer of each residual group for ResNet, and match pairs of the layers of same spatial size. Following this convention, we evaluate two hand-crafted configurations (single, one-to-one) for prior methods and a new configurations (all-to-all) for our methods: (a) *single*: use a pair of the last feature in the source model and a layer with the same spatial size in the target model, (b) *one-to-one*: connect each layer just before down-scaling in the source model to a target layer of the same spatial size, (c) *all-to-all*: use all pairs of layers just before down-scaling, e.g., between ResNet and VGG architectures, we consider $3 \times 5 = 15$ pairs. For matching features of different spatial sizes, we simply use a bilinear interpolation. These configurations are illustrated in Figure 3. Among various combinations between prior methods and matching configurations, we only report the results of those achieving the meaningful performance gains.

3.2. Evaluation on Various Target Tasks

We first evaluate the effect of learning to transfer what (L2T-w) without learning to transfer where. To this end, we use conventional hand-crafted matching configurations, single and one-to-one, illustrated in Figure 3(a) and 3(b), respectively. For most cases reported in Table 1, L2T-w improves the performance on target tasks compared to the unweighted counterpart (FM): for fine-grained target tasks transferred from ImageNet, the gain of L2T-w over FM is more significant. The results support that our method, learning what to transfer, is more effective when target tasks have specific types of input distributions, e.g., fine-grained classification, while the source model is trained on a general task.

matching (JM) (Srinivas & Fleuret, 2018) as the improvement of LwF+AT+JM over LwF+AT is marginal in our setups.

¹<https://tiny-imagenet.herokuapp.com/>

² In our experimental setup, we reproduce similar relative improvements from the scratch for these baselines as reported in the original papers. We do not report the results of Jacobian




Figure 4. Change of $\lambda^{m,n}$ during training for STL-10 as the target task with TinyImageNet as the source task. We plot mean and standard deviation of $\lambda^{m,n}$ of all samples for every 10 epochs.

Table 1. Classification accuracy (%) of transfer learning from TinyImageNet (32×32) or ImageNet (224×224) to CIFAR-100, STL-10, Caltech-UCSD Bird 200 (CUB200), MIT Indoor Scene Recognition (MIT67), Stanford 40 Actions (Stanford40) and Stanford Dogs datasets. For TinyImageNet, ResNet32 and VGG9 are used as a source and target model, respectively, and ResNet34 and ResNet18 are used for ImageNet.

Source task	TinyImageNet		ImageNet				
	Target task	CIFAR-100	STL-10	CUB200	MIT67	Stanford40	Stanford Dogs
Scratch		67.69 \pm 0.22	65.18 \pm 0.91	42.15 \pm 0.75	48.91 \pm 0.53	36.93 \pm 0.68	58.08 \pm 0.26
LwF		69.23 \pm 0.09	68.64 \pm 0.58	45.52 \pm 0.66	53.73 \pm 2.14	39.73 \pm 1.63	66.33 \pm 0.45
AT (one-to-one)		67.54 \pm 0.40	74.19 \pm 0.22	57.74 \pm 1.17	59.18 \pm 1.57	59.29 \pm 0.91	69.70 \pm 0.08
LwF+AT (one-to-one)		68.75 \pm 0.09	75.06 \pm 0.57	58.90 \pm 1.32	61.42 \pm 1.68	60.20 \pm 1.34	72.67 \pm 0.26
FM (single)		69.40 \pm 0.67	75.00 \pm 0.34	47.60 \pm 0.31	55.15 \pm 0.93	42.93 \pm 1.48	66.05 \pm 0.76
FM (one-to-one)		69.97 \pm 0.24	76.38 \pm 1.18	48.93 \pm 0.40	54.88 \pm 1.24	44.50 \pm 0.96	67.25 \pm 0.88
L2T-w (single)		70.27 \pm 0.09	74.35 \pm 0.92	51.95 \pm 0.83	60.41 \pm 0.37	46.25 \pm 3.66	69.16 \pm 0.70
L2T-w (one-to-one)		70.02 \pm 0.19	76.42 \pm 0.52	56.61 \pm 0.20	59.78 \pm 1.90	48.19 \pm 1.42	69.84 \pm 1.45
L2T-ww (all-to-all)		70.96\pm0.61	78.31\pm0.21	65.05\pm1.19	64.85\pm2.75	63.08\pm0.88	78.08\pm0.96

