How to trust unlabeled data? Instance Credibility Inference for Few-Shot Learning

Yikai Wang, Li Zhang, Yuan Yao, and Yanwei Fu

Abstract—Deep learning based models have excelled in many computer vision task and appear to surpass humans performance. However, these models require an avalanche of expensive human labeled training data and many iterations to train their large number of parameters. This severely limits their scalability to the real-world long-tail distributed categories, some of which are with a large number of instances, but with only a few manually annotated. Learning from such extremely limited labeled examples is known as Few-shot learning (FSL). Different to prior arts that leverage meta-learning or data augmentation strategies to alleviate this extremely data-scarce problem, this paper presents a statistical approach, dubbed Instance Credibility Inference (ICI) to exploit the support of unlabeled instances for few-shot visual recognition. Typically, we repurpose the self-taught learning paradigm to predict pseudo-labels of unlabeled instances by an initial classifier trained from the *few shot* and then select the most confident ones to augment the training set to re-train the classifier. To do so, we construct a (Generalized) Linear Model (LM/GLM) with incidental parameters to model the mapping from (un-)labeled features to their (pseudo-)labels, in which the sparsity of the incidental parameters indicates the credibility of corresponding pseudo-labeled instance. We rank the credibility of pseudo-labels of unlabeled instances along the regularization path of their corresponding incidental parameters, and the most trustworthy pseudo-labeled examples are preserved as the augmented labeled instances. This process is repeated until all the unlabeled samples are iteratively included in the expanded training set. Theoretically, under mild conditions of restricted eigenvalue, irrepresentability, and large error, our approach is guaranteed to collect *all* the correctly-predicted pseudo-labeled instances from the noisy pseudo-labeled set. Extensive experiments under two few-shot settings show that our approach can establish new state of the art on four widely used few-shot visual recognition benchmark datasets including *mini*ImageNet, *tiered*ImageNet, CIFAR-FS, and CUB. Code and models are released at https://github.com/Yikai-Wang/ICI-FSL.

Index Terms—Few-Shot Learning, Incidental Parameters, Regularization Path, Semi-Supervised Learning, Self-Taught Learning.

✦

1 INTRODUCTION

LEARNING from one or a few examples is an important ability for humans. For example, children have no problem of forming the concept of "giraffe" by only taking a glance from a picture in a book, or hearing its description as looking like a deer with a long neck [1]. In contrast, the most successful recognition systems [2], [3], [4], [5] still highly rely on an avalanche of labeled training data. This thus increases the burden in rare data collection (*e.g.* accident data in the autonomous driving scenario) and expensive data annotation (*e.g.* disease data for medical diagnose), and more fundamentally limits their scalability to open-ended learning of the long tail categories in the real-world.

Motivated by these observations, there has been a recent resurgence of research interest in few-shot learning [6], [7], [8], [9]. It aims to recognize new objects with extremely limited training data for each category. Basically, a few-shot learning model has the chance to access the source/base dataset with many labeled training instances for model training and then is able to generalize to a disjoint but relevant target/novel dataset with only scarce labeled data.

A key challenge for few-shot learning is how to transfer the learned knowledge to new tasks. The simplest strategy is fine-tuning [10], utilizing the limited training instances to update the learned models. Practically, it inevitably causes severely overfitting as one or a few instances are insufficient to model the data distributions of the novel classes. Data augmentation and regularization techniques [11], [12] can alleviate overfitting in such a limited-data regime, but they do not solve it. Several recent efforts are made in leveraging learning to learn, or meta-learning [13] paradigm by simulating the few-shot scenario in the training process [6], [7], [8], [9], [14], [15], [16], [17], [18]. However, Chen *et al.* [19] empirically argues that such a learning paradigm often results in inferior performance compared to a simple baseline with a linear classifier coupled with a deep feature extractor. This phenomenon is also verified in [20].

In real-world applications, unlabeled instances are easier and cheaper to obtain, comparing to the labeled instances which usually require expensive human annotation. Potentially we could utilize the unlabeled instances to alleviate the data-scarce problem and help learn the few-shot model. Specifically, two types of strategies resort to model the data distribution of novel category beyond traditional *inductive* few-shot learning: (i) semi-supervised few-shot learning (SSFSL) [21], [22], [23] supposes that we can utilize unlabeled data (about ten times more than labeled data) to help to learn the model; furthermore, (ii) transductive inference [24] for few-shot learning (TFSL) [21], [25] assumes we can access to all the test data, rather than evaluate them one by one in the inference process. In other words, the

---

- *Yanwei Fu is the corresponding author.*
- *Yikai Wang and Yanwei Fu are with the School of Data Science, Fudan University, and Shanghai Key Lab of Intelligent Information Processing, Fudan University. Yanwei Fu is also with the MOE Frontiers Center for Brain Science, Fudan University. E-mail: {yikaiwang19, yanweifu}@fudan.edu.cn*
- *Li Zhang is with the Department of Engineering Science, University of Oxford. E-mail: lz@robots.ox.ac.uk*
- *Yuan Yao is with the Department of Mathematics, Hong Kong University of Science and Technology. E-mail: yuany@ust.hk*
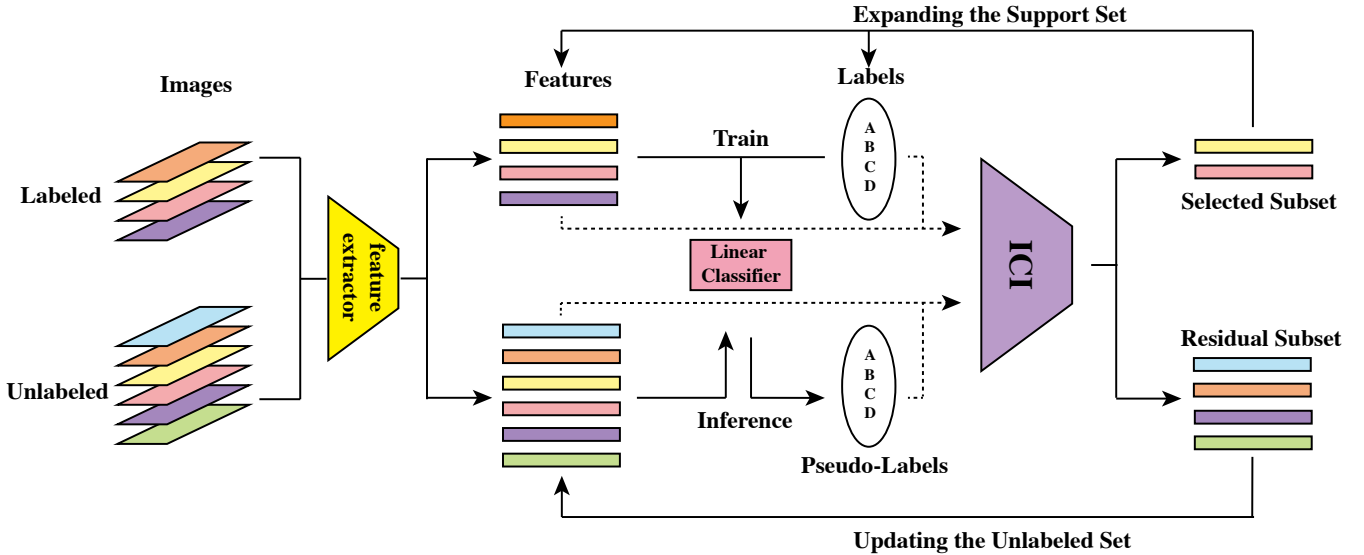
Fig. 1: The inference process of our proposed framework. We extract features of each labeled and unlabeled instance, train a linear classifier with the support set, provide pseudo-label for the unlabeled instances, and use ICI to select the most trustworthy worth subset to expand the support set. This process is repeated until all unlabeled data are included in the support set.

few-shot learning model can utilize the data distributions of testing examples.

Self-taught learning [26] is one of the most straightforward ways to leverage the information of unlabeled data. Typically, a trained classifier infers the pseudo labels of unlabeled data, which are further taken to update the classifier. Nevertheless, the inferred pseudo-labels may not be always trustworthy; the wrongly labeled instances may jeopardize the performance of the classifier. Specifically, in the few-shot scenario, both the labeled and unlabeled instances are limited. Thus the conventional assumptions utilized in semi-supervised learning methods, such as continuity, cluster, and manifold assumptions, are difficult to satisfy which eventually result in a noisy pseudo-labeled set. It is thus essential to investigate the labeling confidence of each unlabeled instance.

To this end, we present a statistical approach, dubbed Instance Credibility Inference (ICI) to exploit the distribution support of unlabeled instances for few-shot learning. Specifically, we first train a simple linear classifier (*e.g.* logistic regression, or linear support vector machine) with the labeled few-shot examples and use it to infer the pseudo-labels for the unlabeled instances. Then a subset of pseudo-labeled instances is selected according to their credibility measured by the proposed ICI to augment the support set. The simple classifier thus can be progressively updated (re-trained) by the expanded support set and further to infer pseudo-labels for the unlabeled data. This process is repeated until all the unlabeled instances are iteratively selected to expand the support set, *i.e.* the pseudo-label of each unlabeled instance is converged. The schematic illustration is shown in Figure 1.

Basically, we re-purpose the standard self-taught learning algorithm by our proposed ICI algorithm. How to select the pseudo-labeled data and exclude the wrongly-predicted samples, *i.e.*, excluding the noise introduced by

the self-taught learning strategy? Our intuition is that the credibility criteria can neither solely rely on the manifold structure of the feature space (*e.g.* instances that are close to labeled instances under a certain distance metric) nor the label space (*e.g.* probability provided by the classifier). Instead, we propose to solve the hypothesis of (generalized) linear models (*i.e.* linear regression or logistic regression) by progressively increasing the sparsity of the data-dependent incidental parameter [27] until it vanishes. Thus we can credit each pseudo-labeled instance by the sparsity of the corresponding incidental parameter. We prove that under the mild conditions of restricted eigenvalue, irrepresentability, and strong signal, our proposed method is able to collect *all* the correctly-predicted pseudo-labeled instances. We conduct extensive experiments on major few-shot learning benchmark datasets to validate the effectiveness of our proposed algorithm.

Contributions. The contributions of this work are as follows.

(i) We present a statistical approach, dubbed Instance Credibility Inference (ICI) to exploit the distribution support of unlabeled instances for few-shot learning. Specifically, our model iteratively selects the pseudo-labeled instances according to its credibility measured by the proposed ICI for classifier training.

(ii) We re-purpose the standard self-taught learning algorithm [26] by our proposed ICI. To measure the credibility of each pseudo-labeled instance, we solve the LM/GLM hypothesis by increasing the sparsity of the incidental parameter [27] and regard the sparsity level as the credibility for each pseudo-labeled instance.

(iii) Under the mild conditions of restricted eigenvalue, irrepresentability, and large error, we can prove that our method collects *all* the correctly-predicted pseudo-labeled instances.

(iv) Extensive experiments under two few-shot settings

show that our approach is able to establish new state-of-the-arts on four widely used few-shot learning benchmark datasets including *mini*ImageNet, *tiered*ImageNet, CIFAR-FS, and CUB.

Extensions. A preliminary version of this work was published in [28]. We have extended our conference version as follows.

(i) We provide the theoretical analysis (sufficient condition and necessary condition) of ICI to answer the question that *under what conditions can ICI find the correctly-predicted instances*?

(ii) We show that our ICI can be extended to generalized linear models, in particular, a *logistic regression model with sparse incidental parameters*. Particularly we show in our experiments the effectiveness of such a logistic regression model with sparsity regularization for ICI.

(iii) Inspired by [29], we introduce a variant to better handle the few-shot scenario in the extreme scenario (one-shot setting). By training the classifier based on the mean feature prototype instead of the instances, we could reduce the variance in the 1-shot case and achieve better performance.

(iv) We extend Instance Credibility Inference (ICI) to a more general form to better exploit the distribution support of unlabeled instances for few-shot learning. By providing each instance a prior weight, our model could potentially provide credibility case by case instead of under the assumption of equal chance to be right-predicted.

