
Gradient-Based Meta-Learning with Learned Layerwise Metric and Subspace

Yoonho Lee¹ Seungjin Choi¹

Abstract

Gradient-based meta-learning methods leverage gradient descent to learn the commonalities among various tasks. While previous such methods have been successful in meta-learning tasks, they resort to simple gradient descent during meta-testing. Our primary contribution is the *MT-net*, which enables the meta-learner to learn on each layer’s activation space a subspace that the task-specific learner performs gradient descent on. Additionally, a task-specific learner of an *MT-net* performs gradient descent with respect to a meta-learned distance metric, which warps the activation space to be more sensitive to task identity. We demonstrate that the dimension of this learned subspace reflects the complexity of the task-specific learner’s adaptation task, and also that our model is less sensitive to the choice of initial learning rates than previous gradient-based meta-learning methods. Our method achieves state-of-the-art or comparable performance on few-shot classification and regression tasks.

1. Introduction

While recent deep learning methods achieve superhuman performance on various tasks including image classification (Krizhevsky et al., 2012) or playing games (Mnih et al., 2015), they can only do so using copious amounts of data and computational resources. In many problems of interest, learners may not have such luxuries. *Meta-learning* (Schmidhuber, 1987; Schmidhuber et al., 1997; Thrun & Pratt, 1998) methods are a potential solution to this problem; these methods leverage information gathered from prior learning experience to learn more effectively in novel tasks. This line of research typically casts learning as a two-level process, each with a different scope. The *meta-learner*

operates on the level of tasks, gathering information from several instances of task-specific learners. A *task-specific learner*, on the other hand, operates on the level of data-points, and incorporates the meta-learner’s knowledge in its learning process.

Model-agnostic meta-learning (MAML) (Finn et al., 2017) is a meta-learning method that directly optimizes the gradient descent procedure of task-specific learners. All task-specific learners of MAML share initial parameters, and a meta-learner optimizes these initial parameters such that gradient descent starting from such initial parameters quickly yields good performance. An implicit assumption in having the meta-learner operate in the same space as task-specific learners is that the two different scopes of learning require equal degrees of freedom.

Our primary contribution is the MT-net (Figure 1), a neural network architecture and task-specific learning procedure. An MT-net differs from previous gradient-based meta-learning methods in that the meta-learner determines a subspace and a corresponding metric that task-specific learners can learn in, thus setting the degrees of freedom of task-specific learners to an appropriate amount. Note that the activation space of the cell shown in Fig.1(b) is 3-dimensional. Because the task-specific learners can only change weights that affect two of the three intermediate activations, task-specific learning only happens on a subspace with 2 degrees of freedom. Additionally, meta-learned parameters \mathbf{T} alter the geometry of the activation space (Fig.1(c)) of task-specific parameters so that task-specific learners are more sensitive to change in task.

2. Background

2.1. Problem Setup

We briefly explain the meta-learning problem setup which we apply to few-shot tasks.

The problems of k -shot regression and classification are as follows. In the training phase for a meta-learner, we are given a (possibly infinite) set of tasks $\{\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3, \dots\}$. Each task provides a training set and a test set $\{\mathcal{D}_{\mathcal{T}_i, \text{train}}, \mathcal{D}_{\mathcal{T}_i, \text{test}}\}$. We assume here that the training set $\mathcal{D}_{\mathcal{T}_i, \text{train}}$ has k examples per class, hence the name k -shot learning. A particular task $\mathcal{T} \in \{\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3, \dots\}$

¹Department of Computer Science and Engineering, Pohang University of Science and Technology, Korea. Correspondence to: Yoonho Lee <einet89@postech.ac.kr>, Seungjin Choi <seungjin@postech.ac.kr>.

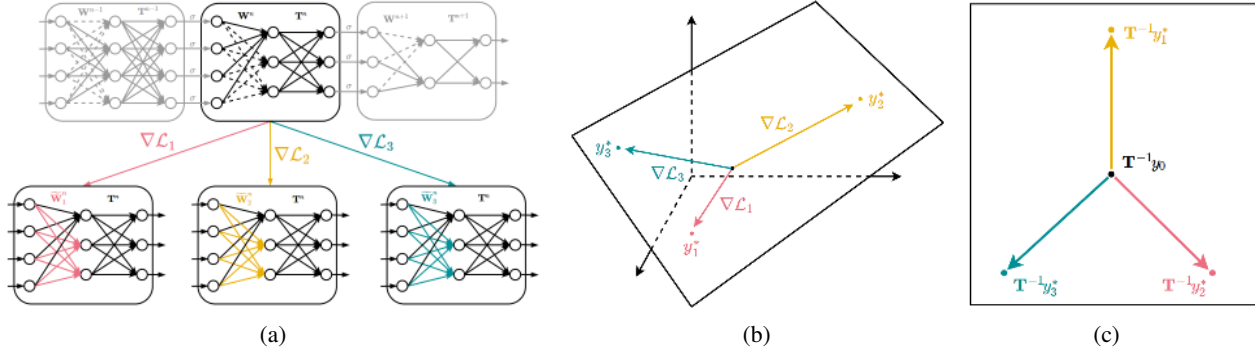


Figure 1. Task-specific learning in an MT-net. (a) A cell (rounded rectangle) consists of two layers. In addition to initial weights (black), the meta-learner specifies weights to be changed (dotted lines) by task-specific learners (colored). (b) Activation of this cell has 3 dimensions, but activation of task-specific learners only change within a subspace (white plane). (c) The value of \mathbf{T} affects task-specific learning so that gradients of \mathbf{W} are sensitive to task identity. Best seen in color.

