# Supplementary Material for Interventional Few-Shot Learning 

Zhongqi Yue ${ }^{1,3}$, Hanwang Zhang ${ }^{1}$, Qianru Sun ${ }^{2}$, Xian-Sheng Hua ${ }^{3}$<br>${ }^{1}$ Nanyang Technological University, ${ }^{2}$ Singapore Management University, ${ }^{3}$ Alibaba Group yuez0003@ntu.edu.sg, hanwangzhang@ntu.edu.sg,<br>qianrusun@smu.edu.sg, xiansheng.hxs@alibaba-inc.com

This supplementary material is organized as follows:

- Section A.1 details our analysis in Section 2.3 by showing many-shot learning converges to true causal effect through instrumental variable (IV);
- Section A. 2 gives the derivation for the backdoor adjustment formula in Eq. (1);
- Section A. 3 presents the detailed derivation for the NWGM approximation used in Eq. (3) and (4);
- Section A.4 includes the algorithms for adding IFSL to fine-tuning and meta-learning;
- Section A.5 shows the implementation details for pre-training (Section A.5.1), fine-tuning (Section A.5.2) and meta-learning (Section A.5.3;
- Section A. 6 includes additional experimental results on Conventional Acc (Section A.6.1), Hardness-Specific Acc (Section A.6.2), CAM-Acc (Section A.6.3) and cross-domain evaluation (Section A.6.4).


## A. 1 Instrumental Variable

In this section, we will show that in our causal graph for many-shot learning, the sampling ID $I$ is essentially an instrumental variable for $X \rightarrow Y$ that achieves $P(Y \mid I) \approx P(Y \mid d o(X))$. Before introducing instrumental variable, we first formally define $d$-separation [7], which gives a criterion to study the dependencies between nodes (data variables) in any structural causal model.
d-separation. A set of nodes $Z$ blocks a path $p$ if and only if 1) $p$ contains a chain $A \rightarrow B \rightarrow C$ or a fork $A \leftarrow B \rightarrow C$ and the middle node $B$ is in $Z ; 2) p$ contains a collider $A \rightarrow B \leftarrow C$ such that the middle node $B$ and its descendants are not in $Z$. If conditioning on $Z$ blocks every path between $X$ and $Y$, we say $X$ and $Y$ are $d$-separated conditional on $Z$, i.e., $X$ and $Y$ are independent given $Z$ $(X \Perp Y \mid Z)$.

Instrumental Variable. For a structual causal model $\mathcal{G}$, a variable Z is an instrumental variable (IV) to $X \rightarrow Y$ by satisfying the graphical criteria [9]: 1) $\left.(Z \Perp Y)_{\mathcal{G}_{X}} ; 2\right)(Z \not \Perp X)_{\mathcal{G}}$, where $\mathcal{G}_{\bar{X}}$ is the manipulated graph where all incoming arrows to node $X$ are deleted. For the SCM of many-shot learning in Figure 4(a), it is easy to see that $I$ satisfies both criteria and therefore it is an IV for $X \rightarrow Y$. However, in the few-shot SCM in Figure 4(b), the paths $I \leftarrow X \leftarrow D \rightarrow C \rightarrow Y$ and $I \leftarrow X \rightarrow C \rightarrow Y$ are not blocked in $\mathcal{G}_{\bar{X}}$, which means the first criterion is not met $(I \nVdash Y)_{\mathcal{G}_{X}}$ and $I$ is not an instrumental variable in the few-shot learning case.
Instrumental variable can help find the true causal effect even in the presence of confounder. This is due to the collider junction that makes the IV and confounder independent ( $I \Perp D$ in Figure 4(a)). To see this, we will first consider a simplified case of Figure 4(a) where each causal link represents a linear relationship and we aim to find the true causal effect from $X \rightarrow Y$ through linear regression. Without loss of generality, let $I, X, Y$ take the value of real number. Denote $r_{I X}, r_{X Y}$, and $r_{I Y}$ as the slope of regression line between $I$ and $X, X$ and $Y, I$ and $Y$ respectively. Notice
that $r_{X Y}$ is spurious as it is contaminated by the backdoor path $X \leftarrow D \rightarrow C \rightarrow Y$. However, since the path $I \rightarrow X \leftarrow D \rightarrow C \rightarrow Y$ is blocked due to collider at $X, r_{I Y}$ is free from confounding bias. Therefore $r_{I Y} / r_{I X}$ gives the true causal effect from $X \rightarrow Y$. Similarly, in the classification case of many-shot learning, a classifier is trained to maximize the conditional probability on the IV $P(Y \mid I)$. As the ID-sample matching $I \rightarrow X$ is deterministic, the classifier eventually learns to predict based on the true causal relationship $X \rightarrow Y$. Yet in the complex case of image classification, it is unreasonable to assume linear relationships between variables. In the nonlinear case, it is shown in [2] that observations on IV provide a bound for the true causal effect. This means that learning based on $P(Y \mid I)$ provides an approximation to the true causal effect, i.e. $P(Y \mid I) \approx P(Y \mid \operatorname{do}(X))$.