2 RELATED WORK

2.1 Semi-supervised learning

Semi-supervised learning (SSL) aims to improve the learning performance with both labeled and unlabeled instances. Basic assumptions in semi-supervised learning include continuity, cluster, and manifold assumptions. Conventional approaches focus on finding decision boundaries with both labeled and unlabeled data [24], [30], [31], and avoiding to learn the "wrong" knowledge from the unlabeled data [32] based on specific hypothesis. Recently, semi-supervised learning with deep learning models use consistency regularization [33], moving average technique [34] and adversarial perturbation regularization [35] to train the model with a large amount of unlabeled data. Though seems similar, semi-supervised learning and semi-supervised few-shot learning aim to solve different problems. The major difference is that the unlabeled data in the few-shot scenario is still limited, making the classical assumptions harder to achieve in the few-shot scenario. Thus semi-supervised FSL is a more difficult problem. Additionally, traditional semi-supervised learning methods usually suffer the problem of overfitting which indicates that SSL methods cannot be naively utilized to solve the few-shot learning problem.

2.2 Self-taught learning

Self-taught learning [26], also known as self-training [36], is a traditional semi-supervised strategy of utilizing unlabeled data to improve the performance of classifiers [37], [38]. Self-taught learning algorithms often start by training an initial recognition model and infer the pseudo-labels of unlabeled instances, then the pseudo-labeled instances are taken to re-train the recognition model with specific strategies [39]. Deep learning based self-taught learning strategy includes (i) directly training the neural network with both labeled instances and pseudo-labeled instances [39], (ii) utilizing mix-up images between labeled instances and pseudo-labeled instances to synthesis training instances with less noise [40], (iii) utilizing indirect ways to infer the pseudo-label of unlabeled instances (for example use label propagation constructed on the nearest-neighbor graph and select the trustworthy subset based on the entropy [41]), and (iv) methods that introducing inductive bias (*e.g.* adding a cluster assumption on the feature space and re-weight the pseudo-labeled instances based on this assumption [42]) One of the key points in self-taught learning algorithms is how to reduce the noise introduced by the imperfect recognition models. Different from previous works, we measure the credibility of each pseudo-labeled instance by a statistical algorithm. Only the most trustworthy subset is employed to re-train the recognition model jointly with the labeled instances.

2.3 Few-shot learning

Few-shot learning aims to recognize novel visual categories from very few labeled examples. Recent efforts are mainly following the meta-learning strategy. That is, by simulating the few-shot scenario in the training process, algorithms will learning to learn with limited data. Generally, existing works on few-shot learning can be roughly categorized into three groups. (i) The first group of method focus on learning robust and discriminative distance metrics, including weighted nearest neighbor classifier (*e.g.* Matching Network [9]), finding robust prototype for each class (*e.g.* Prototypical Network [7]), learning specific metric for each task (*e.g.* TADAM [14]), and learning a metric via neural networks [8]. (ii) The second group of methods aim to find the optimal initialization parameters that could rapidly adapt to specific task, including Meta-Critic [15], MAML [6], Meta-SGD [16], Reptile [17], and LEO [18]. (iii) Recently, the third group of methods start to appear which are based on data augmentation strategies. It aims to synthesize the informative instances to alleviate the limited-data problem in the image level [11] or the feature level [12]. Additionally, SNAIL [43] utilizes the sequence modeling to create a new framework. The proposed statistical algorithm is orthogonal and potentially beneficial to these algorithms – it is always worth increasing the training set by utilizing the unlabeled data with confidently predicted labels.

2.4 Few-shot learning with unlabeled data

Recent work starts to tackle few-shot learning with additional unlabeled instances. Compared with the traditional inductive setting, algorithms trained with unlabeled instances have the chance to handle a more trustworthy empirical distribution. Specifically, some recent efforts [21], [22] enable unlabeled data coming from the same data pool and train the algorithm in the semi-supervised few-shot learning settings. Furthermore, transductive inference is also introduced to alleviate the limited-data problem. For example, LST [23] uses the self-taught learning strategy in the transductive inference setting and trains the model in a

meta-learning manner. Different from these methods, this paper presents a conceptually simple statistical approach derived from self-taught learning. Unlike previous meta-learning algorithms which usually has pre-training, meta-training, and meta-test process [23], our approach only modifies the inference process and does not have to meta-train the architecture. Besides, our approach empirically and significantly improves the performance of FSL on several benchmark datasets, by only using very simple classifiers, *e.g.*, logistic regression, or Support Vector Machine (SVM).

2.5 Incidental parameters

Incidental parameters problem [44] was tackled by the penalized estimation algorithms [45]. It assumes the existence of sparse data-dependent parameters in the estimation models. For example, the linear regression model with incidental parameters follows $y_i = x_i^\top \beta + \gamma_i + \varepsilon_i$, where (x_i, y_i) denotes data input, β is the traditional coefficients, ε_i denotes the random noise and γ_i is the introduced data-dependent incidental parameters. Prior arts solve this problem by estimating the coefficients which are robust against the incidental parameters [27], [44], [46], [47], [48]. Fu *et al.* [49] introduce the incidental parameter in robust ranking task. In this paper, we propose to solve the few-shot learning problem based on the intuition that the incidental parameters indicate the credibility of pseudo-labeled instances. We do so by utilizing a weak estimation of coefficients to enlarge the influence of incidental parameters and transfer a "*classification model with incidental parameters*" into a normal "*classification*" model whose coefficients are the former incidental parameters. Then we estimate the incidental parameters along the regularization path to get the credibility of the corresponding instance. We further provide the theoretical properties of ICI, which, to the best of our knowledge, has not been studied in the previous work.

3 METHODOLOGY

3.1 Problem formulation

Here we define the few-shot learning problem mathematically. We are provided a base category set and a novel category set, denoted as $\mathcal{C}_{base}$ and $\mathcal{C}_{novel}$, respectively. The two category sets have no common category[1], *i.e.*, $\mathcal{C}_{base} \bigcap \mathcal{C}_{novel} = \emptyset$. Within each category set, we have a corresponding dataset, denoted as $\mathcal{D}_{base} = \{(\boldsymbol{I}_i, y_i), y_i \in \mathcal{C}_{base}\}$ and $\mathcal{D}_{novel} = \{(\boldsymbol{I}_i, y_i), y_i \in \mathcal{C}_{novel}\}$, respectively. With the above notations, few-shot learning algorithms aim to train on $\mathcal{D}_{base}$ and contain the capacity of rapidly adapt to $\mathcal{D}_{novel}$ with access to only one or a few labeled instances per class.

For evaluation, we adopt the standard *c-way-m-shot* classification as [9] on $\mathcal{D}_{novel}$. Specifically, in each episode, we randomly sample c classes to construct our category pool $\mathcal{C}$, that is $\mathcal{C} \sim \mathcal{C}_{novel}, |\mathcal{C}| = c$; and s and q labeled images per class are randomly sampled in $\mathcal{C}$ to construct the support set $\mathcal{S}$ and the query set $\mathcal{Q}$, respectively. Thus we have $|\mathcal{S}| = c \times s$

---

1. Note that here and below we ignore another validation set for model selection since we could regard it as the novel set that is accessible in the training process.

---

and $|\mathcal{Q}| = c \times q$. The classification accuracy is averaged on query sets $\mathcal{Q}$ of many meta-testing episodes. In addition, we have unlabeled data of novel categories $\mathcal{U}_{novel} = \{\boldsymbol{I}_u\}$.

3.2 Self-taught learning from unlabeled data

We recap the self-taught learning formalism [26] to tackle few-shot learning problem with unlabeled data. Particularly, denote $f(\cdot)$ as the feature extractor trained on $\mathcal{D}_{base}$. In one episode, one can train a supervised classifier $g(\cdot)$ on the support set $\mathcal{S}$, and pseudo-labeling unlabeled data, $\hat{y}_i = g(f(\boldsymbol{I}_u))$ with corresponding confidence p_i. The most confident unlabeled instances will be further taken as additional data of corresponding classes in the support set $\mathcal{S}$. Thus we obtain the updated supervised classifier $g(\cdot)$. To this end, few-shot classifier acquires additional training instances, and thus its performance can be improved.

However, it is problematic if directly utilizing self-taught learning in few-shot cases. Particularly, the supervised classifier $g(\cdot)$ is only trained by a few instances. The unlabeled instances with high confidence may not be correctly categorized, and the classifier will be updated by some wrong instances. Even worse, one can not assume the unlabeled instances follows the same class labels or generative distribution as the labeled data. Noisy instances or outliers may also be utilized to update the classifiers. To this end, we propose a systematical algorithm: Instance Credibility Inference (ICI) to reduce the noise.

3.3 Instance credibility inference (ICI)

To measure the credibility of predicted labels over unlabeled data, we introduce a hypothesis of linear model by regressing each instance from feature to label spaces. Particularly, given n instances of c classes, $\mathcal{S} = \{(\boldsymbol{I}_i, y_i, \boldsymbol{x}_i), y_i \in \mathcal{C}_{novel}\}$, where y_i is the ground truth when $\boldsymbol{I}_i$ comes from the support set, or the pseudo-label when $\boldsymbol{I}_i$ comes from the unlabeled set, we employ a simple linear regression model to "predict" the class label,

$$\boldsymbol{y}_i = \boldsymbol{x}_i^\top \boldsymbol{\beta} + \boldsymbol{\gamma}_i + \boldsymbol{\varepsilon}_i, \tag{1}$$

where $\boldsymbol{\beta} \in \mathbb{R}^{d \times c}$ is the coefficient matrix; $\boldsymbol{x}_i \in \mathbb{R}^{d \times 1}$ is the feature vector of instance i; $\boldsymbol{y}_i$ is c dimension one-hot vector denoting the class label of instance i. Note that to facilitate the computations, we employ Locally Linear Embedding (LLE) [50] to reduce the dimension of extracted features $f(\boldsymbol{I}_i)$ to d. $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$ is the Gaussian noise of zero mean and σ variance. Inspired by incidental parameters [27], we introduce $\gamma_{i,j}$ to amend the chance of instance i belonging to class j. The larger $\|\gamma_{i,j}\|$, the higher difficulty in attributing instance i to class j.

Consider the linear regression model for all instances, we are solving the problem of

$$\underset{\boldsymbol{\beta},\boldsymbol{\gamma}}{\operatorname{argmin}} \sum_{i=1}^{n} w_i \left[\left\| \boldsymbol{y}_i - \boldsymbol{x}_i^\top \boldsymbol{\beta} + \boldsymbol{\gamma}_i \right\|_2^2 + \lambda R(\boldsymbol{\gamma}_i) \right], \tag{2}$$

where w_i indicates the importance of regression on instance i and $R(\cdot)$ is the sparsity penalty, *e.g.* $R(\boldsymbol{\gamma}) = \sum_{i=1}^{n} \sum_{j=1}^{c} |\gamma_{i,j}|$. By re-writing Eq. (2) in a matrix form, we are thus solving the problem of

$$(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) = \underset{\boldsymbol{\beta},\boldsymbol{\gamma}}{\operatorname{argmin}} \left\| \sqrt{\boldsymbol{W}} (\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{\beta} - \boldsymbol{\gamma}) \right\|_F^2 + \lambda R(\boldsymbol{\gamma}), \tag{3}$$

where $\boldsymbol{W}$ is the diagonal matrix, $\boldsymbol{W} = \mathrm{diag}\,(w_1,\ldots,w_n)$ and $\sqrt{\boldsymbol{W}} = \mathrm{diag}\,(\sqrt{w_1},\ldots,\sqrt{w_n})$. $\|\cdot\|_{\mathrm{F}}^2$ denotes the Frobenius norm. $\boldsymbol{Y} = [\boldsymbol{y}_i] \in \mathbb{R}^{n\times c}$ and $\boldsymbol{X} = [\boldsymbol{x}_i^\top] \in \mathbb{R}^{n\times d}$ indicate label and feature input respectively. $\boldsymbol{\gamma} = [\boldsymbol{\gamma}_i] \in \mathbb{R}^{n\times c}$ is the incidental matrix. λ is the coefficient of the penalty term $R\,(\cdot)$. To solve Eq. (3), we find the derivative with respect to $\boldsymbol{\beta}$ and make it equal to 0, then we have