is assumed to be drawn from the distribution of tasks $p(\mathcal{T})$. Given a task $\mathcal{T} \sim p(\mathcal{T})$, the task-specific model $f_{\theta_{\mathcal{T}}}$ (our work considers a feedforward neural network) is trained using the dataset $\mathcal{D}_{\mathcal{T},train}$ and its corresponding loss $\mathcal{L}_{\mathcal{T}}(\theta_{\mathcal{T}}, \mathcal{D}_{\mathcal{T},train})$. Denote by $\tilde{\theta}_{\mathcal{T}}$ parameters obtained by optimizing $\mathcal{L}_{\mathcal{T}}(\theta_{\mathcal{T}}, \mathcal{D}_{\mathcal{T},train})$. Then, the meta-learner f_{θ} is updated using the feedback from the collection of losses $\{\mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathcal{T}}, \mathcal{D}_{\mathcal{T},test})\}_{\mathcal{T} \sim p(\mathcal{T})}$, where the loss of each task is evaluated using the test data $\mathcal{D}_{\mathcal{T},test}$. Given a new task \mathcal{T}_{new} (not considered during meta-training), the meta-learner helps the model $f_{\theta_{\mathcal{T}_{new}}}$ to quickly adapt to the new task \mathcal{T}_{new} , by warm-starting the gradient updates.

2.2. Model-Agnostic Meta-Learning

We briefly review model-agnostic meta-learning (MAML) (Finn et al., 2017), emphasizing commonalities and differences between MAML and our method. MAML is a meta-learning method that can be used on any model that learns using gradient descent. This method is loosely inspired by fine-tuning, and it learns initial parameters of a network such that the network’s loss after a few (usually $1 \sim 5$) gradient steps is minimized.

Consider a model with parameters θ . MAML alternates between the two updates (1) and (2) to determine initial parameters θ for task-specific learners to warm-start the gradient descent updates, such that new tasks can be solved using a small number of examples. Each task-specific learner updates its parameters by gradient descent (1) using the loss evaluated with the training data $\{\mathcal{D}_{\mathcal{T},train}\}$. The meta-optimization across tasks (2) is performed such that the parameters θ are updated using the loss evaluated with $\{\mathcal{D}_{\mathcal{T},test}\}$. Note that during meta-optimization (2), the gradient is computed with respect to initial parameters θ but the test loss is computed with respect to task-specific parameters $\tilde{\theta}_{\mathcal{T}}$.

$$\tilde{\theta}_{\mathcal{T}} \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}}(\theta, \mathcal{D}_{\mathcal{T},train}) \quad (1)$$

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \left(\sum_{\mathcal{T} \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathcal{T}}, \mathcal{D}_{\mathcal{T},test}) \right), \quad (2)$$

where $\alpha > 0$ and $\beta > 0$ are learning rates and the summation in (2) is computed using minibatches of tasks sampled from $p(\mathcal{T})$.

Intuitively, a well-learned initial parameter θ is close to some local optimum for every task $\mathcal{T} \sim p(\mathcal{T})$. Furthermore, the update (1) is sensitive to task identity in the sense that $\tilde{\theta}_{\mathcal{T}_1}$ and $\tilde{\theta}_{\mathcal{T}_2}$ have different behaviors for different tasks $\mathcal{T}_1, \mathcal{T}_2 \sim p(\mathcal{T})$.

Recent work has shown that gradient-based optimization is a universal learning algorithm (Finn & Levine, 2017), meaning that any learning algorithm can be approximated up to arbitrary accuracy using some parameterized model and gradient descent. Thus, no expressiveness is lost by only considering gradient-based learners as in (1). Note that since MAML operates using a single fixed model, one may have to go through trial and error to find such a good model.

Our method is similar to MAML in that our method also differentiates through gradient update steps to optimize performance after fine-tuning. However, while MAML assumes a fixed model, our method actually chooses a subset of its weights to fine-tune. In other words, it (meta-)learns which model is most suitable for the task at hand. Furthermore, whereas MAML learns with standard gradient descent, a subset of our method’s parameters effectively ‘warp’ the parameter space of the parameters to be learned during meta-testing to enable faster learning.




Figure 2. A diagram of the adaptation process of a Transformation Network (T-net). Blue values are meta-learned and shared across all tasks. Orange values are different for each task.

3. Meta-Learning Models

We present our two models in this section: Transformation Networks (T-net) and Mask Transformation Networks (MT-net), both of which are trained with gradient-based meta-learning. A T-net learns a metric in its activation space; this metric informs each task-specific learner’s update direction and step size. An MT-net additionally learns which subset of its weights to update for task-specific learning. Therefore, an MT-net learns to automatically assign one of two roles (task-specific or task-mutual) to each of its weights.

3.1. T-net

We consider a model $f_\theta(\cdot)$ with parameters θ . This model consists of L cells, where each cell is parameterized* as $\mathbf{T}\mathbf{W}$:

$$f_\theta(\mathbf{x}) = \mathbf{T}^L \mathbf{W}^L (\sigma(\mathbf{T}^{L-1} \mathbf{W}^{L-1} (\dots \sigma(\mathbf{T}^1 \mathbf{W}^1 \mathbf{x}))))), \quad (3)$$

where $\mathbf{x} \in \mathbb{R}^D$ is an input, and $\sigma(\cdot)$ is a nonlinear activation function. T-nets get their name from transformation matrices (\mathbf{T}) because the linear transformation defined by a \mathbf{T} plays a crucial role in meta-learning. Note that a cell has the same expressive power as a linear layer. Model parameters θ are therefore a collection of \mathbf{W} ’s and \mathbf{T} ’s, i.e.,

$$\theta = \left\{ \underbrace{\mathbf{W}^1, \dots, \mathbf{W}^L}_{\theta_{\mathbf{W}}}, \underbrace{\mathbf{T}^1, \dots, \mathbf{T}^L}_{\theta_{\mathbf{T}}} \right\}.$$

Transformation parameters $\theta_{\mathbf{T}}$, which are shared across task-specific models, are determined by the meta-learner. All task-specific learners share the same initial $\theta_{\mathbf{W}}$ but update to

*For convolutional cells, \mathbf{W} is a convolutional layer with some size and stride and \mathbf{T} is a 1×1 convolution that doesn’t change the number of channels