## A. 2 Derivation of Backdoor Adjustment for the Proposed Causal Graph

We will show the derivation of the backdoor adjustment for the causal graph in Figure 3(a) using the three rules of do-calculus [8].
For a causal directed acyclic graph $\mathcal{G}$, let $X, Y, Z$ and $W$ be arbitrary disjoint sets of nodes. We use $\mathcal{G}_{\bar{X}}$ to denote the manipulated graph where all incoming arrows to node $X$ are deleted. Similarly $\mathcal{G}_{\underline{X}}^{X}$ represents the graph where outgoing arrows from node $X$ are deleted. We use lower case $x, y, z$ and $w$ for specific values taken by each set of nodes: $X=x, Y=y, Z=z$ and $W=w$. For any interventional distribution compatible with $\mathcal{G}$, we have the following three rules:

Rule 1 Insertion/deletion of observations:

$$
\begin{equation*}
P(y \mid d o(x), z, w)=P(y \mid d o(x), w), \operatorname{if}(Y \Perp Z \mid X, W)_{\mathcal{G}_{\bar{X}}} \tag{A1}
\end{equation*}
$$

Rule 2 Action/observation exchange:

$$
\begin{equation*}
P(y \mid d o(x), d o(z), w)=P(y \mid d o(x), z, w), \operatorname{if}(Y \Perp Z \mid X, W)_{\mathcal{G}_{\bar{X} \underline{Z}}} \tag{A2}
\end{equation*}
$$

Rule 3 Insertion/deletion of actions:

$$
\begin{equation*}
P(y \mid d o(x), d o(z), w)=P(y \mid d o(x), w), \operatorname{if}(Y \Perp Z \mid X, W)_{\overline{\mathcal{G}_{\overline{X Z(W)}}}}, \tag{A3}
\end{equation*}
$$

where $Z(W)$ is the set of nodes in $Z$ that are not ancestors of any $W$-node in $\mathcal{G}_{\bar{X}}$.
In our causal graph, the desired interventional distribution $P(Y \mid d o(X=\mathbf{x}))$ can be derived by:

$$
\begin{align*}
P(Y \mid d o(\mathbf{x})) & =\sum_{d} P(Y \mid d o(X=\mathbf{x}), D=d) P(D=d \mid d o(X=\mathbf{x}))  \tag{A4}\\
& =\sum_{d} P(Y \mid d o(X=\mathbf{x}), D=d) P(D=d)  \tag{A5}\\
& =\sum_{d} P(Y \mid X=\mathbf{x}, D=d) P(D=d)  \tag{A6}\\
& =\sum_{d} \sum_{c} P(Y \mid X=\mathbf{x}, D=d, C=c) P(C=c \mid X=\mathbf{x}, D=d) P(D=d)  \tag{A7}\\
& =\sum_{d} P(Y \mid X=\mathbf{x}, D=d, C=g(\mathbf{x}, d)) P(D=d) \tag{A8}
\end{align*}
$$

where Eq. A4, and Eq. A7) follow the law of total probability; Eq. A5) uses Rule 3 given $D \Perp X$ in $\mathcal{G}_{\bar{X}}$; Eq. A6) uses Rule 2 to change the intervention term to observation as $(Y \Perp X \mid D)$ in $\mathcal{G}_{\underline{X}}$; Eq. (A8) is because in our causal graph, $C$ takes a deterministic value given by function $g(\mathbf{x}, d)$. This reduces summation over all values of $C$ in Eq. A7, to a single probability measure in Eq. A8).

## A. 3 Derivation of NWGM Approximation

We will show the derivation of NWGM approximation used in Eq. (3) and (4). In a $K$-way FSL problem, let $f(\cdot)$ be a classifier function that calculates logits for $K$ classes and $\sigma$ be the softmax function over $K$ classes. The approximation effectively moves the outer expectation inside the classifier function: $\mathbb{E}[\sigma(f(\cdot))] \approx \sigma(f(\mathbb{E}[\cdot]))$.