Next, instead of using hand-crafted matching pairs of layers, we also learn where to transfer starting from all matching pairs illustrated in Figure 3(c). The proposed final scheme in our paper, learning to transfer what and where (L2T-ww), often improves the performance significantly compared to the hand-crafted matching (L2T-w). As a result, L2T-ww achieves the best accuracy for all cases (with large margin) reported in Table 1, e.g., on the CUB200 dataset, we attain 10.4% relative improvement compared to the second best baseline.

Figure 3(d) shows the amounts $\lambda^{m,n}$ of transfer between pairs of layers after learning transfer from TinyImageNet to STL-10. As shown in the figure, our method transfers knowledge to higher layers in the target model: $\lambda^{2,5} = 1.40$, $\lambda^{1,5} = 2.62$, $\lambda^{3,4} = 2.88$, $\lambda^{2,4} = 0.74$. The amounts $\lambda^{m,n}$ of other pairs are smaller than 0.1, except $\lambda^{1,2} = 0.21$. Clearly, those matching pairs are not trivial to find by hand-crafted tuning, which justifies that our method for learning where to transfer is useful. Furthermore, since

our method outputs sample-wise $\lambda^{m,n}$, amounts of transfer are adjusted more effectively compared to fixed matching pairs over all the samples. For example, amounts of transfer from source features $S^1(x)$ have relatively smaller variance over the samples (Figure 4(a)) compared to the those of $S^3(x)$ (Figure 4(c)). This is because higher-level features are more task-specific while lower-level features are more task-agnostic. It evidences that meta-networks g_ϕ adjust the amounts of transfer for each sample considering the relationship between tasks and the levels of abstractions of features.

3.3. Experiments on Limited-Data Regimes

When a target task has a small number of labeled samples for training, transfer learning can be even more effective. To evaluate our method (L2T-ww) on such limited-data scenario, we use CIFAR-10 as a target task dataset by reducing the number of samples. We use $N \in \{50, 100, 250, 500, 1000\}$ training samples for each class,

Table 2. Classification accuracy (%) of VGG9 on STL-10 transferred from multiple source models. The first source model is ResNet32 trained on TinyImageNet. The additional source model is one of three: ResNet20 trained on TinyImageNet, another ResNet32 trained on TinyImageNet, and ResNet32 trained on CIFAR-10. We report the performance of the target model transferred from a single source model and two source models.

First source	TinyImageNet (ResNet32)				
	Second source	None	TinyImageNet (ResNet20)	TinyImageNet (ResNet32)	CIFAR-10 (ResNet32)
Scratch		65.18 \pm 0.91	65.18 \pm 0.91	65.18 \pm 0.91	65.18 \pm 0.91
LwF		68.64 \pm 0.58	68.56 \pm 2.24	68.05 \pm 2.12	69.51 \pm 0.63
AT		74.19 \pm 0.22	73.24 \pm 0.12	73.78 \pm 1.16	73.99 \pm 0.51
LwF+AT		75.06 \pm 0.57	74.72 \pm 0.46	74.77 \pm 0.30	74.41 \pm 1.51
FM (single)		75.00 \pm 0.34	75.83 \pm 0.56	75.99 \pm 0.11	74.60 \pm 0.73
FM (one-to-one)		76.38 \pm 1.18	77.45 \pm 0.48	77.69 \pm 0.79	77.15 \pm 0.41
L2T-ww (all-to-all)		78.31\pm0.21	79.35\pm0.41	79.80\pm0.52	80.52\pm0.29