Algorithm 1 Transformation Networks (T-net)

Require: $p(\mathcal{T})$
Require: α, β

- 1: randomly initialize θ
- 2: **while** not done **do**
- 3: Sample batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$
- 4: **for all** \mathcal{T}_j **do**
- 5: **for** $i = 1, \dots, L$ **do**
- 6: Compute $\tilde{\mathbf{W}}_{\mathcal{T}}$ according to (4)
- 7: **end for**
- 8: $\tilde{\theta}_{\mathbf{W}, \mathcal{T}_j} = \{\tilde{\mathbf{W}}_{\mathcal{T}_j}^1, \dots, \tilde{\mathbf{W}}_{\mathcal{T}_j}^L\}$
- 9: **end for**
- 10: $\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_j \mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathbf{W}, \mathcal{T}_j}, \theta_{\mathbf{T}}, \mathcal{D}_{\mathcal{T}_j, \text{test}})$
- 11: **end while**

different values since each uses their corresponding train set $\mathcal{D}_{\mathcal{T}, \text{train}}$. Thus we denote such (adjusted) parameters for task \mathcal{T} as $\tilde{\theta}_{\mathbf{W}, \mathcal{T}}$. Though they may look similar, \mathcal{T} denotes a task while \mathbf{T} denotes a transformation matrix.

Given a task \mathcal{T} , each \mathbf{W} is adjusted with the gradient update

$$\tilde{\mathbf{W}}_{\mathcal{T}} \leftarrow \mathbf{W} - \alpha \nabla_{\mathbf{W}} \mathcal{L}_{\mathcal{T}}(\theta_{\mathbf{W}}, \theta_{\mathbf{T}}, \mathcal{D}_{\mathcal{T}, \text{train}}). \quad (4)$$

Again, $\tilde{\theta}_{\mathbf{W}, \mathcal{T}}$ is defined as $\{\tilde{\mathbf{W}}_{\mathcal{T}}^1, \dots, \tilde{\mathbf{W}}_{\mathcal{T}}^L\}$. Using the task-specific learner $\tilde{\theta}_{\mathbf{W}, \mathcal{T}}$, the meta-learner improves itself with the gradient update

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \left(\sum_{\mathcal{T} \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathbf{W}, \mathcal{T}}, \theta_{\mathbf{T}}, \mathcal{D}_{\mathcal{T}, \text{test}}) \right). \quad (5)$$

$\alpha > 0$ and $\beta > 0$ are learning rate hyperparameters. We show our full algorithm in Algorithm 1.

To evaluate on a new task \mathcal{T}_* , we do the following. We compute task-specific parameters $\tilde{\theta}_{\mathbf{W}, \mathcal{T}_*}$ using (4), starting from the meta-learned initial value $\theta_{\mathbf{W}}$. We report the loss of task-specific parameters $\tilde{\theta}_{\mathbf{W}, \mathcal{T}_*}$ on the test set $\mathcal{D}_{\mathcal{T}_*, \text{test}}$.

We now briefly examine a single cell:

$$\mathbf{y} = \mathbf{T}\mathbf{W}\mathbf{x},$$

where \mathbf{x} is the input to the cell and \mathbf{y} its output. The squared length of a change in output $\Delta \mathbf{y} = \mathbf{y}^* - \mathbf{y}_0$ is calculated as

$$\|\Delta \mathbf{y}\|^2 = ((\Delta \mathbf{W})\mathbf{x})^\top (\mathbf{T}^\top \mathbf{T}) ((\Delta \mathbf{W})\mathbf{x}), \quad (6)$$

where $\Delta \mathbf{W}$ is similarly defined as $\mathbf{W}^* - \mathbf{W}_0$. We see here that the magnitude of $\Delta \mathbf{y}$ is determined by the interaction between $(\Delta \mathbf{W})\mathbf{x}$ and $\mathbf{T}^\top \mathbf{T}$. Since a task-specific learner performs gradient descent only on \mathbf{W} and not \mathbf{T} , the change in \mathbf{y} resulting from (4) is guided by the meta-learned value $\mathbf{T}^\top \mathbf{T}$. We provide a more precise analysis of this behavior in Section 4.




Figure 3. A diagram of the adaptation process of a Mask Transformation Network (MT-net). Blue values are meta-learned and shared across all tasks. Orange values are different for each task.

3.2. MT-net

The MT-net is built on the same feedforward model (3) as the T-net:

$$f_{\theta}(\mathbf{x}) = \mathbf{T}^L \mathbf{W}^L (\sigma(\mathbf{T}^{L-1} \mathbf{W}^{L-1} (\dots \sigma(\mathbf{T}^1 \mathbf{W}^1 \mathbf{x}))))). \quad (7)$$

The MT-net differs from the T-net in the binary mask applied to the gradient update to determine which parameters are to be updated. The update rule for task-specific parameters $\tilde{\mathbf{W}}_{\mathcal{T}}$ is given by

$$\tilde{\mathbf{W}}_{\mathcal{T}} \leftarrow \mathbf{W} - \alpha \mathbf{M} \odot \nabla_{\mathbf{W}} \mathcal{L}(\theta_{\mathbf{W}}, \theta_{\mathbf{T}}, \mathcal{D}_{\mathcal{T}, \text{train}}), \quad (8)$$

where \odot is the Hadamard (elementwise) product between matrices of the same dimension. \mathbf{M} is a binary gradient mask which is sampled each time the task-specific learner encounters a new task. Each row of \mathbf{M} is either an all-ones vector $\mathbf{1}$ or an all-zeros vector $\mathbf{0}$. We parameterize the probability of row j in \mathbf{M} being $\mathbf{1}$ with a scalar variable ζ_j :

$$\begin{aligned} \mathbf{M} &= [\mathbf{m}_1, \dots, \mathbf{m}_n]^{\top}, \\ \mathbf{m}_j^{\top} &\sim \text{Bern}\left(\frac{\exp(\zeta_j)}{\exp(\zeta_j) + 1}\right) \mathbf{1}^{\top}, \end{aligned} \quad (9)$$

where $\text{Bern}(\cdot)$ denotes the Bernoulli distribution. Each logit ζ acts on a row of a weight matrix \mathbf{W} , so weights that contribute to the same immediate activation are updated or not updated together.