We will first show the derivation for moving the expectation inside softmax function, i.e., $\mathbb{E}[\sigma(f(\cdot))] \approx \sigma(\mathbb{E}[f(\cdot)])$. Without loss of generality, the backdoor adjustment formula in Eq. (3) and Eq. (4) can be written as:

$$
\begin{equation*}
P(Y=y \mid d o(X=\mathbf{x}))=\sum_{d \in D} \sigma\left(f_{y}(\mathbf{x} \oplus \mathbf{c})\right) P(d) \tag{A9}
\end{equation*}
$$

where $D$ represents the set of stratifications, $f_{y}$ is the classifier logit for class $y, \mathbf{c}=g(\mathbf{x}, d)$ is the feature concatenated to $\mathbf{x}$ in Eq. (3) and (4) and $P(d)$ is the prior for each stratificaction.
It is shown in [1] that Eq. A9) can be approximated by the Normalized Weighted Geometric Mean (NWGM) as:

$$
\begin{align*}
\sum_{d \in D} \sigma\left(f_{y}(\mathbf{x} \oplus \mathbf{c})\right) P(d) & \approx N W G M_{d \in D}\left(\sigma\left(f_{y}(\mathbf{x} \oplus \mathbf{c})\right)\right)  \tag{A10}\\
& =\frac{\prod_{d}\left[\exp \left(f_{y}(\mathbf{x} \oplus \mathbf{c})\right)\right]^{P(d)}}{\sum_{i=1}^{K} \prod_{d}\left[\exp \left(f_{i}(\mathbf{x} \oplus \mathbf{c})\right)\right]^{P(d)}}  \tag{A11}\\
& =\frac{\exp \left(\sum_{d} f_{y}(\mathbf{x} \oplus \mathbf{c}) P(d)\right)}{\sum_{i=1}^{K} \exp \left(\sum_{d} f_{i}(\mathbf{x} \oplus \mathbf{c}) P(d)\right)}  \tag{A12}\\
& =\sigma\left(\mathbb{E}_{d}\left[f_{y}(\mathbf{x} \oplus \mathbf{c})\right]\right) \tag{A13}
\end{align*}
$$

where Eq. A10 follows [1], Eq. A11) follows the definition of NWGM, Eq. A12 is because $\exp (a)^{b}=\exp (a b)$.
Next we will show the derivation for linear, cosine and $k$-NN classifier to further move the expectation inside the classifier function, i.e., $\sigma(\mathbb{E}[f(\cdot)])=\sigma(f(\mathbb{E}[\cdot]))$.
For the linear classifier, $f(\mathbf{x} \oplus \mathbf{c})=\mathbf{W}_{1} \mathbf{x}+\mathbf{W}_{2} \mathbf{c}$, where $\mathbf{W}_{1}, \mathbf{W}_{2} \in \mathbb{R}^{K \times N}$ denote the learnable weight, $N$ is the feature dimension, which is the same for $\mathbf{x}$ and $\mathbf{c}$ in Eq. (3) and (4). The bias term is dropped as it does not impact our analysis. Now the expectation can be further moved inside the classifier function through:

$$
\begin{align*}
\left.\sum_{d} f(\mathbf{x} \oplus \mathbf{c})\right) P(d) & =\sum_{d}\left(\mathbf{W}_{1} \mathbf{x}+\mathbf{W}_{2} \mathbf{c}\right) P(d)  \tag{A14}\\
& =\mathbf{W}_{1} \mathbf{x}+\sum_{d} \mathbf{W}_{2} \mathbf{c} P(d)  \tag{A15}\\
& =f\left(\mathbf{x} \oplus \sum_{d} \mathbf{c} P(d)\right) \tag{A16}
\end{align*}
$$

where Eq. A15 is because the feature vector $\mathbf{x}$ is the same for all $d$ and $\mathbb{E}_{d}[\mathbf{x}]=\mathbf{x}$.
For the cosine classifier, $f(\mathbf{x} \oplus \mathbf{c})=\left(\mathbf{W}_{1} \mathbf{x}+\mathbf{W}_{2} \mathbf{c}\right) /\|\mathbf{x} \oplus \mathbf{c}\|\|\mathbf{W}\|$, where $\mathbf{W} \in \mathbb{R}^{K \times 2 N}$ is the concatenation of $\mathbf{W}_{1}$ and $\mathbf{W}_{2}$. In the special case where $\mathbf{x}$ and $\mathbf{c}$ are unit vector, $\|\mathbf{x} \oplus \mathbf{c}\|$ is $\sqrt{2}$ and the cosine classifier function reduces to a linear combination of terms involving only $\mathbf{x}$ and only $\mathbf{c}$. From there, the analysis for linear classifier follows and we have $\sigma(\mathbb{E} f(\cdot))=\sigma(f(\mathbb{E} \cdot))$ for cosine classifier. In the general case where $\mathbf{x}$ and $\mathbf{c}$ are not unit vector, moving the expectation inside cosine classifier function is an approximation $\sigma(\mathbb{E}[f(\cdot)]) \approx \sigma(f(\mathbb{E}[\cdot]))$.
For the $k$-NN classifier, our implementation calculates class centroids using the mean feature of the $K$ support sets and then uses the nearest centroid for prediction (1-NN). Specifically, let $\mathbf{x}$ be a feature vector and $\mathbf{x}^{\prime}$ be the $i$ th class centroid, $i \in\{1, \ldots, K\}$. The logit for class $i$ is given by $f_{i}(\mathbf{x})=-\left\|\mathbf{x}-\mathbf{x}^{\prime}\right\|^{2}$. It is shown in [11] that $k$-NN classifier that uses squared Euclidean distance to generate logits is equivalent to a linear classifier with a particular parameterization. Therefore, our analysis on linear classifier follows for $k-\mathrm{NN}$.
In summary, the derivation of $\mathbb{E}[\sigma(f(\cdot))] \approx \sigma(f(\mathbb{E}[\cdot]))$ is a two-stage process where we first move the expectation inside the softmax function and then further move it inside the classifier function.