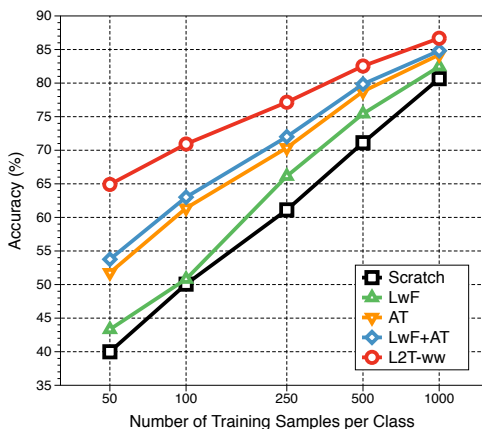


Figure 5. Transfer from TinyImageNet to CIFAR-10 with varying numbers of training samples per class in CIFAR-10. x -axis is plotted in logarithmic scale.

and compare the performance of learning from scratch, LwF, AT, LwF+AT and L2T-ww. The results are reported in Figure 5. They show that ours achieves significant more improvements compared to other baselines, when the volume of the target dataset is smaller. For example, in the case of $N = 50$, our method achieves 64.91% classification accuracy, while the baselines, LwF+AT, AT, LwF and scratch show 53.76%, 51.76%, 43.32% and 39.99%, respectively. Observe that ours needs only 50 samples per class to achieve similar accuracy of LwF with 250 samples per class.

3.4. Experiments on Multi-Source Transfer

In practice, one can have multiple pre-trained source models with various source datasets. Transfer from multiple sources may potentially provide more knowledge for learning a target task, however, using them simultaneously could require more hand-crafted configurations of transfer, such

as balancing the transfer from many sources or choosing different pairs of layers depending on the source models. To evaluate the effects of using multiple source models, we consider the scenarios transferred from two source models simultaneously, where the models are different architectures (ResNet20, ResNet32) or trained on different datasets (TinyImageNet, CIFAR-10). In Table 2, we report the results of ours (L2T-ww) and other transfer methods on a target task STL-10 with 9-layer VGG as a target model architecture.

Our method consistently improves the target model performance over more informative transitions (from left to right in Table 2) on sources, i.e., when using a larger source model (ResNet20 \rightarrow ResNet32) or using a different second source dataset (TinyImageNet \rightarrow CIFAR-10). This is not the case for all other methods. In particular, compare the best performance of each method transferred from two TinyImageNet models and TinyImageNet+CIFAR-10 models as sources. Then, one can conclude that ours is the only one that effectively aggregates the heterogeneous source knowledge, i.e., TinyImageNet+CIFAR-10. It shows the importance of choosing the right configurations of transfer when using multiple source models, and confirms that ours can automatically decide the useful configuration from many possible candidate pairs for transfer.

3.5. Visualization

With learning what to transfer, our weighted feature matching will allocate larger attention to task-related channels of feature maps. To visualize the attention used in knowledge transfer, we compare saliency maps (Simonyan et al., 2014) for unweighted (FM) and weighted (L2T-w) matching between the last layers of source and target models. Saliency maps can be computed as follows:

$$M_{i,j} = \max_c \left| \frac{\partial \mathcal{L}_{\text{wfm}}^{m,n}(\theta|x, w^{m,n})}{\partial x_{c,i,j}} \right|$$



Figure 6. More (the second column) and less (the third column) activated pixels in the saliency maps of L2T-w compared to unweighted feature matching (FM) on images of (a) CUB200 and (b) Stanford Dogs datasets. When computing saliency maps, we use normalized gradients. One can observe that the higher activated pixels induced by L2T-w tend to correspond to where task-specific objects are, while less activated location spread over entire location.

where x is an image, c is a channel of the image, e.g., RGB, and $(i, j) \in \{1, 2, \dots, H\} \times \{1, 2, \dots, W\}$ is a pixel position. For the unweighted case, we use uniform weights. On the other hand, for the weighted case, we use the outputs $w^{m,n} = f_{\phi}^{m,n}(S^m(x))$ of meta-networks learned by our meta-training scheme. Figure 6 shows which pixels are more or less activated in the saliency map of L2T-w compared to FM. As shown in the figure, pixels containing task-specific objects (birds or dogs) are more activated when using L2T-w, while background pixels are less activated. It means that the weights $w^{m,n}$ make knowledge of the source model be more task-specific, consequently it can improve transfer learning.