$$\hat{\boldsymbol{\beta}} = \left(\boldsymbol{X}^\top \boldsymbol{W} \boldsymbol{X} \right)^\dagger \boldsymbol{X}^\top \boldsymbol{W} \left(\boldsymbol{Y} - \boldsymbol{\gamma} \right), \qquad (4)$$

where $(\cdot)^\dagger$ denotes the Moore-Penrose pseudo-inverse. Note that (1) we are interested in utilizing $\boldsymbol{\gamma}$ to measure the credibility of each instance along its regularization path, rather than estimating $\hat{\boldsymbol{\beta}}$, since the linear regression model is not good enough for classification in general; (2) the $\hat{\boldsymbol{\beta}}$ also relies on the estimation of $\boldsymbol{\gamma}$. To this end, we take Eq. (4) into Eq. (3) and solve the problem as

$$\operatorname*{argmin}_{\boldsymbol{\gamma}\in\mathbb{R}^{n\times c}} \left\| \sqrt{\boldsymbol{W}} \left(\boldsymbol{Y} - \boldsymbol{H} \left(\boldsymbol{Y} - \boldsymbol{\gamma} \right) - \boldsymbol{\gamma} \right) \right\|_{\mathrm{F}}^2 + \lambda R\,(\boldsymbol{\gamma}), \quad (5)$$

where $\boldsymbol{H} = \boldsymbol{X} \left(\boldsymbol{X}^\top \boldsymbol{W} \boldsymbol{X} \right)^\dagger \boldsymbol{X}^\top \boldsymbol{W}$. We further define $\tilde{\boldsymbol{X}} = \sqrt{\boldsymbol{W}} \left(\boldsymbol{I} - \boldsymbol{H} \right)$ and $\tilde{\boldsymbol{Y}} = \tilde{\boldsymbol{X}} \boldsymbol{Y}$. Then the above equation can be simplified as

$$\hat{\boldsymbol{\gamma}} = \operatorname*{argmin}_{\boldsymbol{\gamma}\in\mathbb{R}^{n\times c}} \left\| \tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{X}} \boldsymbol{\gamma} \right\|_{\mathrm{F}}^2 + \lambda R\,(\boldsymbol{\gamma}), \qquad (6)$$

which is a multi-response regression problem.

Particularly, we regard $\hat{\boldsymbol{\gamma}}$ as a function of λ. When λ changes from 0 to ∞, the sparsity of $\hat{\boldsymbol{\gamma}}$ is increased until all of its elements are forced to vanish. Further, we use the penalty $R\,(\boldsymbol{\gamma})$ to encourage $\boldsymbol{\gamma}$ vanishes row by row, *i.e.*, instance by instance. For example, $R\,(\boldsymbol{\gamma}) = \sum_{i=1}^n \sum_{j=1}^c |\gamma_{i,j}|$ or $R\,(\boldsymbol{\gamma}) = \sum_{i=1}^n \|\boldsymbol{\gamma}_i\|_2$ or the weighted variants. Moreover, the penalty will tend to vanish the subset of $\tilde{X}$ with the lowest deviations, indicating less discrepancy between the prediction and the ground truth. Hence we could rank the pseudo-labeled data by the *smallest* λ value when the corresponding $\hat{\gamma}_i$ vanishes. As shown in one toy example of Figure 2, the $\hat{\boldsymbol{\gamma}}$ value of the instance denoted by the red line vanishes first, and thus it is the most trustworthy sample by our algorithm.

We seek the best subset by checking the regularization path, *i.e.* $\hat{\boldsymbol{\gamma}}(\lambda)$ as λ varies, which can be easily configured by a block coordinate descent algorithm implemented in Glmnet [51]. Specifically, we have a theoretical value of $\lambda_{max} = \max_i \left\| \tilde{\boldsymbol{X}}_{\cdot i}^\top \tilde{\boldsymbol{Y}} \right\|_2 / n$ [51] to guarantee the solution of Eq. (6) all 0. Then we can get a list of λs from 0 to λ_{max}. We solve a specific Eq. (6) with each λ, and get the regularization path of $\boldsymbol{\gamma}$ along the way.

To further improve the robustness of ICI, in each iteration, we limit the number of candidate pseudo-labeled instances as the $2 \times a$, where a is the number of images we will select in this iteration and the pre-candidate list is provided by the most confidence instances ranked by the linear classifier used, *i.e.* linear support vector machine or logistic regression in our experiments

3.4 Extension to logistic regression

In the above section, we develop ICI with a linear model. But the basic idea of measuring credibility of pseudo-labeled
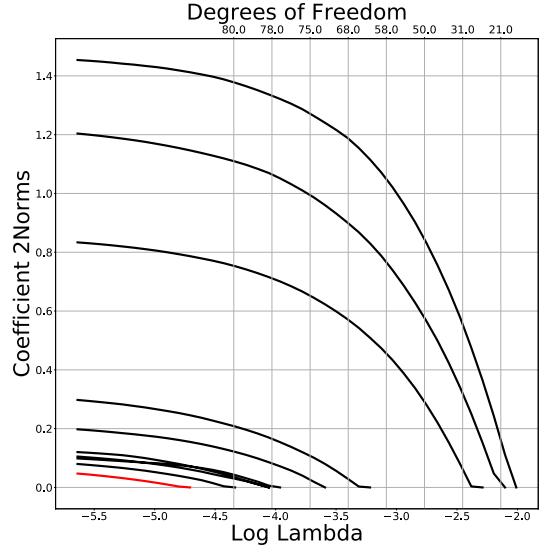

Fig. 2: Regularization path of λ on ten samples. Red line is corresponding to the most trustworthy sample suggested by our ICI algorithm.

Algorithm 1 Inference process of our algorithm.

Input: support data$\{(\boldsymbol{X}_i, \boldsymbol{y}_i)\}_{i=1}^{c\times s}$, query data $\boldsymbol{X}_t = \{\boldsymbol{X}_j\}_{j=1}^M$, unlabeled data $\boldsymbol{X}_u = \{\boldsymbol{X}_k\}_{k=1}^U$
Initialization: support set $(\boldsymbol{X}_s, \boldsymbol{Y}_s) = \{(\boldsymbol{X}_i, \boldsymbol{y}_i)\}_{i=1}^{c\times s}$, feature matrix $\boldsymbol{X}_{c\times s+U,d} = [\boldsymbol{X}_s; \boldsymbol{X}_u]$, classifier
Repeat:
Train classifier using $(\boldsymbol{X}_s, \boldsymbol{Y}_s)$;
Get pseudo-label $\boldsymbol{Y}_u$ for $\boldsymbol{X}_u$ by classifier;
Rank $(\boldsymbol{X}, \boldsymbol{Y}) = (\boldsymbol{X}, [\boldsymbol{Y}_s; \boldsymbol{Y}_u])$ by ICI;
Select a subset $(\boldsymbol{X}_{\mathrm{sub}}, \boldsymbol{Y}_{\mathrm{sub}})$ into $(\boldsymbol{X}_s, \boldsymbol{Y}_s)$;
Until Converged.
Inference:
Train classifier using $(\boldsymbol{X}_s, \boldsymbol{Y}_s)$;
Get pseudo-label $\boldsymbol{Y}_t$ for $\boldsymbol{X}_t$ by classifier;
Output: inference labels $\boldsymbol{Y}_t = \{\hat{\boldsymbol{y}}_j\}_{j=1}^M$

---

instance as the sparsity level of the corresponding incidental parameters along the regularization path is general and not limited in the linear model. To show this, in this section we extend ICI with generalized linear models, particularly, the logistic regression model.

Recall that we have $\boldsymbol{Y} = [\boldsymbol{y}_i] \in \mathbb{R}^{n\times c}$ and $\boldsymbol{X} = [\boldsymbol{x}_i^\top] \in \mathbb{R}^{n\times d}$ as our label matrix and feature matrix, respectively. We use $\boldsymbol{\beta} \in \mathbb{R}^{d\times c}$ as the coefficient matrix and $\boldsymbol{\gamma} = [\boldsymbol{\gamma}_i] \in \mathbb{R}^{n\times c}$ as the incidental matrix. Then our logistic model with incidental parameters can be formed as

$$\hat{\boldsymbol{Y}}_{i,c} = \frac{\exp\left(\boldsymbol{X}_{i\cdot} \boldsymbol{\beta}_{\cdot c} + \boldsymbol{\gamma}_{i,c} \right)}{\sum_{l=1}^C \exp\left(\boldsymbol{X}_{i\cdot} \boldsymbol{\beta}_{\cdot l} + \boldsymbol{\gamma}_{i,l} \right)}. \qquad (7)$$

This could be reformulated into a standard logistic regression model with sparsity regularization. Specifically, we define $\tilde{\boldsymbol{X}} = (\boldsymbol{X}, \boldsymbol{I}) \in \mathbb{R}^{n\times(d+n)}$ and $\tilde{\boldsymbol{\beta}} = (\boldsymbol{\beta}, \boldsymbol{\gamma})^\top \in \mathbb{R}^{(d+n)\times c}$, in which $\boldsymbol{I}$ is the identity matrix. Then we have

$$\tilde{\boldsymbol{X}}_{i\cdot} \tilde{\boldsymbol{\beta}}_{\cdot c} = (\boldsymbol{X}_{i\cdot}, \boldsymbol{I}_{i\cdot})(\boldsymbol{\beta}_{\cdot c}, \boldsymbol{\gamma}_{\cdot c})^\top = \boldsymbol{X}_{i\cdot} \boldsymbol{\beta}_{\cdot c} + \boldsymbol{\gamma}_{i,c}. \qquad (8)$$

Fig. 3: Selected images of each iteration in an inference episode on *mini*ImageNet. The averaged test accuracy is on the left, while the test accuracy of each class is put at the bottom of the corresponding images in each iteration. In each iteration, the right-predicted instances of each class are placed on the left, and vice versa on the right. For each class, we select 5 images at most. Some classes are expanded with smaller than that number because, in that iteration, the residual pseudo-labeled instances belonging to that class is smaller than 5. Then we simply select them all.

Hence we could reformulate Eq. (7) as

$$\hat{\boldsymbol{Y}}_{i,c} = \frac{\exp\left(\tilde{\boldsymbol{X}}_{i\cdot}\tilde{\boldsymbol{\beta}}_{\cdot c}\right)}{\sum_{l=1}^{C} \exp\left(\tilde{\boldsymbol{X}}_{i\cdot}\tilde{\boldsymbol{\beta}}_{\cdot l}\right)}, \tag{9}$$

which is exactly a logistic regression model. Our objective is the penalized negative log-likelihood function:

$$\operatorname*{argmin}_{\tilde{\boldsymbol{\beta}}=(\boldsymbol{\beta},\boldsymbol{\gamma})^{\top}} -\frac{1}{n}\sum_{i=1}^{n}\left(\sum_{l=1}^{c}\boldsymbol{Y}_{i,l}\left(\tilde{\boldsymbol{X}}_{i,\cdot}\tilde{\boldsymbol{\beta}}_{\cdot,l}\right) - \log\left(\sum_{l=1}^{c}e^{\tilde{\boldsymbol{X}}_{i,\cdot}\tilde{\boldsymbol{\beta}}_{\cdot,l}}\right)\right)$$
$$+ \lambda_1 R\left(\boldsymbol{\beta}\right) + \lambda_2 R\left(\boldsymbol{\gamma}\right). \tag{10}$$

Like in the linear regression scenario, we could extend the logistic regression model to a weighted model by re-weight each instance in the objective function as follows,

$$-\frac{1}{n}\sum_{i=1}^{n}w_i\left(\sum_{l=1}^{c}\boldsymbol{Y}_{i,l}\left(\tilde{\boldsymbol{X}}_{i,\cdot}\tilde{\boldsymbol{\beta}}_{\cdot,l}\right) - \log\left(\sum_{l=1}^{c}e^{\tilde{\boldsymbol{X}}_{i,\cdot}\tilde{\boldsymbol{\beta}}_{\cdot,l}}\right)\right)$$
$$+ \lambda_1 R\left(\boldsymbol{\beta}\right) + \lambda_2 R\left(\boldsymbol{\gamma}\right). \tag{11}$$