## A. 4 Algorithms for Fine-tuning and Meta-Learning with IFSL

In this section, we will briefly revisit the settings of fine-tuning and meta-learning and introduce how to integrate IFSL into them.

In fine-tuning, the goal is to train a classifier $\theta$ conditioned on the current support set $\mathcal{S}=$ $\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n_{s}}$, where $\mathbf{x}_{i}$ is the feature generated by $\Omega$ for $i$ th sample, $y_{i}$ is the ground-truth label for $i$ th sample and $n_{s}$ is the support set size. This is achieved by first predicting the support label $\hat{y}$ using the classifier $P(y \mid \mathbf{x} ; \theta)$. Then with the predicted label $\hat{y}$ and ground-truth label $y$, one can calculate a loss $\mathcal{L}(\hat{y}, y)$ (usually cross-entropy loss) to update the classifier parameter, e.g.through stochastic gradient descent. Adding IFSL to fine-tuning is simple: 1) Pick an adjustment strategy introduced in Section 3. Each implementation defines the set of pre-trained knowledge stratifications $D$, function form of $g(X, D)$, function form of $P(Y \mid X, D, C)$ and the prior $P(D) ; 2)$ The classifier prediction is now based on $P(Y \mid d o(X) ; \theta)$. The process of fine-tuning with IFSL is summarized in Algorithm 1 Note that for the non-parametric $k$-NN classifier, the fine-tuning process is not applicable. When adding IFSL to $k-\mathrm{NN}$, each sample is represented by the adjusted feature instead of original feature $\mathbf{x}$. Please refer to the classifier inputs in Eq. (2), (3) and (4) for the exact form of adjusted feature.
In meta-learning, the goal is to learn the additional "learning behavior" parameterized by $\phi$ using training episodes $\left\{\left(\mathcal{S}_{i}, \mathcal{Q}_{i}\right)\right\}$ sampled from training dataset $\mathcal{D}$. The classifier in meta-learning makes predictions by additionally conditioning on the learning behavior, written as $P_{\phi}(y \mid \mathbf{x} ; \theta)$. Within each episode, $\theta$ is first fine-tuned on the support set $\mathcal{S}_{i}$. Then the fine-tuned model is tested on the query set $\mathcal{Q}_{i}$ to obtain the loss $\mathcal{L}_{\phi}\left(\mathcal{S}_{i}, \mathcal{Q}_{i}\right)$ (e.g.using cross-entropy loss). Finally the loss is used to update $\phi$ using an optimizer. It is also easy to integrate IFSL into meta-learning by only changing the classifier from $P_{\phi}(y \mid \mathbf{x} ; \theta)$ to $P_{\phi}(y \mid d o(\mathbf{x}) ; \theta)$. The flow of meta-learning with IFSL is presented in Algorithm 2. Firstly notice that the initialization of $\theta$ in each task may depend on $\phi$ or $\mathcal{S}_{i}$. For example, in MAML [4] $\phi$ essentially defines an initialization of model parameters, and in LEO [10] the initial classifier parameter is generated conditioned on $\phi$ and $\mathcal{S}_{i}$. Secondly, although the fine-tuning of $\theta$ largely follows Algorithm [1] some meta-learning methods additionally utilize meta-knowledge $\phi$. For example, in SIB the gradients for updating $\theta$ are predicted by $\phi$ using unlabelled query features instead of calculated from $\mathcal{L}(\hat{y}, y)$ as in Algorithm 1

```
Algorithm 1: Fine-tuning + IFSL
Input : \(D\), Support set \(\mathcal{S}=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n_{s}}\)
Output:Fine-tuned classifier parameters \(\theta\)
Initialize \(\theta\);
while not converged do
    for \(i=1, \ldots, n_{s} \mathbf{d o}\)
        for \(d \in D\) do
                Calculate \(c=g(\mathbf{x}, d)\);
                Obtain \(P\left(Y \mid \mathbf{x}_{i}, c, d ; \theta\right), P(d)\)
            Prediction \(\hat{y}_{i}=P(y \mid d o(\mathbf{x}) ; \theta)\);
            Update \(\theta\) using \(\mathcal{L}\left(\hat{y}_{i}, y_{i}\right)\)
return \(\theta\)
```