We approximately differentiate through the Bernoulli sampling of masks using the Gumbel-Softmax estimator (Jang et al., 2017; Maddison et al., 2017):

$$g_1, g_2 \sim \text{Gumbel}(0, 1), \quad (10)$$

$$\mathbf{m}_j^{\top} \leftarrow \frac{\exp\left(\frac{\zeta_j + g_1}{c}\right)}{\exp\left(\frac{\zeta_j + g_1}{c}\right) + \exp\left(\frac{g_2}{c}\right)} \mathbf{1}^{\top}, \quad (11)$$

Algorithm 2 Mask Transformation Networks (MT-net)

Require: $p(\mathcal{T})$
Require: α, β

- 1: randomly initialize θ
- 2: **while** not done **do**
- 3: Sample batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$
- 4: **for all** \mathcal{T}_j **do**
- 5: **for** $i = 1, \dots, L$ **do**
- 6: Sample binary mask \mathbf{M}^i according to (11)
- 7: Compute $\tilde{\mathbf{W}}_{\mathcal{T}_j}^i$ according to (8)
- 8: **end for**
- 9: $\tilde{\theta}_{\mathbf{W}, \mathcal{T}_j} = \{\tilde{\mathbf{W}}_{\mathcal{T}_j}^1, \dots, \tilde{\mathbf{W}}_{\mathcal{T}_j}^L\}$
- 10: **end for**
- 11: $\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_j \mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathbf{W}, \mathcal{T}}, \theta_{\mathbf{T}}, \theta_{\zeta}, \mathcal{D}_{\mathcal{T}, \text{test}})$
- 12: **end while**

where c is a temperature hyperparameter. This reparameterization allows us to directly backpropagate through the mask. At the limit of $c \rightarrow 0$, (11) follows the behavior of (9).

As in T-nets, we denote the collection of altered weights as $\tilde{\theta}_{\mathbf{W}, \mathcal{T}} = \{\tilde{\mathbf{W}}_{\mathcal{T}}^1, \dots, \tilde{\mathbf{W}}_{\mathcal{T}}^L\}$. The meta-learner learns all parameters θ :

$$\theta = \left\{ \underbrace{\mathbf{W}^1, \dots, \mathbf{W}^L}_{\theta_{\mathbf{W}}}, \underbrace{\mathbf{T}^1, \dots, \mathbf{T}^L}_{\theta_{\mathbf{T}}}, \underbrace{\zeta^1, \dots, \zeta^L}_{\theta_{\zeta}} \right\}. \quad (12)$$

As in a T-net, the meta-learner performs stochastic gradient descent on $\mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathbf{W}, \mathcal{T}}, \theta_{\mathbf{T}}, \theta_{\zeta}, \mathcal{D}_{\mathcal{T}, \text{test}})$:

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \left(\sum_{\mathcal{T} \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathbf{W}, \mathcal{T}}, \theta_{\mathbf{T}}, \theta_{\zeta}, \mathcal{D}_{\mathcal{T}, \text{test}}) \right). \quad (13)$$

The full algorithm is shown in Algorithm 2.

We emphasize that the binary mask used for task-specific learning (\mathbf{M}) depends on meta-learned parameter weights (ζ). Since the meta-learner optimizes the loss in a task after a gradient step (8), the matrix \mathbf{M} gets assigned a high probability of having value 1 for weights that are meant to encode task-specific information. Furthermore, since we update \mathbf{M} along with model parameters \mathbf{W} and \mathbf{T} , the meta-learner is incentivized to learn configurations of \mathbf{W} and \mathbf{T} in which there exists a clear divide between task-specific and task-mutual neurons.

4. Analysis

In this section, we provide further analysis of the update schemes of T-nets and MT-nets.

We analyse how the activation space of a single cell of a T-net or MT-net behaves during task-specific learning. More specifically, we make precise how \mathbf{W} encodes a learned curvature matrix. By using such an analysis to reason about a whole network consisting of several cells, we are implicitly approximating the full curvature matrix of the network by a block-diagonal curvature matrix. In this approximation, second-order interactions only occur among weights in the same layer (or cell). Previous works (Heskes, 2000; Martens & Grosse, 2015; Desjardins et al., 2015) have used such an approximation of the curvature of a neural network.

4.1. T-nets Learn a Metric in Activation Space

We consider a cell in a T-net where the pre-activation value \mathbf{y} is given by

$$\mathbf{y} = \mathbf{T}\mathbf{W}\mathbf{x} = \mathbf{A}\mathbf{x}, \quad (14)$$

where $\mathbf{A} = \mathbf{T}\mathbf{W}$ and \mathbf{x} is the input to the cell. We omit superscripts throughout this section.