4. Conclusion

We propose a transfer method based on meta-learning which can transfer knowledge selectively depending on tasks and architectures. Our method transfers more important knowledge for learning a target task, with identifying *what* and *where* to transfer using meta-networks. To learn the meta-networks, we design an efficient meta-learning scheme which requires a few steps in the inner-loop procedure. By doing so, we jointly train the target model and the meta-networks. We believe that our work would shed a new angle for complex transfer learning tasks between heterogeneous or/and multiple network architectures and tasks.

Acknowledgements

This work was supported by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government MSIT (No.2016-0-00563, Research on Adaptive Machine Learning Technology Development for Intelligent Autonomous Digital Companion) and supported by the Engineering Research Center Program through the National Research Foundation of Korea (NRF) funded by the Korean Government MSIT (NRF-2018R1A5A1059921).

References

- Coates, A., Ng, A., and Lee, H. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS 2011)*, 2011.
- Colson, B., Marcotte, P., and Savard, G. An overview of bilevel optimization. *Annals of operations research*, 2007.
- Cui, Y., Song, Y., Sun, C., Howard, A., and Belongie, S. Large scale fine-grained categorization and domain-specific transfer learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2018)*, 2018.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2009)*, 2009.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, 2017.
- Franceschi, L., Donini, M., Frasconi, P., and Pontil, M. Forward and reverse gradient-based hyperparameter optimization. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- Franceschi, L., Frasconi, P., Salzo, S., Grazzi, R., and Pontil, M. Bilevel programming for hyperparameter optimization and meta-learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2016)*, 2016.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. In *Deep Learning and Representation Learning Workshop, Advances in Neural Information Processing Systems 29 (NIPS 2015)*, 2015.
- Khosla, A., Jayadevaprakash, N., Yao, B., and Fei-Fei, L. Novel dataset for fine-grained image categorization. In *The 1st Workshop on Fine-Grained Visual Categorization, the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011)*, 2011.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Krizhevsky, A. and Hinton, G. Convolutional deep belief networks on cifar-10, 2010.
- Li, Z. and Hoiem, D. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- Naik, D. K. and Mammone, R. Meta-neural networks that learn by learning. In *Neural Networks, 1992. IJCNN., International Joint Conference on*, 1992.
- Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 2010.
- Quattoni, A. and Torralba, A. Recognizing indoor scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2009)*, 2009.
- Razavian, A. S., Azizpour, H., Sullivan, J., and Carlsson, S. Cnn features off-the-shelf: an astounding baseline for recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR 2014)*, 2014.
- Romero, A., Ballas, N., Kahou, S. E., Chassang, A., Gatta, C., and Bengio, Y. Fitnets: Hints for thin deep nets. In *The 3rd International Conference on Learning Representations (ICLR 2015)*, 2015.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *The 3rd International Conference on Learning Representations (ICLR 2015)*, 2015.
- Simonyan, K., Vedaldi, A., and Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *The 2nd International Conference on Learning Representations Workshop (ICLR 2014)*, 2014.
- Srinivas, S. and Fleuret, F. Knowledge transfer with Jacobian matching. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, 2018.
- Thrun, S. and Pratt, L. *Learning to learn*. Springer Science & Business Media, 2012.
- Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.

Yao, B., Jiang, X., Khosla, A., Lin, A. L., Guibas, L., and Fei-Fei, L. Human action recognition by learning bases of action attributes and parts. In *The IEEE International Conference on Computer Vision (ICCV 2011)*, 2011.

Zagoruyko, S. and Komodakis, N. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *The 5th International Conference on Learning Representations (ICLR 2017)*, 2017.

Zeiler, M. D. and Fergus, R. Visualizing and understanding convolutional networks. In *The European Conference on Computer Vision (ECCV 2014)*, 2014.