## A. 5 Implementation Details

## A.5.1 Pre-training

Prior to fine-tuning or meta-learning, we pre-trained a deep neural network (DNN) as feature extractor on the train split of a dataset. We use ResNet-10[5] or WRN-28-10[15] as feature extractor backbone. This section will present the architecture and exact training procedure for our backbones.
Network Architecture. The architecture of our ResNet-10 and WRN-28-10 backbone is shown in Figure A1. Specifically, each convolutional layer is described as " $n \times n$ conv, $p$ ", where $n$ is the kernel size and $p$ is the number of output channels. Convolutional layers with " $/ 2$ " have a stride of 2 and are used to perform downsampling. The solid curved lines represent identity shortcuts, and the dotted lines are projection shortcuts implemented by $1 \times 1$ convolutions. The batch normalization and ReLU layers are omitted in Figure A1 to highlight the key structure of the two backbones.

Pre-training Procedure. The networks are trained from scratch with stochastic gradient descent in a fully-supervised manner, i.e., minimizing cross-entropy loss on the train split of a dataset. Specifically the training is conducted on 90 epochs with early stopping using validation accuracy. We used batch
size of 256 and image size of $84 \times 84$. For data augmentation, a random patch is sampled from an image, resized to $84 \times 84$ and randomly flipped along horizontal axis before used for training. The initial learning rate is set to 0.1 and it is scaled down by factor of 10 every 30 epochs. We used the feature pre-processing methods in [14].

## A.5.2 Fine-Tuning

We consider linear, cosine and $k$-NN classifier for our fine-tuning experiments. In a $K$-way FSL problem, the detailed implementations for the classifier function $f(\mathbf{x})$ are:

Linear. $f(\mathbf{x})=\mathbf{W} \mathbf{x}+\mathbf{b}$, where $\mathbf{x}$ is the input feature, $\mathbf{W} \in \mathbb{R}^{K \times N}$ is the learnable weight parameter, $N$ is the feature dimension and $\mathbf{b} \in \mathbb{R}^{K}$ is the learnable bias parameter.

Cosine. $f(\mathbf{x})=\mathbf{W} \mathbf{x} /\|\mathbf{W}\|\|\mathbf{x}\|$, where $\mathbf{W} \in \mathbb{R}^{K \times N}$ is the learnable weight parameter. We implemented cosine classifier without using the bias term.
$k-\mathbf{N N}$. Our implementation of $k$-NN is similar to [11, 14]. For each of the $K$ classes, we first calculated the average support set feature (centroid) denoted as $\mathbf{x}_{i}, i \in$ $\{1, \ldots, K\}$. The classifier output for class $i$ is then given by $f_{i}(\mathbf{x})=-\left\|\mathbf{x}-\mathbf{x}_{i}\right\|^{2}$. Notice that the prediction given by this classifier will be the nearest centroid.

We froze the backbone and used the average pooling layer output of $\Omega$ to learn the classifier. The output logits from classifier functions are normalized using softmax to gen-


Figure A1: The architecture of our backbones: (a) ResNet-10 [5]; (b) WRN-28 10 [15]. erate probability output $P(y \mid \mathbf{x})$. For linear and cosine classifier, we followed [3] and trained the classifier for 100 iteration with a batch size of 4 . For fine-tuning baseline, we set the learning rate as $1 \times 10^{-2}$ and weight decay as $1 \times 10^{-3}$. For IFSL, we set the learning rate as $5 \times 10^{-3}$ and weight decay as $1 \times 10^{-3}$. $k$-NN classifier is non-parametric and can be initialized directly from support set.