A standard feedforward network resorts to the gradient of a loss function $\mathcal{L}_{\mathcal{T}}$ (which involves a particular task $\mathcal{T} \sim p(\mathcal{T})$) with respect to the parameter matrix \mathbf{A} , to update model parameters. In such a case, a single gradient step yields

$$\begin{aligned} \mathbf{y}^{\text{new}} &= (\mathbf{A} - \alpha \nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}}) \mathbf{x} \\ &= \mathbf{y} - \alpha \nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}} \mathbf{x}. \end{aligned} \quad (15)$$

The update of a T-net (4) results in the following new value of \mathbf{y} :

$$\begin{aligned} \mathbf{y}^{\text{new}} &= \mathbf{T} (\mathbf{T}^{-1} \mathbf{A} - \alpha \nabla_{\mathbf{T}^{-1} \mathbf{A}} \mathcal{L}_{\mathcal{T}}) \mathbf{x} \\ &= \mathbf{y} - \alpha (\mathbf{T} \mathbf{T}^{\top}) \nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}} \mathbf{x}, \end{aligned} \quad (16)$$

where \mathbf{T} is determined by the meta-learner. Thus, in a T-net, the incremental change of \mathbf{y} is proportional to the negative of the gradient $(\mathbf{T} \mathbf{T}^{\top}) \nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}}$, while the standard feedforward net resorts to a step proportional to the negative of $\nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}}$. Task-specific learning in the T-net is guided by a full rank metric in each cell’s activation space, which is determined by each cell’s transformation matrix \mathbf{T} . This metric $(\mathbf{T} \mathbf{T}^{\top})^{-1}$ warps (scaling, rotation, etc.) the activation space of the model so that in this warped space, a single gradient step with respect to the loss of a new task yields parameters that are well suited for that task.

4.2. MT-nets Learn a Subspace with a Metric

We now consider MT-nets and analyze what their update (8) means from the viewpoint of $\mathbf{y} = \mathbf{T}\mathbf{W}\mathbf{x} = \mathbf{A}\mathbf{x}$.

MT-nets can restrict its task-specific learner to any subspace of its gradient space:

Proposition 1. *Fix \mathbf{x} and \mathbf{A} . Let $\mathbf{y} = \mathbf{T}\mathbf{W}\mathbf{x}$ be a cell in an MT-net and let ζ be its corresponding mask parameters. Let \mathbf{U} be a d -dimensional subspace of \mathbb{R}^n ($d \leq n$). There exist configurations of \mathbf{T} , \mathbf{W} , and ζ such that the span of $\mathbf{y}^{\text{new}} - \mathbf{y}$ is \mathbf{U} while satisfying $\mathbf{A} = \mathbf{T}\mathbf{W}$.*

Proof. See Appendix B. \square

This proposition states that \mathbf{W} , \mathbf{T} , and ζ have sufficient expressive power to restrict updates of \mathbf{y} to any subspace. Note that this construction is only possible because of the transformation \mathbf{T} ; if we only had binary masks \mathbf{M} , we would only be able to restrict gradients to axis-aligned subspaces.

In addition to learning a subspace that we project gradients onto (\mathbf{U}), we are also learning a metric in this subspace. We first provide an intuitive exposition of this idea.

We unroll the update of an MT-net as we did with T-nets in (16):

$$\begin{aligned} \mathbf{y}^{\text{new}} &= \mathbf{T} ((\mathbf{T}^{-1} \mathbf{A} - \alpha \mathbf{M} \odot \nabla_{\mathbf{T}^{-1} \mathbf{A}} \mathcal{L}_{\mathcal{T}}) \mathbf{x}) \\ &= \mathbf{y} - \alpha \mathbf{T} (\mathbf{M} \odot (\mathbf{T}^{\top} \nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}})) \mathbf{x} \\ &= \mathbf{y} - \alpha \mathbf{T} (\mathbf{M}_{\mathbf{T}} \odot \mathbf{T}^{\top}) \nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}} \mathbf{x} \\ &= \mathbf{y} - \alpha (\mathbf{T} \odot \mathbf{M}_{\mathbf{T}}^{\top}) (\mathbf{M}_{\mathbf{T}} \odot \mathbf{T}^{\top}) \nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}} \mathbf{x}. \end{aligned} \quad (17)$$

Where $\mathbf{M}_{\mathbf{T}}$ is an $m \times m$ matrix which has the same columns as \mathbf{M} . Let’s denote $\mathbf{T}_M = \mathbf{M}_{\mathbf{T}} \odot \mathbf{T}^{\top}$. We see that the update of a task-specific learner in an MT-net performs the update $\mathbf{T}_M^{\top} \mathbf{T}_M \nabla_{\mathbf{A}} \mathcal{L}_{\mathcal{T}}$. Note that $\mathbf{T}_M^{\top} \mathbf{T}_M$ is an $n \times n$ matrix that only has nonzero elements in rows and columns where \mathbf{m} is 1. By setting appropriate ζ , we can view $\mathbf{T}_M^{\top} \mathbf{T}_M$ as a full-rank $d \times d$ metric tensor.

This observation can be formally stated as:

Proposition 2. *Fix \mathbf{x} , \mathbf{A} , and a loss function $\mathcal{L}_{\mathcal{T}}$. Let $\mathbf{y} = \mathbf{T}\mathbf{W}\mathbf{x}$ be a cell in an MT-net and let ζ be its corresponding mask parameters. Let \mathbf{U} be a d -dimensional subspace of \mathbb{R}^n , and $g(\cdot, \cdot)$ a metric tensor on \mathbf{U} . There exist configurations of \mathbf{T} , \mathbf{W} , and ζ such that the vector $\mathbf{y}^{\text{new}} - \mathbf{y}$ is in the steepest direction of descent on $\mathcal{L}_{\mathcal{T}}$ with respect to the metric $g(\cdot, \cdot)$.*

Proof. See Appendix B. \square

Therefore, not only can MT-nets project gradients of task-specific learners onto a subspace of the pre-activation (\mathbf{y}) space, they can also learn a metric in that subspace and thereby learning a low-dimensional linear embedding of the activation space. The MT-net update (8) is gradient descent in this low-dimensional embedding, so the meta-objective shown in (13) is minimized when gradient descent in this embedding requires few steps to converge and is sensitive to task identity.

5. Related Work

A successful line of research in few-shot learning uses feed-forward neural networks as learners. These approaches learn update rules (Ravi & Larochelle, 2017; Li & Malik, 2016; Andrychowicz et al., 2016) or directly generate weights (Ha et al., 2016). A related research direction is to learn initial parameters (Finn et al., 2017) while fixing the learning rule to gradient descent, or additionally learning learning rates for each weight (Li et al., 2017). (Grant et al., 2018) interprets such gradient-based meta-learning as hierarchical bayesian inference, and (Finn & Levine, 2017) states that such methods are expressive enough to approximate any learning algorithm.