The algorithm for solving Eq. (10) and Eq. (11) is well established [51], [52], [53], [54]. Note that unlike the linear regression version where we can calculate a closed-form solution for $\boldsymbol{\beta}$, here the penalty of $\boldsymbol{\beta}$ is necessary or we will not achieve a unique solution, *i.e.* the solution is ill-posed [55]. For example, assume that we have a large enough λ_2 to vanish all elements of $\boldsymbol{\gamma}$. Then the problem degenerates to the normal logistic regression with the coefficient $\boldsymbol{\beta}$. Suppose we have an optimal solution $\boldsymbol{\beta}^*$, and we replace the k-th row $\boldsymbol{\beta}_{k,\cdot}^*$ by $\boldsymbol{\beta}_{k,\cdot}^* + \varepsilon \mathbf{1}^{\top}$ where ε is some scalar. Then we have

$$\hat{\boldsymbol{Y}}_{i,l}\mid_{\boldsymbol{\beta}_{k,\cdot}^* + \varepsilon\mathbf{1}^{\top}} = \frac{\exp\left(\tilde{\boldsymbol{X}}_{i\cdot}\boldsymbol{\beta}_{\cdot c}^* + x_{i,k}\varepsilon\right)}{\sum_{l=1}^{C}\exp\left(\tilde{\boldsymbol{X}}_{i\cdot}\boldsymbol{\beta}_{\cdot l}^* + x_{i,k}\varepsilon\right)}$$
$$= \frac{\exp\left(\tilde{\boldsymbol{X}}_{i\cdot}\boldsymbol{\beta}_{\cdot c}^*\right)}{\sum_{l=1}^{C}\exp\left(\tilde{\boldsymbol{X}}_{i\cdot}\boldsymbol{\beta}_{\cdot l}^*\right)} \tag{12}$$
$$= \hat{\boldsymbol{Y}}_{i,l}\mid_{\boldsymbol{\beta}_{k,\cdot}^*}.$$

Hence, to get a unique solution, we must provide some penalty on $\boldsymbol{\beta}$.

We use a partial Newton algorithm [51] to solve this optimization problem. Similar to the linear regression model, we use a list of λs to calculate the regularization path of $\boldsymbol{\gamma}$. The following maximal value of λ guarantees the solution $\hat{\boldsymbol{\gamma}}$ vanishes, $\lambda_{\max} = \max_i \left\|\tilde{\boldsymbol{X}}_{\cdot,i}^{\top}\left(\boldsymbol{Y} - \boldsymbol{P}_0\right)\right\|_2$ [51] where $\boldsymbol{P}_0$ is the sample proportions in each class.

3.5 Self-taught learning with ICI

The proposed ICI can thus be easily integrated to improve the self-taught learning algorithm. Particularly, the initialized classifier can predict the pseudo-labels of unlabeled instances; and we further employ the ICI algorithm to select the most confident subset of unlabeled instances, to update the classifier. The whole algorithm can be iteratively updated, as summarized in Algorithm 1. We also show a qualitative result in an inference episode in Fig. 3.

Intuitively, ICI focuses on fitting a line using the observations $(\boldsymbol{x}_i, \boldsymbol{y}_i)_{i=1}^{n}$ which contains outliers. Start from the labeled instances, we search the most possible inliers from the pseudo-labeled instances in each iteration. When we solve the line along the regularization path (from λ_{max} to λ_{min}), the estimated line will approach the more linear-separable subset, resulting in $\|\boldsymbol{\gamma}_i\| = 0$ for instances in this subset while $\|\boldsymbol{\gamma}_i\| > 0$ for others. Then we could use the linear-separable subset to improve the linear classifier. Furthermore, the fitted line cannot provide the right label for those outliers, hence the re-train process and re-infer process are essential to transfer outliers to inliers.

In practice, we propose two ways to utilize ICI: (1) Simply regard each instance as a data point for training the classifier, which is the most straightforward way and will be denoted with (c). (2) Inspired by [29], we find a prototype in the feature space for each class, and use these prototypes to train the classifier. These type of classification will be denoted with (p). When doing this, the selected pseudo-labeled data is utilized to update the prototype of

the corresponding class. Practically, we simply calculate the averaged feature of the few-shot samples (and the pseudo-labeled samples) as the prototype. Note that we still utilize the prototypes to train a classifier, rather than provide prediction directly based on the distance in the feature space.

4 IDENTIFIABILITY OF INSTANCE CREDIBILITY INFERENCE

In this part, we provide a theory for identifiability of ICI with linear regression model. Our theory is based on the model selection consistency for a linear regression with ℓ_1-sparsity regularization [56], [57]. Here our purpose is to answer the question of *under which conditions can we find the right-predicted instances?*

Recall that our intuition is that $\gamma_{i,j}$ can be regarded as the correction of the chance that instance i belonging to class j. Suppose γ^* is the ground truth. If the pseudo-labeled instance i is right-predicted, then we have $\gamma_{i,j}^* = 0, \forall j \in \{1, \ldots, c\}$. On the contrary, if the instance is wrongly predicted, then we should have $\gamma_{i,j}^* \neq 0$ for some j.

We start with reformulate the derivation process from Eq. (3) to Eq. (6) by another decoupled representation of solving $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$. For simplicity, we replace the notation of $\sqrt{W}Y, \sqrt{W}X, \sqrt{W}\gamma$ by Y, X, γ, respectively. Thus we are solving the problem of

$$\left(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}\right) = \operatorname*{argmin}_{\boldsymbol{\beta}, \boldsymbol{\gamma}} \|Y - X\boldsymbol{\beta} - \boldsymbol{\gamma}\|_F^2 + \lambda R\left(\boldsymbol{\gamma}\right). \quad (13)$$

Note that the above simplification will not change the support set of $\boldsymbol{\gamma}$, i.e.$\{(i,j) \mid \gamma_{i,j} \neq 0\}$ remains unchanged. We conduct the singular vector decomposition of X as $X = U\Sigma V^\top$, where $U \in \mathbb{R}^{n \times n}$, $\Sigma \in \mathbb{R}^{n \times d}$, $V \in \mathbb{R}^{d \times d}$. Recall that d is set as the reduced dimension from the original feature, hence we have $d \ll n$. Thus we could divide U into $U = [U_1, U_2]$ where U_1 is an orthogonal basis of the column space of X. Then we have $U^\top U = UU^\top = I$ and $U_2^\top X = 0$. Hence

$$
\begin{aligned}
&\|Y - X\boldsymbol{\beta} - \boldsymbol{\gamma}\|_F^2 \\
&= \left\|U^\top \left(Y - X\boldsymbol{\beta} - \boldsymbol{\gamma}\right)\right\|_F^2 \\
&= \left\|U_1^\top Y - U_1^\top X\boldsymbol{\beta} - U_1^\top \boldsymbol{\gamma}\right\|_F^2 + \left\|U_2^\top Y - U_2^\top \boldsymbol{\gamma}\right\|_F^2.
\end{aligned}
\quad (14)
$$

Again, we find the derivative with respect to $\boldsymbol{\beta}$ and make it equal to 0, then we have

$$\hat{\boldsymbol{\beta}} = \left(X^\top X\right)^\dagger X^\top \left(Y - \boldsymbol{\gamma}\right). \quad (15)$$

Take Eq. (15) into Eq. (14), we are now solving the problem of

$$L\left(\boldsymbol{\gamma}\right) = \left\|U_2^\top Y - U_2^\top \boldsymbol{\gamma}\right\|_F^2 + \lambda R\left(\boldsymbol{\gamma}\right). \quad (16)$$

Eq. (16) is equivalent to Eq. (6) but provides another interpretation that the incidental parameters (with a projection) try to find a sparse approximation of $U_2^\top Y$. Based on this, we could provide the answer of *under which condition could we recover the true support set of $\boldsymbol{\gamma}$?*

Formally, let $S = \operatorname{supp}(\boldsymbol{\gamma}^*)$ and $\hat{S} = \operatorname{supp}(\hat{\boldsymbol{\gamma}})$, where $\boldsymbol{\gamma}^*$ is the realistic prediction error, $\hat{\boldsymbol{\gamma}}$ is the estimator provided by our algorithm, and $\operatorname{supp}(X) = \{(i, \cdot) \mid X_{i,j} \neq$

0 for some $j\}$. For simplicity, we replace the notation of $U_2^\top$ and $U_2^\top Y$ by U and Y, respectively. Furthermore, denote U_S (U_{S^c}) as the column vectors of U whose index are in S (S^c), respectively. We are solving the problem of

$$\min_{\boldsymbol{\gamma}} \|Y - U\boldsymbol{\gamma}\|_F^2 + \lambda R\left(\boldsymbol{\gamma}\right), \quad (17)$$

where $Y \in \mathbb{R}^{(n-d) \times c}, U \in \mathbb{R}^{(n-d) \times n}, \boldsymbol{\gamma} \in \mathbb{R}^{n \times c}$. We consider the scenario $R\left(\boldsymbol{\gamma}\right) = \sum_{i=1}^n \sum_{j=1}^c |\gamma_{i,j}|$ for simplicity. Assume that $Y = U\boldsymbol{\gamma}^* + U\boldsymbol{\varepsilon} = U_S \boldsymbol{\gamma}_S^* + U\boldsymbol{\varepsilon}$, where $\varepsilon_{i,j}$ are independent Gaussian random variables with zero mean and standard deviation σ. Further let $\mu_U = \max_{(i,j) \in S^c} |U_{i,j}|$. We give three assumptions:
(C1: Restricted eigenvalue)

$$\lambda_{\min}\left(U_S^\top U_S\right) = C_{\min} > 0. \quad (18)$$

(C2: Irrepresentability) For some constant $\eta \in (0, 1]$,

$$\left\|U_{S^c}^\top U_S \left(U_S^\top U_S\right)^{-1} \operatorname{sign}\left(\boldsymbol{\gamma}_S^*\right)\right\|_{\max} \leq 1 - \eta. \quad (19)$$

(C3: Large error)

$$\boldsymbol{\gamma}_{\min} := \min_{(i,j) \in S} \left|\gamma_{i,j}^*\right| > h\left(\lambda, \eta, U, \boldsymbol{\gamma}^*\right), \quad (20)$$

where

$$h\left(\lambda, \eta, U, \boldsymbol{\gamma}^*\right) = \frac{\lambda\eta}{\sqrt{C_{\min}}\mu_U} + \lambda \left\|\left(U_S^\top U_S\right)^{-1} \operatorname{sign}\left(\boldsymbol{\gamma}_S^*\right)\right\|_{\max}. \quad (21)$$

Based on these conditions, we could provide the following theorems:

Theorem 1 (Sufficiency). *Let*

$$\lambda \geq \frac{2\sigma\sqrt{\mu_U}}{\eta}\sqrt{\log cn}.$$

Then with probability greater than

$$1 - 2cn \exp\left\{-\frac{\lambda^2 \eta^2}{2\sigma^2 \mu_U}\right\} \geq 1 - 2\left(cn\right)^{-1},$$

Eq. (17) has a unique solution $\hat{\boldsymbol{\gamma}}$ satisfies the following properties:

1) *If C1 and C2 hold, the wrong-predicted instances indicated by ICI has no false positive error, i.e. $\hat{S} \subseteq S$, and*

$$\|\hat{\boldsymbol{\gamma}}_S - \boldsymbol{\gamma}_S^*\|_{\max} \leq h\left(\lambda, \eta, U, \boldsymbol{\gamma}^*\right);$$

2) *If C1, C2, and C3 hold, ICI will identify all the correctly-predicted instances, i.e. $\hat{S} = S$ (in fact $\operatorname{sign}(\hat{\boldsymbol{\gamma}}) = \operatorname{sign}(\boldsymbol{\gamma}^*)$).*

Theorem 2 (Necessity). *Assume that C1 holds.*

1) *If C2 is violated, i.e.*

$$\left\|U_{S^c}^\top U_S \left(U_S^\top U_S\right)^{-1} \operatorname{sign}\left(\boldsymbol{\gamma}_S^*\right)\right\|_{\max} \geq 1,$$

then $\mathbb{P}\left(\hat{S} \subseteq S\right) < \frac{1}{2}$. In other words, with probability more than $1/2$ a solution $\hat{\boldsymbol{\gamma}}$ of Eq. (17) has false positive errors.