## A.5.3 Meta-Learning

MAML. MAML [4] aims to learn an initialization of network parameters such that it can be finetuned within a few steps to solve a variety of few-shot classification tasks. When using pre-trained network with MAML, it has been shown that learning initialization of the backbone can lead to unsatisfactory performance [3, 12]. Therefore in our experiment, we froze the backbone and appended a 2-layer MLP with ReLU activation in between the hidden layers and a linear classifier after the average pooling layer of $\Omega$. The hidden dimension of the layers in MLP is the same as output dimension of $\Omega$ ( 512 for ResNet-10 and 640 for WRN-28-10). The initialization of MLP and the linear classifier is meta-learnt using MAML. For hyper-parameters, we set the inner loop learning rate $\alpha=0.01$, the outer loop learning rate $\beta=0.01$ and the number of adaptation steps as 20 . For IFSL, we adopted the same hyper-parameter setting and set $n=8$ for feature-wise and combined adjustment. Implementation-wise, we adopted the released cod ${ }^{1}$ [from [3] and performed experiments on MAML without using first-order approximation. Following the implementation in [3], the model was trained on 10,000 randomly sampled tasks with model selection using validation accuracy. We used 2,000 randomly sampled tasks for validation and testing.
MTL. MTL [12] learns scaling and shifting parameters at each convolutional layer of the backbone. We used the MTL implementation released by the autho ${ }^{2}$ which adopts linear classifier. We integrated our ResNet-10 and WRN-28-10 backbones into the released code. The learning rate for scaling and shifting weights $\phi_{S S}$ and initial classifier parameters was set to $1 \times 10^{-4}$ uniformly. We set the inner loop learning rate for classifier as $1 \times 10^{-2}$ and the inner loop update step as 100 . For IFSL, we

[^0]adopted the same hyper-parameter setting and set $n=8$ for feature-wise and combined adjustment. We trained the MTL model on 10,000 randomly sampled tasks with model selection using validation accuracy and used 2,000 randomly sampled tasks for validation and testing. We used 3 RTX 2080 Ti for MTL experiments on WRN-28-10 backbone.
LEO. LEO [10] learns to generate classifier parameters conditioned on support set and the generated parameters are further fine-tuned within each FSL task. Our experiments were conducted on the released code of $\mathrm{LEQ}^{3}$ using linear classifier. Following author's implementation, we saved the center cropped features from our pre-trained backbones and used the saved features to train LEO. For baseline, we used the hyper-parameter settings released by the author. For IFSL, we set $n=8$ for feature-wise and combined adjustment and halved the outer loop learning rate compared to baseline. The model was trained up to 100,000 randomly sampled tasks from training split with early stopping using validation accuracy. We used 2,000 randomly sampled tasks for validation and testing.
Matching Net. Matching Net (MN) [13] is a metric-based method that learns a distance kernel function for $k$-NN. We used the Matching Net implementation in [3]. The implementation follows the setup in [13] and uses LSTM-based fully conditional embedding. We set the learning rate as 0.01 uniformly. For IFSL, we used $n=16$ for feature-wise and combined adjustment. The model was trained using 10,000 randomly sampled tasks with model selection using validation accuracy. We used 2,000 randomly sampled tasks for validation and testing.

SIB. SIB [6] initializes classifier from support set and generates gradients conditioned on unlabelled query set features to update classifier parameters. We followed the SIB implementation released by the author ${ }^{4}$ which uses cosine classifier. In the transductive setting, the query set size is set to 15 . In the inductive setting, we used only 1 query sample randomly selected from the $K$ classes in each episode. In terms of hyper-parameter settings, we took 3 synthetic gradient steps $(K=3)$ for all our experiments. For baseline, the learning rate for SIB network and classifier was set to $1 \times 10^{-3}$ following author's implementation. For IFSL, we set the learning rate to $5 \times 10^{-4}$ and used $n=4$ for feature-wise and combined adjustment. In both transductive and inductive settings, we meta-trained SIB using 50,000 randomly sampled tasks with model selection using validation accuracy. We used 2,000 randomly sampled tasks for validation and testing.

## A. 6 Additional Results

In this section, we include additional results on 1) Conventional Acc in Table A1 supplementary to Table 1; 2) Hardness-Specific Acc in Figure A2 for miniImageNet and Figure A3 for tieredImageNet, supplementary to Figure 5; 3) CAM-Acc in Table A2 supplementary to Figure 6; 4) Cross-Domain Evaluation in Table A3 supplementary to Table 3.

[^1]
## A.6.1 Conventional Acc

Table A1: Supplementary to Table 1. Acc (\%) and 95\% confidence intervals averaged over 2000 5way FSL tasks before and after applying three proposed implementations of adjustment. Specifically, " $f t$ " refers to feature-wise adjustment, "cl" refers to class-wise adjustment and " $f t+c l$ " refers to combined adjustment.


## A.6.2 Hardness-Specific Acc



Figure A2: Supplementary to Figure 5. Hardness-specific Acc of 5-shot fine-tuning and meta-learning on miniImageNet.


Figure A3: Supplementary to Figure 5. Hardness-specific Acc of 5-shot fine-tuning and meta-learning on tieredImageNet.