Our work is closely related to this line of research. Unlike previous work, MT-nets learn how many degrees of freedom the task-specific learner should have at meta-test time. Additionally, while MT-nets learn update rules, these update rules are directly embedded in the network itself instead of being stored in a separate model.

Distance metric learning (Xing et al., 2003; Weinberger et al., 2006) methods learn a distance function between datapoints. Similarly, MT-nets learn a full metric matrix. Whereas those methods required constrained optimization techniques to enforce that the learned matrix represents a metric, our parameterization allows us to directly learn such a metric using gradient descent. Recently, neural networks have been used to learn a metric between images (Koch et al., 2015; Vinyals et al., 2016; Snell et al., 2017), achieving state-of-the-art performance on few-shot classification benchmarks. Unlike these methods, we learn a metric in feature space instead of input space. Our method applies to a larger class of problems including regression and reinforcement learning, since all MT-nets require is a differentiable loss function.

Another line of research in few-shot learning is to use a recurrent neural network (RNN) as a learner (Santoro et al., 2016; Munkhdalai & Yu, 2017). Here, the meta-learning algorithm is gradient descent on an RNN, and the learning algorithm is the update of hidden cells. The (meta-learned) weights of the RNN specify a learning strategy, which processes training data and uses the resulting hidden state vector to make decisions about test data. A recent work that uses temporal convolutions for meta-learning (Mishra et al., 2018) is also closely related to this line of research.

6. Experiments

We performed experiments to answer:

- Do our novel components (**TW**, **M** etc) improve meta-learning performance? (6.1)
- Is applying a mask **M** row-wise actually better than

| Models | 5-shot | 10-shot | 20-shot |
|-----------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| MAML ¹ | 1.07 ± 0.11 | 0.71 ± 0.07 | 0.50 ± 0.05 |
| Meta-SGD ¹ | 0.88 ± 0.14 | 0.53 ± 0.09 | 0.35 ± 0.06 |
| M-net-full | 0.91 ± 0.09 | 0.63 ± 0.07 | 0.38 ± 0.04 |
| M-net | 0.88 ± 0.09 | 0.60 ± 0.06 | 0.41 ± 0.04 |
| T-net | 0.83 ± 0.08 | 0.56 ± 0.06 | 0.38 ± 0.04 |
| MT-net-full | 0.81 ± 0.08 | 0.51 ± 0.05 | 0.35 ± 0.04 |
| MT-net | 0.76 ± 0.09 | 0.49 ± 0.05 | 0.33 ± 0.04 |

Table 1. Loss on sine wave regression. Networks were meta-trained using 10-shot regression tasks. Reported losses were calculated after adaptation using various numbers of examples.

¹ Reported by (Li et al., 2017).

applying one parameter-wise? (6.1)

- To what degree does **T** alleviate the need for careful tuning of step size α ? (6.2)
- In MT-nets, does learned subspace dimension reflect the difficulty of tasks? (6.3)
- Can T-nets and MT-nets scale to large-scale meta-learning problems? (6.4)

Most of our experiments were performed by modifying the code accompanying (Finn et al., 2017), and we follow their experimental protocol and hyperparameters unless specified otherwise.

6.1. Toy Regression Problem

We start with a K -shot regression problem and compare results to previous meta-learning methods (Finn et al., 2017; Li et al., 2017). The details of our regression task are the same as (Li et al., 2017). Each individual task is to regress from the input x to the output y of a sine function

$$y(x) = A \sin(wx + b) \quad (18)$$

For each task, A, w, b are sampled uniformly from $[0.1, 5.0]$, $[0.8, 1.2]$, and $[0, \pi]$, respectively. Each task consists of $K \in \{5, 10, 20\}$ training examples and 10 testing examples. We sample x uniformly from $[-5.0, 5.0]$ for both train and test sets. Our regressor architecture has two hidden cells each with activation size 40. After every **T** is a ReLU nonlinearity. The loss function is the mean squared error (MSE) between the regressor’s prediction $f(x)$ and the true value $y(x)$. We used Adam (Kingma & Ba, 2015) as our meta-optimizer with a learning rate of $\beta = 10^{-3}$. Task-specific learners used step size $\alpha = 10^{-2}$. We initialize all ζ to 0, all **T** as identity matrices, and all **W** as truncated normal matrices with standard deviation 10^{-2} . While we trained our meta-learner with $K = 10$ examples, we tested using various numbers of examples ($K \in \{5, 10, 20\}$).




Figure 4. 10-shot regression tasks to sets of polynomials of various degrees. MT-nets choose to update a larger fraction of weights as the set of tasks gets more complex.

| α | MAML | T-net | MT-net |
|----------|--------------------|-----------------------------------|-----------------------------------|
| 10 | 171.92 ± 25.04 | 4.08 ± 0.30 | 4.18 ± 0.30 |
| 1 | 5.81 ± 0.49 | 4.15 ± 0.30 | 0.61 ± 0.07 |
| 0.1 | 1.05 ± 0.11 | 0.68 ± 0.06 | 0.54 ± 0.05 |
| 0.01 | 0.71 ± 0.07 | 0.56 ± 0.06 | 0.49 ± 0.05 |
| 0.001 | 0.82 ± 0.08 | 0.59 ± 0.06 | 0.59 ± 0.06 |
| 0.0001 | 2.54 ± 0.19 | 0.62 ± 0.06 | 0.72 ± 0.07 |

Table 2. Loss on 10-shot sine wave regression. T-nets and MT-nets are most robust to change in step size α . This is due to the meta-learned matrix \mathbf{T} inside each cell, which alters the effective step size.