2) *If C3 is violated, then $\mathbb{P}\left(\operatorname{sign}(\hat{\boldsymbol{\gamma}}) = \operatorname{sign}(\boldsymbol{\gamma}^*)\right) < \frac{1}{2}$, or equivalently, with probability more than $1/2$ there exists some Gaussian noise such that a solution $\hat{\boldsymbol{\gamma}}$ of Eq. (17) has a false negative error.*

The proof is given in the Appendix. These two theorems show that our algorithm could find the right-predicted pseudo-labeled instances under specific conditions. Practically, it may be hard for us to choose a reasonable λ to satisfy the three conditions since we could not know γ_S^* in advance. Specifically, in the tasks of both semi-supervised and transductive few-shot learning concerned in this paper, one can not assume knowing the know γ_S^*. Hence, we use the iterative strategy to search along the solution path to select the instances automatically.

5 EXPERIMENTS

Datasets. Our experiments are conducted on several widely few-shot learning benchmark datasets for general object recognition and fine-grained classification, including *mini*ImageNet [58], *tiered*ImageNet [22], CIFAR-FS [59] and CUB [60]. *mini*ImageNet consists of 100 classes with 600 labeled instances in each category. We follow the split proposed by [58], using 64 classes as the base set to train the feature extractor, 16 classes as the validation set, and report performance on the novel set which consists of 20 classes. *tiered*ImageNet is a larger dataset compared with *mini*ImageNet, and its categories are selected with hierarchical structure to split base and novel datasets semantically. We follow the split introduced in [22] with base set of 20 superclasses (351 classes), validation set of 6 superclasses (97 classes) and novel set of 8 superclasses (160 classes). Each class contains 1281 images on average. **CUB** is a fine-grained dataset of 200 bird categories with 11788 images in total. Following the previous setting in [61], we use 100 classes as the base set, 50 for validation, and 50 as the novel set. To make a fair comparison, we crop all images with the bounding box provided by [62]. **CIFAR-FS** is a dataset with lower-resolution images derived from CIFAR-100 [63] . It contains 100 classes with 600 instances in each class. We follow the split given by [59], using 64 classes to construct the base set, 16 for validation, and 20 as the novel set.

Experimental setup. Unless otherwise specified, we use the following settings and implementation in the experiments for our approach to make a fair comparison. As in [14], [43], [64], we use ResNet-12 [65] with 4 residual blocks as the feature extractor in our experiments. Each block consists of three 3×3 convolutional layers, each of which followed by a BatchNorm layer and a LeakyReLu(0.1) activation. At the end of each block, a 2×2 max-pooling layer is utilized to reduce the output size. The number of filters in each block is 64, 128, 256 and 512 respectively. Specifically, referring to [64], we adopt the Dropout [66] in the first two block to vanish 10% of the output, and adopt DropBlock [67] in the latter two blocks to vanish 10% of output at channel level. Finally, an average-pooling layer is employed to produce the input feature embedding. We select 90% images from each training class (*e.g.*, 64 categories for *mini*ImageNet) to construct our training set for training the feature extractor and use the remaining 10% as the validation set to select the best model. We use SGD with momentum as the optimizer to train the feature extractor from scratch. Momentum factor and L_2 weight decay is set to 0.9 and $1e-4$, respectively. All inputs are resized to 84×84. We set the initial learning rate of 0.1, decayed

by 10 after every 30 epochs. The total training epochs is 120 epochs. In all of our experiments, we normalize the feature with L_2 norm and reduce the feature dimension to $d=5$ using LLE [50]. Our model and all baselines are evaluated over 600 episodes with 15 test samples from each class. Practically, we assume the uniform weight for the instances in Eq. 3.

5.1 Semi-supervised few-shot learning

Settings. In the inference process, the unlabeled data from the corresponding category pool is utilized to help FSL. In our experiments, we report the following settings of SSFSL: (1) we use 15 unlabeled samples for each class, the same as TFSL, to compare our algorithm in SSFSL and TFSL settings; (2) we use 30 unlabeled samples in 1-shot task and 50 unlabeled samples in 5-shot task, the same as current SSFSL approaches [23]; We denote these as 15/15 and 30/50 in Table 1. Note that CUB is a fine-grained dataset and does not have so sufficient samples in each class, so we simply choose 5 as support set, 15 as query set and other samples as unlabeled set (about 39 samples on average) in the 5-shot task in the latter setting. For all settings, we select 5 samples for every class in each iteration. The process is finished when at most 15/15, 25/45 unlabeled instances are selected in total, respectively. Further, we utilize Logistic Regression (denoted as *LR*) and linear Support Vector Machine (denoted as *SVM*) to show the robustness of ICI against different linear classifiers.

Competitors. We compare our algorithm with current approaches in SSFSL. TPN [21] uses labeled support set and unlabeled set to propagate label to one query sample each time. LST [23] also uses self-taught learning strategy to pseudo-label data and select confident ones, but they do this by a neural network trained in the meta-learning manner for many iterations. Other approaches include Masked Soft k-Means [22] and a combination of MTL with TPN and Masked Soft k-Means reported by LST.

Results. The results are shown in Table 1 where denoted as Semi. in the first column. Analyzing from the experimental results, we can find that: (1) Compare SSFSL with TFSL with the same number of unlabeled data, we can see that our SSFSL results are only reduced by a little or even beat TFSL results, which indicates that the information we got from the unlabeled data are robust and we can indeed handle the true distribution with unlabeled data practically. (2) The more unlabeled data we get, the better performance we have. Thus we can learn more knowledge with more unlabeled data almost consistently using a linear classifier (*e.g.*logistic regression). When lots of unlabeled data are accessible, ICI achieves state-of-the-art in all experiments even compared with competitors which use bigger network and higher-resolution inputs. (3) Compared with other SSFSL approaches, ICI also achieves varying degrees of improvements in almost all tasks and datasets. These results further indicate the robustness of our algorithm. Compare logistic regression with SVM, the robustness of ICI still holds.

5.2 Transductive few-shot learning

Settings. In transductive few-shot learning setting, we have the chance to access the query data in the inference

Setting	Model	*mini*ImageNet		*tiered*ImageNet		CIFAR-FS		CUB	
		1shot	5shot	1shot	5shot	1shot	5shot	1shot	5shot
In.	Baseline* [19]	$51.75_{\pm0.80}$	$74.27_{\pm0.63}$	-	-	-	-	$65.51_{\pm0.87}$	$82.85_{\pm0.55}$
	Baseline++* [19]	$51.87_{\pm0.77}$	$75.68_{\pm0.63}$	-	-	-	-	$67.02_{\pm0.90}$	$83.58_{\pm0.54}$
	MatchingNet* [9]	$52.91^{1}_{\pm0.88}$	$68.88^{1}_{\pm0.69}$	-	-	-	-	$72.36^{1}_{\pm0.90}$	$83.64^{1}_{\pm0.60}$
	ProtoNet* [7]	$54.16^{1}_{\pm0.82}$	$73.68^{1}_{\pm0.65}$	-	-	72.20^{3}	83.50^{3}	$71.88^{1}_{\pm0.91}$	$87.42^{1}_{\pm0.48}$
	MAML* [6]	$49.61^{1}_{\pm0.92}$	$65.72^{1}_{\pm0.77}$	-	-	-	-	$69.96^{1}_{\pm1.01}$	$82.70^{1}_{\pm0.65}$
	RelationNet* [8]	$52.48^{1}_{\pm0.86}$	$69.83^{1}_{\pm0.68}$	-	-	-	-	$67.59^{1}_{\pm1.02}$	$82.75^{1}_{\pm0.58}$
	adaResNet [68]	56.88	71.94	-	-	-	-	-	-
	TapNet [69]	61.65	76.36	63.08	80.26	-	-	-	-
	CTM† [70]	64.12	80.51	68.41	84.28	-	-	-	-
	MetaOptNet [64]	64.09	80.00	65.81	81.75	72.60	84.30	-	-
Tran.	TPN [21]	59.46	75.65	58.68^{4}	74.26^{4}	65.89^{4}	79.38^{4}	-	-
	TEAM* [25]	60.07	75.90	-	-	70.43	81.25	80.16	87.17
Semi.	MSkM + MTL	62.10^{2}	73.60^{2}	68.6^{2}	81.00^{2}	-	-	-	-
	TPN + MTL	62.70^{2}	74.20^{2}	72.10^{2}	83.30^{2}	-	-	-	-
	MSkM [22]	50.40	64.40	52.40	69.90	-	-	-	-
	TPN [21]	52.78	66.42	55.70	71.00	-	-	-	-
	LST [23]	70.10	78.70	77.70	85.20	-	-	-	-
In.	SVM(c)	$54.46_{\pm0.775}$	$74.62_{\pm0.650}$	$67.51_{\pm0.896}$	$84.67_{\pm0.618}$	$60.94_{\pm0.877}$	$80.32_{\pm0.678}$	$75.84_{\pm0.764}$	$89.36_{\pm0.451}$
	SVM(p)	$54.46_{\pm0.775}$	$75.06_{\pm0.649}$	$67.51_{\pm0.896}$	$84.92_{\pm0.637}$	$60.94_{\pm0.877}$	$80.55_{\pm0.685}$	$75.84_{\pm0.764}$	$90.06_{\pm0.457}$
	LR(c)	$56.06_{\pm0.773}$	$75.43_{\pm0.639}$	$69.02_{\pm0.877}$	$85.37_{\pm0.612}$	$62.25_{\pm0.871}$	$81.01_{\pm0.676}$	$76.16_{\pm0.750}$	$90.44_{\pm0.437}$
	LR(p)	$56.06_{\pm0.773}$	$75.41_{\pm0.634}$	$69.02_{\pm0.877}$	$85.32_{\pm0.616}$	$62.25_{\pm0.871}$	$80.79_{\pm0.668}$	$76.16_{\pm0.750}$	$89.95_{\pm0.461}$
Tran.	SVM + ICIR(c)	$62.53_{\pm1.059}$	$78.87_{\pm0.683}$	$76.90_{\pm1.095}$	$88.22_{\pm0.619}$	$70.03_{\pm1.139}$	$83.98_{\pm0.726}$	$85.39_{\pm0.846}$	$92.48_{\pm0.404}$
	SVM + ICIR(p)	$63.39_{\pm1.083}$	$78.00_{\pm0.698}$	$77.41_{\pm1.134}$	$86.92_{\pm0.651}$	$71.05_{\pm1.191}$	$82.59_{\pm0.748}$	$87.08_{\pm0.885}$	$92.18_{\pm0.420}$
	SVM + ICIC(c)	$61.72_{\pm1.055}$	$79.03_{\pm0.693}$	$76.32_{\pm1.086}$	$88.12_{\pm0.614}$	$69.10_{\pm1.120}$	$84.25_{\pm0.711}$	$84.64_{\pm0.867}$	$92.41_{\pm0.385}$
	SVM + ICIC(p)	$62.41_{\pm1.096}$	$78.00_{\pm0.686}$	$77.01_{\pm1.157}$	$86.97_{\pm0.655}$	$70.37_{\pm1.176}$	$82.75_{\pm0.737}$	$86.44_{\pm0.908}$	$92.16_{\pm0.414}$
	LR + ICIR(c)	$68.01_{\pm1.101}$	$79.66_{\pm0.669}$	$80.97_{\pm1.061}$	$88.48_{\pm0.597}$	$74.94_{\pm1.152}$	$84.37_{\pm0.727}$	$88.88_{\pm0.728}$	$92.96_{\pm0.387}$
	LR + ICIR(p)	$68.70_{\pm1.104}$	$79.07_{\pm0.670}$	$81.51_{\pm1.069}$	$87.89_{\pm0.618}$	$75.46_{\pm1.155}$	$83.38_{\pm0.736}$	$89.01_{\pm0.747}$	$92.22_{\pm0.408}$
	LR + ICIC(c)	$67.05_{\pm1.088}$	$79.69_{\pm0.657}$	$80.58_{\pm1.044}$	$88.46_{\pm0.602}$	$74.11_{\pm1.118}$	$84.50_{\pm0.718}$	$88.07_{\pm0.752}$	$92.84_{\pm0.382}$
	LR + ICIC(p)	$68.04_{\pm1.089}$	$78.96_{\pm0.655}$	$81.16_{\pm1.050}$	$87.76_{\pm0.618}$	$74.80_{\pm1.137}$	$83.32_{\pm0.718}$	$88.39_{\pm0.765}$	$92.21_{\pm0.409}$
Semi. 15/15	SVM + ICIR(c)	$61.58_{\pm1.022}$	$77.02_{\pm0.692}$	$76.84_{\pm1.145}$	$87.37_{\pm0.636}$	$69.88_{\pm1.148}$	$82.66_{\pm0.709}$	$84.94_{\pm0.838}$	$91.51_{\pm0.433}$
	SVM + ICIR(p)	$62.49_{\pm1.078}$	$77.23_{\pm0.696}$	$77.36_{\pm1.201}$	$86.97_{\pm0.645}$	$71.11_{\pm1.152}$	$82.40_{\pm0.738}$	$86.53_{\pm0.897}$	$92.13_{\pm0.423}$
	SVM + ICIC(c)	$60.74_{\pm1.014}$	$77.02_{\pm0.689}$	$76.06_{\pm1.132}$	$87.30_{\pm0.640}$	$69.11_{\pm1.156}$	$82.72_{\pm0.704}$	$84.42_{\pm0.857}$	$91.54_{\pm0.421}$
	SVM + ICIC(p)	$61.45_{\pm1.094}$	$77.28_{\pm0.701}$	$76.64_{\pm1.189}$	$87.00_{\pm0.641}$	$70.46_{\pm1.188}$	$82.32_{\pm0.766}$	$86.20_{\pm0.902}$	$92.22_{\pm0.420}$
	LR + ICIR(c)	$66.75_{\pm1.081}$	$78.16_{\pm0.683}$	$81.21_{\pm1.087}$	$88.22_{\pm0.617}$	$74.38_{\pm1.126}$	$83.58_{\pm0.721}$	$88.16_{\pm0.763}$	$92.47_{\pm0.417}$
	LR + ICIR(p)	$68.07_{\pm1.100}$	$77.93_{\pm0.679}$	$81.79_{\pm1.098}$	$87.64_{\pm0.636}$	$75.52_{\pm1.126}$	$82.68_{\pm0.734}$	$88.22_{\pm0.797}$	$91.80_{\pm0.434}$
	LR + ICIC(c)	$66.10_{\pm1.069}$	$78.25_{\pm0.681}$	$80.50_{\pm1.115}$	$88.27_{\pm0.602}$	$73.45_{\pm1.122}$	$83.67_{\pm0.711}$	$87.53_{\pm0.814}$	$92.47_{\pm0.407}$
	LR + ICIC(p)	$67.43_{\pm1.100}$	$77.85_{\pm0.677}$	$81.31_{\pm1.100}$	$87.59_{\pm0.639}$	$74.82_{\pm1.118}$	$82.83_{\pm0.716}$	$88.00_{\pm0.802}$	$91.74_{\pm0.433}$
Semi. 30/50	SVM + ICIR(c)	$65.24_{\pm1.102}$	$79.76_{\pm0.679}$	$78.91_{\pm1.135}$	$89.29_{\pm0.582}$	$72.39_{\pm1.190}$	$85.53_{\pm0.688}$	$86.73_{\pm0.815}$	$92.45_{\pm0.440}$
	SVM + ICIR(p)	$63.73_{\pm1.140}$	$78.64_{\pm0.691}$	$78.60_{\pm1.184}$	$87.67_{\pm0.638}$	$71.96_{\pm1.234}$	$83.07_{\pm0.753}$	$87.21_{\pm0.923}$	$92.41_{\pm0.459}$
	SVM + ICIC(c)	$64.01_{\pm1.065}$	$80.05_{\pm0.662}$	$78.17_{\pm1.151}$	$89.16_{\pm0.589}$	$71.31_{\pm1.192}$	$85.33_{\pm0.694}$	$86.34_{\pm0.813}$	$92.52_{\pm0.432}$
	SVM + ICIC(p)	$62.67_{\pm1.152}$	$78.57_{\pm0.697}$	$77.68_{\pm1.216}$	$87.51_{\pm0.658}$	$70.72_{\pm1.283}$	$83.11_{\pm0.756}$	$86.86_{\pm0.956}$	$92.43_{\pm0.448}$
	LR + ICIR(c)	$71.07_{\pm1.136}$	$81.05_{\pm0.653}$	$\mathbf{84.44}_{\pm0.995}$	$\mathbf{89.99}_{\pm0.546}$	$\mathbf{78.02}_{\pm1.140}$	$85.96_{\pm0.689}$	$\mathbf{90.72}_{\pm0.682}$	$\mathbf{93.34}_{\pm0.423}$
	LR + ICIR(p)	$\mathbf{71.11}_{\pm1.151}$	$79.88_{\pm0.661}$	$84.35_{\pm1.004}$	$88.52_{\pm0.617}$	$77.85_{\pm1.122}$	$83.79_{\pm0.749}$	$90.15_{\pm0.707}$	$92.31_{\pm0.447}$
	LR + ICIC(c)	$70.29_{\pm1.114}$	$\mathbf{81.25}_{\pm0.652}$	$84.14_{\pm1.001}$	$89.89_{\pm0.554}$	$77.68_{\pm1.135}$	$\mathbf{85.97}_{\pm0.709}$	$90.42_{\pm0.684}$	$93.32_{\pm0.404}$
	LR + ICIC(p)	$70.54_{\pm1.151}$	$79.84_{\pm0.674}$	$83.70_{\pm1.029}$	$88.57_{\pm0.615}$	$77.36_{\pm1.130}$	$83.72_{\pm0.751}$	$89.92_{\pm0.702}$	$92.13_{\pm0.456}$