## A.6.3 CAM-Acc

Table A2: Supplementary to Figure 6. CAM-Acc (\%) on fine-tuning and meta-learning. We used combined adjustment for IFSL.


## A.6.4 Cross-Domain Evaluation

Table A3: Supplementary to Table 3. Acc (\%) and 95\% confidence interval averaged over 2000 5-way FSL tasks on cross-domain evaluation. Specifically, "ft" refers to feature-wise adjustment, "cl" refers to class-wise adjustment and " $f t+c l$ " refers to combined adjustment.

| Method |  |  | ResNet-10 |  | WRN-28-10 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | 5-shot | 1-shot | 5-shot | 1-shot |
| 而 | Linear |  | $58.84 \pm 0.41$ | $42.25 \pm 0.42$ | $62.12 \pm 0.40$ | $42.89 \pm 0.41$ |
|  |  | $f t$ | $60.12 \pm 0.39$ | $42.30 \pm 0.41$ | $63.13 \pm 0.39$ | $43.39 \pm 0.40$ |
|  |  | cl | $60.51 \pm 0.40$ | $42.43 \pm 0.42$ | $62.95 \pm 0.39$ | $44.21 \pm 0.40$ |
|  |  | $f t+c l$ | $60.65 \pm 0.39$ | $45.14 \pm 0.40$ | $64.15 \pm 0.38$ | $45.64 \pm 0.39$ |
|  | Cosine |  | $58.30 \pm 0.39$ | $40.47 \pm 0.40$ | $60.21 \pm 0.39$ | $42.12 \pm 0.39$ |
|  |  | $f t$ | $58.32 \pm 0.39$ | $41.01 \pm 0.40$ | $61.16 \pm 0.38$ | $42.35 \pm 0.41$ |
|  |  | cl | $58.68 \pm 0.39$ | $40.67 \pm 0.41$ | $61.87 \pm 0.40$ | $43.23 \pm 0.40$ |
|  |  | $f t+c l$ | $60.23 \pm 0.38$ | $42.78 \pm 0.40$ | $62.49 \pm 0.38$ | $45.12 \pm 0.39$ |
|  | $k$-NN |  | $57.18 \pm 0.40$ | $38.44 \pm 0.37$ | $59.31 \pm 0.41$ | $40.53 \pm 0.42$ |
|  |  | $f t$ | $59.44 \pm 0.39$ | $43.49 \pm 0.40$ | $62.48 \pm 0.39$ | $45.68 \pm 0.43$ |
|  |  | cl | $58.37 \pm 0.39$ | $43.20 \pm 0.41$ | $62.04 \pm 0.39$ | $45.36 \pm 0.40$ |
|  |  | $f t+c l$ | $59.59 \pm 0.40$ | $43.45 \pm 0.40$ | $62.45 \pm 0.40$ | $45.72 \pm 0.40$ |
|  | MAML [4] |  | $51.09 \pm 0.43$ | $37.20 \pm 0.46$ | $55.04 \pm 0.42$ | $39.06 \pm 0.47$ |
|  |  | $f t$ | $54.95 \pm 0.44$ | $37.34 \pm 0.47$ | $59.57 \pm 0.44$ | $39.25 \pm 0.46$ |
|  |  | cl | $53.62 \pm 0.43$ | $38.13 \pm 0.47$ | $56.80 \pm 0.45$ | $40.32 \pm 0.48$ |
|  |  | $f t+c l$ | $56.71 \pm 0.46$ | $40.36 \pm 0.46$ | $60.89 \pm 0.45$ | $42.16 \pm 0.47$ |
|  | LEO 10 |  | $56.52 \pm 0.46$ | $39.21 \pm 0.53$ | $56.66 \pm 0.48$ | $41.45 \pm 0.54$ |
|  |  | $f t$ | $56.77 \pm 0.48$ | $39.72 \pm 0.54$ | $62.95 \pm 0.47$ | $45.46 \pm 0.55$ |
|  |  | cl | $56.73 \pm 0.47$ | $40.12 \pm 0.55$ | $56.90 \pm 0.47$ | $41.93 \pm 0.56$ |
|  |  | $f t+c l$ | $61.27 \pm 0.46$ | $42.79 \pm 0.52$ | $63.30 \pm 0.47$ | $43.81 \pm 0.56$ |
|  | MTL [12] |  | $56.61 \pm 0.42$ | $41.56 \pm 0.43$ | $56.89 \pm 0.41$ | $43.15 \pm 0.44$ |
|  |  | $f t$ | $61.34 \pm 0.41$ | $42.90 \pm 0.43$ | $63.49 \pm 0.40$ | $45.28 \pm 0.44$ |
|  |  | cl | $60.62 \pm 0.41$ | $42.87 \pm 0.42$ | $62.94 \pm 0.40$ | $45.57 \pm 0.43$ |
|  |  | $f t+c l$ | $62.39 \pm 0.40$ | $44.51 \pm 0.43$ | $65.00 \pm 0.40$ | $46.67 \pm 0.43$ |
|  | MN (13] |  | $53.39 \pm 0.46$ | $40.34 \pm 0.56$ | $53.08 \pm 0.45$ | $42.04 \pm 0.57$ |
|  |  | $f t$ | $54.22 \pm 0.46$ | $40.62 \pm 0.57$ | $54.97 \pm 0.47$ | $42.52 \pm 0.58$ |
|  |  | cl | $53.72 \pm 0.47$ | $40.42 \pm 0.56$ | $53.43 \pm 0.45$ | $42.19 \pm 0.56$ |
|  |  | $f t+c l$ | $56.03 \pm 0.45$ | $41.68 \pm 0.54$ | $58.69 \pm 0.44$ | $43.58 \pm 0.56$ |
|  | SIB [6](transductive) |  | $60.60 \pm 0.46$ | $45.87 \pm 0.55$ | $62.60 \pm 0.49$ | $49.16 \pm 0.58$ |
|  |  | $f t$ | $61.12 \pm 0.45$ | $46.64 \pm 0.55$ | $63.15 \pm 0.47$ | $49.78 \pm 0.56$ |
|  |  | cl | $60.70 \pm 0.46$ | $46.14 \pm 0.56$ | $63.02 \pm 0.48$ | $49.43 \pm 0.57$ |
|  |  | $f t+c l$ | $62.07 \pm 0.44$ | $47.07 \pm 0.53$ | $64.07 \pm 0.49$ | $50.71 \pm 0.54$ |
|  | SIB (6] (inductive) |  | $59.06 \pm 0.42$ | $41.48 \pm 0.43$ | $59.94 \pm 0.42$ | $43.27 \pm 0.44$ |
|  |  | $f t$ | $59.45 \pm 0.41$ | $41.98 \pm 0.44$ | $60.33 \pm 0.44$ | $43.61 \pm 0.45$ |
|  |  | cl | $59.32 \pm 0.42$ | $41.67 \pm 0.43$ | $60.46 \pm 0.43$ | $43.52 \pm 0.45$ |
|  |  | $f t+c l$ | $59.89 \pm 0.41$ | $43.20 \pm 0.43$ | $61.45 \pm 0.43$ | $44.27 \pm 0.44$ |