We show results in Table 1. To see if each of our novel components increased meta-learning performance, we also performed the same experiments with variations of MT-nets. An M-net uses a mask \mathbf{M} like an MT-net, but each cell consists of a single matrix \mathbf{W} instead of $\mathbf{T}\mathbf{W}$. A model with “-full” at the end of its name learns a separate mask parameter for each weight of \mathbf{W} instead of sharing a mask among weights that contribute to the same activation. For example, if \mathbf{W} has size 5×10 , the corresponding ζ in an MT-net would be of dimension 5, but in MT-net-full, the dimension of ζ would be 50. MT-nets outperform MAML, meta-SGD, and all variations of MT-nets.

6.2. Robustness to learning rate change

The transformation \mathbf{T} of our method can change the effective step size α . We performed experiments to see how robust our method is to variations in α . We perform the same sinusoid experiment as in section 6.1, but with various step sizes ($\alpha \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10\}$). We evaluate on $K = 10$ training examples, and all other settings are identical to the experiments in section 6.1.

We show losses after adaptation of both MAML and MT-nets in Table 2. We can see that MT-nets are more robust to change in step size. This indicates that as shown in section 4.2, the matrix \mathbf{T} is capable of warping the parameter space to recover from suboptimal step size α .

6.3. Task Complexity and Subspace Dimension

We performed this experiment to see whether the dimension of the learned subspace of MT-nets reflect the underlying complexity of its given set of tasks.

We consider 10-shot regression tasks in which the target function is a polynomial. A polynomial regression meta-task consists of polynomials of the same order with various coefficients. To generate a polynomial of order n ($\sum_{i=0}^n c_i x^i$), we uniformly sampled c_0, \dots, c_n from $[-1, 1]$. We used the same network architecture and hyperparameters as in Section 6.1 and performed 10-shot regression for polynomial orders $n \in \{0, 1, 2\}$. Since the number of free parameters is proportional to the order of the polynomial, we expect higher-order polynomials to require more parameters to adapt to. The fraction of parameters that task-specific learners change is calculated as the expected value of $\frac{e^{-\zeta}}{e^{-\zeta} + 1}$ over all logits ζ .

We show results in Figure 4, and additional results in Appendix C. The number of weights that the meta-learner of an MT-net sets to be altered increases as the task gets more complex. We interpret this as the meta-learner of MT-nets having an effect akin to Occam’s razor: it selects a task-specific model of just enough complexity to learn in a set of tasks. This behavior emerges even though we do not introduce any additional loss terms to encourage such behavior. We think this is caused by the noise inherent in stochastic gradient descent. Since the meta-learner of an MT-net can choose whether or not to perform gradient descent in a particular direction, it is incentivized not to do so in directions that are not model-specific, because doing so would introduce more noise into the network parameters and thus (in expectation) suffer more loss.

6.4. Classification

To compare the performance of MT-nets to prior work in meta-learning, we evaluate our method on few-shot classification on the Omniglot (Lake et al., 2015) and MiniImagenet (Ravi & Larochelle, 2017) datasets. We used the miniIma-

| Models | 5-way 1-shot acc. (%) | 20-way 1-shot acc. (%) |
|---|------------------------------------|----------------------------------|
| Matching Networks (Vinyals et al., 2016) | 98.1 | 93.8 |
| Prototypical Networks (Snell et al., 2017) | 97.4 | 92.0 |
| mAP-SSVM (Triantafillou et al., 2017) | 98.6 | 95.4 |
| MAML (Finn et al., 2017) | 98.7 ± 0.4 | 95.8 ± 0.3 |
| Meta-SGD (Li et al., 2017) | 99.53 ± 0.26 | 95.93 ± 0.38 |
| T-net (ours) | 99.4 ± 0.3 | 96.1 ± 0.3 |
| MT-net (ours) | 99.5 ± 0.3 | 96.2 ± 0.4 |

| Models | 5-way 1-shot acc. (%) |
|--|------------------------------------|
| Matching Networks (Vinyals et al., 2016) ¹ | 43.56 ± 0.84 |
| Prototypical Networks (Snell et al., 2017) ² | 46.61 ± 0.78 |
| mAP-SSVM (Triantafillou et al., 2017) | 50.32 ± 0.80 |
| Fine-tune baseline ¹ | 28.86 ± 0.54 |
| Nearest Neighbor baseline ¹ | 41.08 ± 0.70 |
| meta-learner LSTM (Ravi & Larochelle, 2017) | 43.44 ± 0.77 |
| MAML (Finn et al., 2017) | 48.70 ± 1.84 |
| L-MAML (Grant et al., 2018) | 49.40 ± 1.83 |
| Meta-SGD (Li et al., 2017) | 50.47 ± 1.87 |
| T-net (ours) | 50.86 ± 1.82 |
| MT-net (ours) | 51.70 ± 1.84 |

Table 3. Few-shot classification accuracy on (top) held-out Omniglot characters and (bottom) test split of MiniImagenet. \pm represents 95% confidence intervals.

¹ Reported by (Ravi & Larochelle, 2017).

² Reported results for 5-way 1-shot.

genet splits proposed by (Ravi & Larochelle, 2017) in our experiments.

Our CNN model uses the same architecture as (Finn et al., 2017). The model has 4 modules: each has 3×3 convolutions and 64 filters, followed by batch normalization (Ioffe & Szegedy, 2015). As in (Finn et al., 2017), we used 32 filters per layer in miniImagenet. Convolutions have stride 2×2 on Omniglot, and 2×2 max-pooling is used after batch normalization instead of strided convolutions on MiniImagenet. We evaluate with 3, 5, and 10 gradient steps for Omniglot 5-way, Omniglot 20-way, and miniImagenet 5-way, respectively.

Results are shown in Table 3. MT-nets achieve state-of-the-art or comparable performance on both problems. Several works (Mishra et al., 2018; Munkhdalai & Yu, 2017; Sung et al., 2017) have reported improved performance on MiniImagenet using a significantly more expressive architecture. We only report methods that have equal or comparable expressiveness to the model first described in (Vinyals et al., 2016). Not controlling for network expressivity, the highest reported accuracy so far on 5-way 1-shot miniImagenet classification is 57.02 (Sung et al., 2017).