TABLE 1: The averaged accuracies with 95% confidence intervals over 600 episodes on several datasets. Results with $(\cdot)^{1}$ are reported in [19], with $(\cdot)^{2}$ are reported in [23], with $(\cdot)^{3}$ are reported in [64]. $(\cdot)^{4}$ is our implementation with the official code of [21]. Methods denoted by $(\cdot)^{*}$ denotes ResNet-18 with input size 224×224, while $(\cdot)^{\dagger}$ denotes ResNet-18 with input size 84×84. Our method and other alternatives use ResNet-12 with input size 84×84. **In.** and **Tran.** indicate inductive and transductive setting, respectively. **Semi.** denotes semi-supervised setting where $(\cdot/\cdot)$ shows the number of unlabeled data available in 1-shot and 5-shot experiments. ICIC indicates the logistic regression version of our model, and ICIR indicates the linear regression version. For the classification strategy, (c) denotes the normal classification strategy while (p) denotes averaging all the instances to get the prototype of each class and train the classifier using the prototypes.

stage. Thus the unlabeled set and the query dataset are the same. In our experiments, we select 5 instances for each class in each iteration and repeat our algorithm until all the query samples are included. We also utilize both Logistic Regression and SVM as our classifier, respectively.

Competitors. We compare ICI with current TFSL approaches. TPN [21] constructs a graph and uses label propagation to transfer labels from support samples to query samples and learn their framework in a meta-learning way.

TEAM [25] utilizes class prototypes with a data-dependent metric to inference labels of query samples.

Results. The results are shown in Table 1 where denoted as Tran. in the first column. Experiments cross four benchmark datasets indicate that: (1) Compared with the basic linear classifiers, ICI enjoys consistently improvements, especially in the 1-shot setting where the labeled data is extremely limited and such improvements are robust regardless of utilizing which linear classifiers. Further, compare results

between *mini*ImageNet and *tiered*ImageNet, we can find that the improvement margin is in a similar scale, indicating that the improvement of ICI does not rely on the semantic relationship between base set and novel set. Hence the effectiveness and robustness of ICI are confirmed practically. (2) Compared with current TFSL approaches, ICI also achieves state-of-the-art results.
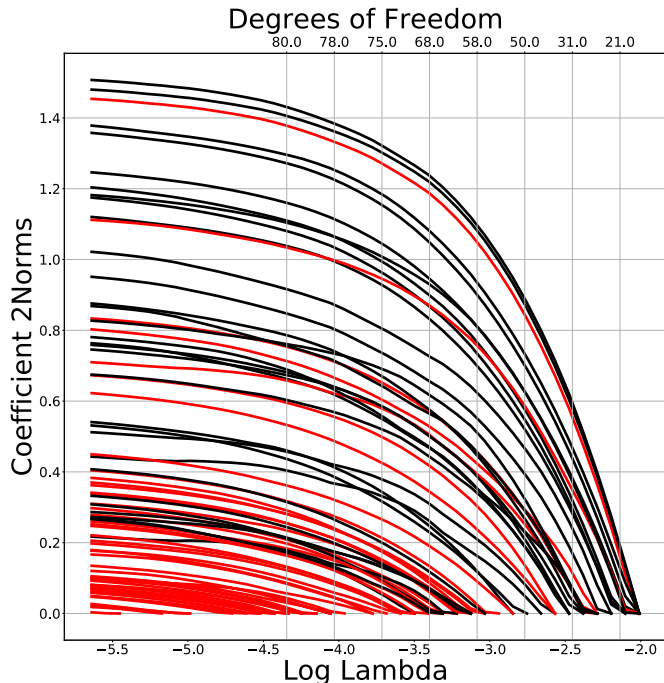
5.3 Ablation study

Fig. 4: Regularization path of λ. Red lines are correct-predicted instances while black lines are wrong-predicted ones. ICI will choose instances in the lower-left subset.

Effectiveness of ICI. To show the effectiveness of ICI, we visualize the regularization path of γ in one episode of the inference process in Figure 4 where red lines are instances that are correct-predicted while black lines are wrong-predicted ones. It is obvious that that most of the correct-predicted instances lie in the lower-left part. Since ICI will select samples whose norm will vanish in a lower λ. We could get more correct-predicted instances than wrong-predicted instances in a high ratio.

Compare to baselines. To further show the effectiveness of ICI, we compare ICI with other sample selection strategies under the self-taught learning pipeline. Here our ICI will not use the pre-candidate list provided by the linear classifier to make a fair comparison. One simple strategy is randomly sampling the unlabeled data into the expanded support set in each iteration, denoted as *ra*. Another is selecting the data based on the confidence given by the classifier, denoted by *co*. In this strategy, the more confident the classifier is to one sample, the more trustworthy that sample is. The last one is replacing our algorithm of computing credibility by choosing the nearest-neighbor of each class in the feature space, denoted as *nn*. In this part, we have

15 unlabeled instances for each class and select 5 to re-train the classifier by different methods for Semi. and Tran. task on *mini*ImageNet. From Table 2, we observe that ICI outperforms all the baselines in all settings.

Model	Tran.		Semi.	
	1shot	5shot	1shot	5shot
LR	$56.06_{\pm0.773}$	$75.43_{\pm0.639}$	$56.06_{\pm0.773}$	$75.43_{\pm0.639}$
+ ra	$59.01_{\pm0.885}$	$76.38_{\pm0.662}$	$59.46_{\pm0.916}$	$76.58_{\pm0.654}$
+ nn	$63.24_{\pm0.946}$	$77.63_{\pm0.636}$	$63.10_{\pm0.928}$	$77.75_{\pm0.664}$
+ co	$63.29_{\pm0.930}$	$77.92_{\pm0.640}$	$63.57_{\pm0.938}$	$77.71_{\pm0.626}$
ICI	$\mathbf{65.32}_{\pm1.010}$	$\mathbf{78.30}_{\pm0.658}$	$\mathbf{64.60}_{\pm0.999}$	$\mathbf{77.96}_{\pm0.653}$

TABLE 2: Compare to baselines on *mini*ImageNet under several settings.