## References

[1] Pierre Baldi and Peter Sadowski. The dropout learning algorithm. Artificial intelligence, 2014. 3
[2] Alexander Balke and Judea Pearl. Bounds on treatment effects from studies with imperfect compliance. Journal of the American Statistical Association, 1997. 2
[3] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In International Conference on Learning Representations, 2019. 5. 6
[4] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In International Conference on Machine Learning, 2017. 4, 5, 7, 8,9, 10
[5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016. 4.5
[6] Shell Xu Hu, Pablo Moreno, Yang Xiao, Xi Shen, Guillaume Obozinski, Neil Lawrence, and Andreas Damianou. Empirical bayes transductive meta-learning with synthetic gradients. In International Conference on Learning Representations, 2020. 6, 7, 8, 9, 10
[7] J. Pearl, M. Glymour, and N.P. Jewell. Causal Inference in Statistics: A Primer. Wiley, 2016. 1
[8] Judea Pearl. Causal diagrams for empirical research. Biometrika, 1995. 2
[9] Judea Pearl. Causality: Models, Reasoning and Inference. Cambridge University Press, 2nd edition, 2009. 1
[10] Andrei A. Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. In International Conference on Learning Representations, 2019. 4, 6, 7, 8, 9,10
[11] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In Advances in Neural Information Processing Systems, 2017. 3.5
[12] Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. Meta-transfer learning for few-shot learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019. 5, 7, 8, 9, 10
[13] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In Advances in Neural Information Processing Systems, 2016. 6, 7, 8, 9, 10
[14] Yan Wang, Wei-Lun Chao, Kilian Q Weinberger, and Laurens van der Maaten. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. arXiv preprint arXiv:1911.04623, 2019. 5
[15] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In British Machine Vision Conference, 2016. 4,5


[^0]:    ${ }^{1}$ https://github.com/wyharveychen/CloserLookFewShot
    ${ }^{2}$ https://github.com/yaoyao-liu/meta-transfer-learning

[^1]:    ${ }^{3}$ https://github.com/deepmind/leo
    ${ }^{4}$ https://github.com/hushell/sib_meta_learn