7. Conclusion

We introduced T-nets and MT-nets. One can transform any feedforward neural network into an MT-net, so any future architectural advances can take advantage of our method. Experiments showed that our method alleviates the need for careful tuning of the learning rate in few-shot learning problems and that the mask M reflects the complexity of the set of tasks it is learning to adapt in. MT-nets also showed state-of-the-art performance in a challenging few-shot classification benchmark (MiniImagenet).

While we think MT-nets are a gradient-based meta-learning method, our analysis has shown that it has some interesting commonalities with optimizer learning methods such as (Ravi & Larochelle, 2017). We will investigate this connection between two seemingly disparate approaches to meta-learning in future work.

One of the biggest weaknesses of deep networks is that they are very data intensive. By learning what to learn when a new task is encountered, we can train networks with high capacity using a small amount of data. We believe that designing effective gradient-based meta-learners will be beneficial not just for the few-shot learning setting, but also machine learning problems in general.

Acknowledgements

SC was supported by Samsung DS Software Center, Samsung Research, and Naver.

References

- Amari, S. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.
- Andrychowicz, M., Denil, M., Gómez Colmenarejo, S., Hoffman, M. W., Pfau, D., Schaul, T., Shillingford, B., De Freitas, N., and Deepmind, G. Learning to learn by gradient descent by gradient descent. *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- Desjardins, G., Simonyan, K., Pascanu, R., et al. Natural neural networks. 2015.
- Edwards, H. and Storkey, A. Towards a Neural Statistician. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Finn, C. and Levine, S. Meta-learning and universality: Deep representations and gradient descent can approximate any learning algorithm, 2017. *Preprint arXiv:1710.11622*.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, Sydney, Australia, 2017.
- Fort, S. Gaussian Prototypical Networks for Few-Shot Learning on Omniglot, 2017. *Preprint arXiv:1708.02735*.
- Garcia, V. and Bruna, J. Few-Shot Learning with Graph Neural Networks, 2017. *Preprint arXiv:1711.04043*.
- Grant, E., Finn, C., Levine, S., Darrell, T., and Griffiths, T. Recasting gradient-based meta-learning as hierarchical bayes. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.
- Hariharan, B. and Girshick, R. B. Low-shot visual object recognition by shrinking and hallucinating features. *Proceedings of the International Conference on Computer Vision (ICCV)*, 2016.
- Heskes, T. On natural learning and pruning in multilayered perceptrons. *Neural Computation*, 2000.
- Ioffe, S. and Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- Jang, E., Gu, S., and Poole, B. Categorical Reparameterization with Gumbel-Softmax. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Kaiser, L., Nachum, O., Roy, A., and Bengio, S. Learning to Remember Rare Events. *ICLR*, 2017.
- Kingma, D. P. and Ba, J. L. ADAM: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015.
- Koch, G., Zemel, R., and Salakhutdinov, R. Siamese neural networks for one-shot image recognition. In *Proceedings of the International Conference on Machine Learning (ICML)*, Lille, France, 2015.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 2015.
- Li, D., Yang, Y., Song, Y.-Z., and Hospedales, T. M. Learning to Generalize: Meta-Learning for Domain Generalization. *Proceedings of the AAAI National Conference on Artificial Intelligence (AAAI)*, 2017.
- Li, K. and Malik, J. Learning to Optimize. *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- Li, Z., Zhou, F., Chen, F., and Li, H. Meta-SGD: Learning to Learn Quickly for Few Shot Learning, 2017. *Preprint arXiv:1707.09835*.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Martens, J. and Grosse, R. Optimizing neural networks with kronecker-factored approximate curvature. 2015.
- Mishra, N., Rohaninejad, M., Chen, X., and Abbeel, P. A simple neural attentive meta-learner. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. *Nature*, 518, 2015.

- Munkhdalai, T. and Yu, H. Meta networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, Sydney, Australia, 2017.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, Toulon, France, 2017.
- Ren, M., Ravi, S., Triantafillou, E., Snell, J., Swersky, K., Tenenbaum, J. B., Larochelle, H., and Zemel, R. S. Meta-learning for semi-supervised few-shot classification. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Robbins, H. and Monro, S. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3): 400–407, 1951.
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., and Lillicrap, T. Meta-learning with memory-augmented neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, New York, NY, USA, 2016.
- Schmidhuber, J. *Evolutionary Principles in Self-Referential Learning*. PhD thesis, Technical University of Munich, 1987.
- Schmidhuber, J., Zhao, J., and Wiering, M. Shifting inductive bias with success-story algorithm, adaptive levin search, and incremental self-improvement. *Machine Learning*, 28:105–130, 1997.
- Snell, J., Swersky, K., and Zemel, R. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems (NIPS)*, volume 30, 2017.
- Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P. H. S., and Hospedales, T. M. Learning to Compare: Relation Network for Few-Shot Learning, 2017. *Preprint arXiv:1711.06025*.
- Thrun, S. and Pratt, L. *Learning to Learn*. Kluwer Academic Publishers, 1998.
- Triantafillou, E., Zemel, R., and Urtasun, R. Few-shot learning through an information retrieval lens. *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Tsai, Y.-H. H. and Salakhutdinov, R. Improving One-Shot Learning through Fusing Side Information, 2017. *Preprint arXiv:1710.08347*.
- Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., and Wierstra, D. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems (NIPS)*, volume 29, 2016.
- Weinberger, K. Q., Blitzer, J., and Saul, L. K. Distance metric learning for large margin nearest neighbor classification. In *Advances in Neural Information Processing Systems (NIPS)*, volume 18, 2006.
- Xing, E. P., Ng, A. Y., Jordan, M. I., and Russel, S. Distance metric learning, with application to clustering with side-information. In *Advances in Neural Information Processing Systems (NIPS)*, volume 15. MIT Press, 2003.