Model	Tran.		Semi.	
	1shot	5shot	1shot	5shot
C(c) w/o	$66.19_{\pm1.079}$	$79.44_{\pm0.663}$	$64.87_{\pm1.075}$	$78.04_{\pm0.701}$
C(c) w/	$67.05_{\pm1.088}$	$79.69_{\pm0.657}$	$66.10_{\pm1.069}$	$78.25_{\pm0.681}$
C(p) w/o	$67.28_{\pm1.081}$	$78.71_{\pm0.665}$	$66.38_{\pm1.097}$	$77.71_{\pm0.698}$
C(p) w/	$68.04_{\pm1.089}$	$78.96_{\pm0.655}$	$67.43_{\pm1.100}$	$77.85_{\pm0.677}$
R(c) w/o	$67.58_{\pm1.119}$	$79.44_{\pm0.682}$	$66.45_{\pm1.112}$	$77.84_{\pm0.698}$
R(c) w/	$68.01_{\pm1.101}$	$79.66_{\pm0.669}$	$66.75_{\pm1.081}$	$78.16_{\pm0.683}$
R(p) w/o	$68.39_{\pm1.115}$	$78.89_{\pm0.674}$	$67.80_{\pm1.118}$	$77.71_{\pm0.694}$
R(p) w/	$68.70_{\pm1.104}$	$79.07_{\pm0.670}$	$68.07_{\pm1.100}$	$77.93_{\pm0.679}$

TABLE 3: Compare to baselines on *mini*ImageNet under several settings.

Effectiveness of the pre-candidate list. In our experiments, we utilize the linear classifier to provide a pre-candidate list to exclude the worst instances. To see the effectiveness, we run experiments on different settings and shots with ICI using or not using the pre-candidate list. Results are shown in Table 3, indicating the robust improvement of using the pre-candidate list.
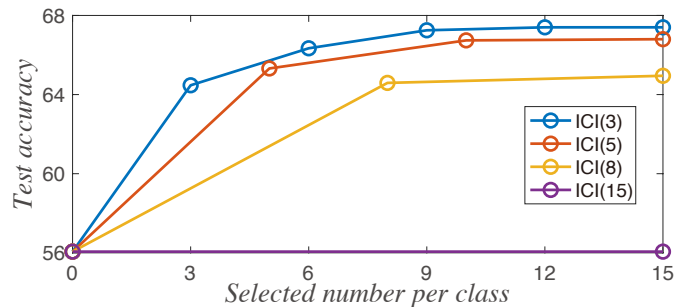

Fig. 5: Variation of accuracy as the selected samples increases over 600 episodes on *mini*ImageNet. "ICI (*n*)": select *n* samples per class in each iteration.

Effectiveness of iterative manner. Our intuition is the proposed ICI learns to generate a set of trustworthy unlabelled data for classifier training. Select all the unlabelled data in *one go* cannot take the distribution, or the credibility of the unlabeled data into account, and thus produce more noise labels to hurt the performance of the model. The classifier thus is trained with its prediction, resulting in no improvements in TFSL setting. We briefly validate this as ICI (15) in Figure 5 whilst ICI obtained better accuracy with

iterative selection manner. For example, select 6 images with two iterations (ICI(3)) is superior to select 8 images in one iteration (ICI(8)).

Acc (%)	0-10	10-20	20-30	30-40	40-50
b/t	0/0	0/0	1/3	16/23	105/125

Acc (%)	50-60	60-70	70-80	80-90	90-100
b/t	193/218	171/189	34/40	2/2	0/0

TABLE 4: We run 600 episodes, with each episode training an initial classifier. We denote "Acc" as the accuracy intervals; and "b/T" as the number of classifiers experienced improvement v.s. total classifiers in this accuracy interval.

Robustness against initial classifier. What are the requirements for the initial linear classifier? Is it necessary to satisfy that the accuracy of the initial linear classifier is higher than 50% or even higher? The answer is no. As long as the initial linear classifier can be trained, theoretically our method should work. It thus is a future open question of how the initial classifier affects. We briefly validate it in Table 4. We run 600 episodes, with each episode training an initial classifier with different classification accuracy. Table 4 shows that most classifiers can get improved by ICI regardless of the initial accuracy (even with an accuracy of 30-40%).

Influence of reduced dimension. In this part, we study the influence of reduced dimension d in our algorithm on 5-way 1-shot *mini*ImageNet experiments. The results with reduced dimension 2, 5, 10, 20, 50, and without dimensionality reduction *i.e.*, $d = 512$, are shown in Table 5. Our algorithm achieves better performance when the reduced dimension is much smaller than the number of instances (*i.e.*, $d \ll n$), which is consistent with the theoretical property [27]. Moreover, we can observe that our model achieves the best accuracy of 66.80% when $d = 5$. Practically, we adopt $d = 5$ in our model.

Normal classification v.s. Prototype classification. Results in Table 1 show that the prototype variant will achieve better performance in the 1-shot task in most cases no matter what the classifier, the task setting, or the dataset. However, this benefit becomes damage when we increase the labeled instance, for example in the 5-shot task. The main reason is that in the 1-shot case, the classifier can only utilize the ground-truth information from one labeled instance, resulting in higher variance for the learning process. Thus using the prototype instead of the pseudo-labeled instance will reduce the variance and potentially reduce the damage of noisy data. When the labeled instances increase, the classifier can learn with less variance. In this case, learning the manifold structure is more important than reducing the variance. Hence using the traditional classification strategy is a better choice.

Influence of dimension reduction algorithms. Furthermore, we study the robustness of ICI to different dimension reduction algorithms. We compare Isomap [71], principal components analysis [72] (PCA), local tangent space alignment [73] (LTSA), multi-dimensional scaling [74] (MDS), locally linear embedding [50] (LLE) and spectral embedding [75] (SE) on 5-way 1-shot *mini*ImageNet experiments.

d	Acc (%)	Alg.	Acc (%)
2	$63.71_{\pm 1.025}$	Isomap [71]	$66.53_{\pm 1.073}$
5	$\mathbf{66.80}_{\pm 1.096}$	PCA [72]	$66.80_{\pm 1.096}$
10	$66.25_{\pm 1.048}$	LTSA [73]	$64.61_{\pm 1.058}$
20	$64.98_{\pm 1.049}$	MDS [74]	$59.99_{\pm 0.941}$
50	$61.54_{\pm 0.980}$	LLE [50]	$67.59_{\pm 1.120}$
512	$57.41_{\pm 0.877}$	SE [75]	$67.70_{\pm 1.117}$

TABLE 5: Influence of dimensionality reduction dimensions and algorithms.

From Table 5 we can observe that ICI is robust across most of the dimensionality reduction algorithms (from LTAS 64.61% to SE 67.7%) except MDS (59.99%). We adopt LLE for dimension reduction in our method.
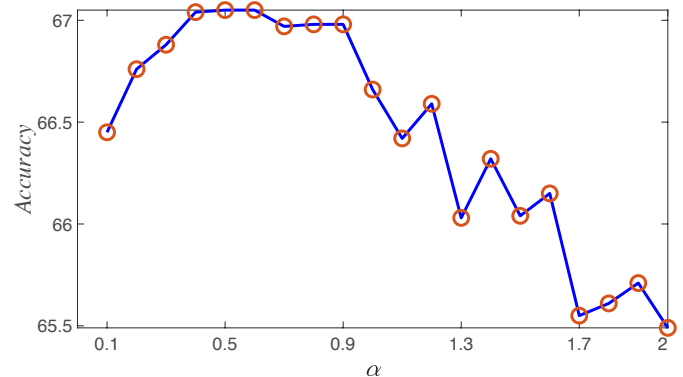

Fig. 6: Validation accuracy with different αs.

Influence of the penalty of logistic regression coefficient in ICI. In Section 3.4, we have shown that the penalty of the logistic regression coefficient is necessary for a unique solution. However, this introduces the hyper-parameters λ_1 and λ_2 which we need to trade-off. Note that since we still aim to find the solution path of γ, which is solved when we use a list of λ_2s. We set $\lambda_1 = \alpha\lambda_2$ for each solution point along the path and search for the best α based on the inference performance on the validation set. Results are shown in Fig. 6, indicating that the performance is maximized when α is set around 0.5. In our experiments, we use $\alpha = 0.5$.

6 CONCLUSION

In this paper, we have proposed a statistical method, called Instance Credibility Inference (ICI) to exploit the distribution support of unlabeled instances for few-shot learning. The proposed ICI effectively select the most trustworthy pseudo-labeled instances according to their credibility to augment the training set. In order to measure the credibility of each pseudo-labeled instance, we propose to solve a hypothesis by increasing the sparsity of the incidental parameters and rank the pseudo-labeled instance with their sparsity degree. Theoretical analysis shows that under conditions of restricted eigenvalue, irrepresentability, and large error, our ICI will find all the correctly-predicted instances from the noisy pseudo-labeled set. Extensive experiments show that our simple approach can establish new state-of-the-arts on four widely used few-shot learning benchmark

datasets including *mini*ImageNet, *tiered*ImageNet, CIFAR-FS, and CUB.

REFERENCES

[1] L. Zhang, T. Xiang, and S. Gong, "Learning a deep embedding model for zero-shot learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2012.

[3] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

[4] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

[5] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[6] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International Conference on Machine Learning*, 2017.

[7] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, 2017.

[8] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, "Learning to compare: Relation network for few-shot learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[9] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," in *Advances in Neural Information Processing Systems*, 2016.

[10] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Advances in Neural Information Processing Systems*, 2014.

[11] Z. Chen, Y. Fu, Y.-X. Wang, L. Ma, W. Liu, and M. Hebert, "Image deformation meta-networks for one-shot learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[12] Z. Chen, Y. Fu, Y. Zhang, Y.-G. Jiang, X. Xue, and L. Sigal, "Multi-level semantic feature augmentation for one-shot learning," *IEEE Transactions on Image Processing*, 2019.

[13] C. Lemke, M. Budka, and B. Gabrys, "Metalearning: a survey of trends and technologies," *Artificial intelligence review*, 2015.

[14] B. Oreshkin, P. R. López, and A. Lacoste, "Tadam: Task dependent adaptive metric for improved few-shot learning," in *Advances in Neural Information Processing Systems*, 2018.

[15] F. Sung, L. Zhang, T. Xiang, T. Hospedales, and Y. Yang, "Learning to learn: Meta-critic networks for sample efficient learning," *arXiv preprint arXiv:1706.09529*, 2017.

[16] Z. Li, F. Zhou, F. Chen, and H. Li, "Meta-sgd: Learning to learn quickly for few-shot learning," *arXiv preprint arXiv:1707.09835*, 2017.

[17] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," *arXiv preprint arXiv:1803.02999*, 2018.

[18] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell, "Meta-learning with latent embedding optimization," *arXiv preprint arXiv:1807.05960*, 2018.

[19] W. Chen, Y. Liu, Z. Kira, Y. F. Wang, and J. Huang, "A closer look at few-shot classification," in *International Conference on Learning Representations*, 2019.

[20] C. Liu, C. Xu, Y. Wang, L. Zhang, and Y. Fu, "An embarrassingly simple baseline to one-shot learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.

[21] Y. Liu, J. Lee, M. Park, S. Kim, E. Yang, S. J. Hwang, and Y. Yang, "Learning to propagate labels: Transductive propagation network for few-shot learning," *arXiv preprint arXiv:1805.10002*, 2018.

[22] M. Ren, E. Triantafillou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel, "Meta-learning for semi-supervised few-shot classification," *arXiv preprint arXiv:1803.00676*, 2018.

[23] Q. Sun, X. Li, Y. Liu, S. Zheng, T.-S. Chua, and B. Schiele, "Learning to self-train for semi-supervised few-shot classification," *arXiv preprint arXiv:1906.00562*, 2019.

[24] T. Joachims, "Transductive inference for text classification using support vector machines," in *International Conference on Machine Learning*, 1999.

[25] L. Qiao, Y. Shi, J. Li, Y. Wang, T. Huang, and Y. Tian, "Transductive episodic-wise adaptive metric for few-shot learning," in *IEEE International Conference on Computer Vision*, 2019.

[26] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng, "Self-taught learning: Transfer learning from unlabeled data," in *International Conference on Machine Learning*, 2007.

[27] J. Fan, R. Tang, and X. Shi, "Partial consistency with sparse incidental parameters," *arXiv preprint arXiv:1210.6950*, 2012.

[28] Y. Wang, C. Xu, C. Liu, L. Zhang, and Y. Fu, "Instance credibility inference for few-shot learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.

[29] Y. Fu, T. M. Hospedales, T. Xiang, and S. Gong, "Learning multi-modal latent attributes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013.

[30] V. Vapnik and V. Vapnik, "Statistical learning theory wiley," *New York*, 1998.

[31] K. P. Bennett and A. Demiriz, "Semi-supervised support vector machines," in *Advances in Neural Information Processing Systems*, 1999.

[32] Y.-F. Li and Z.-H. Zhou, "Towards making unlabeled data never hurt," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2014.

[33] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning." in *International Conference on Learning Representations*, 2017.

[34] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Advances in Neural Information Processing Systems*, 2017.

[35] T. Miayto, A. M. Dai, and I. Goodfellow, "Virtual adversarial training for semi-supervised text classification," 2016.

[36] Q. Xie, E. Hovy, M.-T. Luong, and Q. V. Le, "Self-training with noisy student improves imagenet classification," *arXiv preprint arXiv:1911.04252*, 2019.

[37] M.-R. Amini and P. Gallinari, "Semi-supervised logistic regression," in *ECAI*, 2002.

[38] Y. Grandvalet and Y. Bengio, "Semi-supervised learning by entropy minimization," in *Advances in Neural Information Processing Systems*, 2005.

[39] D.-H. Lee, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *International Conference on Machine Learning workshops*, 2013.

[40] E. Arazo, D. Ortego, P. Albert, N. E. O'Connor, and K. McGuinness, "Pseudo-labeling and confirmation bias in deep semi-supervised learning," *arXiv preprint arXiv:1908.02983*, 2019.

[41] A. Iscen, G. Tolias, Y. Avrithis, and O. Chum, "Label propagation for deep semi-supervised learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[42] W. Shi, Y. Gong, C. Ding, Z. MaXiaoyu Tao, and N. Zheng, "Transductive semi-supervised deep learning using min-max features," in *European Conference on Computer Vision*, 2018.

[43] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel, "A simple neural attentive meta-learner," *arXiv preprint arXiv:1707.03141*, 2017.

[44] J. Neyman and E. L. Scott, "Consistent estimates based on partially consistent observations," *Econometrica: Journal of the Econometric Society*, 1948.

[45] J. Fan and J. Lv, "A selective overview of variable selection in high dimensional feature space," *Statistica Sinica*, 2010.

[46] J. Kiefer and J. Wolfowitz, "Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters," *The Annals of Mathematical Statistics*, 1956.

[47] D. Basu, "On the elimination of nuisance parameters," in *Selected Works of Debabrata Basu*, 2011.

[48] M. Moreira, "A maximum likelihood method for the incidental parameter problem," Tech. Rep., 2008.

[49] Y. Fu, T. M. Hospedales, T. Xiang, J. Xiong, S. Gong, Y. Wang, and Y. Yao, "Robust subjective visual property prediction from crowd-sourced pairwise labels," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015.

[50] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *science*, 2000.

[51] N. Simon, J. Friedman, and T. Hastie, "A blockwise descent algorithm for group-penalized multiresponse and multinomial regression," *arXiv preprint arXiv:1311.6529*, 2013.

[52] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, "Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization," *ACM Transactions on Mathematical Software*, 1997.

[53] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, "Liblinear: A library for large linear classification," *Journal of Machine Learning Research*, 2008.

[54] H.-F. Yu, F.-L. Huang, and C.-J. Lin, "Dual coordinate descent methods for logistic regression and maximum entropy models," *Machine Learning*, 2011.

[55] A. N. Tikhonov and V. Y. Arsenin, "Solutions of ill-posed problems," *New York*, 1977.

[56] P. Zhao and B. Yu, "On model selection consistency of lasso," *Journal of Machine learning research*, vol. 7, no. Nov, pp. 2541–2563, 2006.

[57] M. J. Wainwright, "Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso)," *IEEE transactions on information theory*, 2009.

[58] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," in *International Conference on Learning Representations*, 2017.

[59] L. Bertinetto, J. F. Henriques, P. Torr, and A. Vedaldi, "Meta-learning with differentiable closed-form solvers," in *International Conference on Learning Representations*, 2019.

[60] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The Caltech-UCSD Birds-200-2011 Dataset," California Institute of Technology, Tech. Rep., 2011.

[61] N. Hilliard, L. Phillips, S. Howland, A. Yankov, C. D. Corley, and N. O. Hodas, "Few-shot learning with metric-agnostic conditional embeddings," *arXiv preprint arXiv:1802.04376*, 2018.

[62] E. Triantafillou, R. Zemel, and R. Urtasun, "Few-shot learning through an information retrieval lens," in *Advances in Neural Information Processing Systems*, 2017.

[63] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009.

[64] K. Lee, S. Maji, A. Ravichandran, and S. Soatto, "Meta-learning with differentiable convex optimization," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[65] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, 2015.

[66] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, 2014.

[67] G. Ghiasi, T.-Y. Lin, and Q. V. Le, "Dropblock: A regularization method for convolutional networks," in *Advances in Neural Information Processing Systems*, 2018.

[68] T. Munkhdalai, X. Yuan, S. Mehri, and A. Trischler, "Rapid adaptation with conditionally shifted neurons," *arXiv preprint arXiv:1712.09926*, 2017.

[69] S. W. Yoon, J. Seo, and J. Moon, "Tapnet: Neural network augmented with task-adaptive projection for few-shot learning," *arXiv preprint arXiv:1905.06549*, 2019.

[70] H. Li, D. Eigen, S. Dodge, M. Zeiler, and X. Wang, "Finding task-relevant features for few-shot learning by category traversal," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[71] J. B. Tenenbaum, V. De Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *science*, 2000.

[72] M. E. Tipping and C. M. Bishop, "Probabilistic principal component analysis," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 1999.

[73] Z. Zhang and H. Zha, "Principal manifolds and nonlinear dimensionality reduction via tangent space alignment," *SIAM journal on scientific computing*, 2004.

[74] I. Borg and P. Groenen, "Modern multidimensional scaling: Theory and applications," *Journal of Educational Measurement*, 2003.

[75] M. Belkin and P. Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural computation*, 2003.

APPENDIX
PROOF OF THEOREM 1 AND 2.

Proposition 3. *Assume that $U^\top U$ is invertible. If*

$$\left\| \lambda U_{S^c}^\top U_S \left(U_S^\top U_S\right)^{-1} \hat{V}_S + U_{S^c}^\top \left(I - I_S\right)(U\varepsilon) \right\|_{\max} < \lambda \tag{22}$$

holds for all $\hat{V}_S \in [-1,1]^S$, where $I_S = U_S \left(U_S^\top U_S\right)^{-1} U_S^\top$, then the estimator $\hat{\gamma}$ of Eq. (17) satisfies that

$$\hat{S} = \operatorname{supp}\left(\hat{\gamma}\right) \subseteq \operatorname{supp}\left(\gamma^*\right) = S.$$

Moreover, if the sign consistency

$$\operatorname{sign}\left(\hat{\gamma}_S\right) = \operatorname{sign}\left(\gamma_S^*\right) \tag{23}$$

holds, where

$$\hat{\gamma}_S = \gamma_S^* + \delta_S, \quad \delta_S := \left(U_S^\top U_S\right)^{-1} \left[U_S^\top U\varepsilon - \lambda \operatorname{sign}\left(\gamma_S^*\right)\right].$$

Then $\hat{\gamma}$ is the unique solution of (17) with the same sign as $\hat{\gamma}^*$.

Proof. The duality of Eq. (17) implies that $\hat{\gamma}$ is the unique minimizer if the following holds for some $\hat{V} \in \mathbb{R}^{m \times m}$

$$- U^\top \left(Y - U\hat{\gamma}\right) + \lambda \hat{V} = 0, \tag{24}$$

where for $(i,j) \in \hat{S} := \operatorname{supp}(\hat{\gamma}), \hat{V}_{i,j} = \operatorname{sign}\left(\hat{\gamma}_{i,j}\right)$ and otherwise $\hat{V}_{i,j} \in (-1,1)$.

To ensure that $\hat{S} \subseteq S = \operatorname{supp}\left(\gamma^*\right)$, we just need for $(i,j) \in S^c, \left|\hat{V}_{i,j}\right| < 1$, i.e.

$$\left\| U_{S^c}^\top \left(Y - U_S \hat{\gamma}_S\right) \right\|_{\max} < \lambda, \tag{25}$$

which ensures that $\hat{\gamma}_{i,j} = 0$ for $(i,j) \in S^c$. $\|A\|_{\max} := \max_{i,j} |A_{i,j}|$. On the other hand for $(i,j) \in S$, we obtain the following equation

$$- U_S^\top \left(Y - U_S \hat{\gamma}_S\right) + \lambda \hat{V}_S = 0, \tag{26}$$

which leads to the following solution if $U^\top U$ is invertible

$$\hat{\gamma}_S = \gamma_S^* + \delta_S, \quad \delta_S := \left(U_S^\top U_S\right)^{-1} \left[U_S^\top U\varepsilon - \lambda \hat{V}_S\right]. \tag{27}$$

Plugging (27) into (25) we have

$$\left\| U_{S^c}^\top U\varepsilon - U_{S^c}^\top U_S \left(U_S^\top U_S\right)^{-1} \left[U_S^\top U\varepsilon - \lambda \hat{V}_S\right] \right\|_{\max} < \lambda, \tag{28}$$

or equivalently

$$\left\| \lambda U_{S^c}^\top U_S \left(U_S^\top U_S\right)^{-1} \hat{V}_S + U_{S^c}^\top \left(I - I_S\right)(U\varepsilon) \right\|_{\max} < \lambda, \tag{29}$$

where $I_S = U_S \left(U_S^\top U_S\right)^{-1} U_S^\top$. To ensure the sign consistency, replacing $\hat{V}_S = \operatorname{sign}\left(\gamma_S^*\right)$ in the inequality above leads to the final result. $\square$

Lemma 4. *Assume that ε is i.i.d. Gaussian with*

$$\mathbb{E}\left[\varepsilon\varepsilon^\top\right] \le \sigma^2 I.$$

Then with probability at least

$$1 - 2cn \exp\left(-\frac{\lambda^2 \eta^2}{2\sigma^2 \max_{(i,j) \in S^c} U_{i,j}^2}\right) \tag{30}$$

there holds

$$\left\| U_{S^c}^\top \left(I - I_S\right)(U\varepsilon) \right\|_{\max} \le \lambda\eta \tag{31}$$

and

$$\left\| \left(U_S^\top U_S\right)^{-1} U_S^\top U\varepsilon \right\|_{\max} \le \frac{\lambda\eta}{\sqrt{C_{\min}} \max_{(i,j) \in S^c} |U_{i,j}|}. \tag{32}$$

Proof. Let $Z_c = U_{S^c}^\top \left(I - I_S\right)(U\varepsilon)$, for each $(i,j) \in S^c$ whose variance can be bounded by

$$\operatorname{Var}\left(Z_c^{(i,j)}\right) \le \sigma^2 U_{\cdot,i}^\top \left(I - I_S\right)^2 U_{\cdot,i} \le \sigma^2 \max_{(i,j) \in S^c} U_{i,j}^2.$$

Hoeffding inequality implies that

$$\mathbb{P}\left(\left\| U_{S^c}^\top \left(I - I_S\right)(U\varepsilon) \right\|_{\max} \ge t\right)$$
$$\le 2c(n-m) \exp\left(-\frac{t^2}{2\sigma^2 \max_{(i,j) \in S^c} U_{i,j}^2}\right).$$

Setting $t = \lambda\eta$ leads to the result.

Now let $Z = \left(U_S^\top U_S\right)^{-1} U_S^\top U\varepsilon$, then we have

$$\operatorname{Var}(Z) = \left(U_S^\top U_S\right)^{-1} U_S^\top U \mathbb{E}\left[\varepsilon\varepsilon^\top\right] U^\top U_S \left(U_S^\top U_S\right)^{-1}$$
$$\le \sigma^2 \left(U_S^\top U_S\right)^{-1}$$
$$\le \frac{\sigma^2}{C_{\min}} I$$

Then

$$\mathbb{P}\left(\left\| \left(U_S^\top U_S\right)^{-1} U_S^\top U\varepsilon \right\|_{\max} \ge t\right) \le 2cm \exp\left(-\frac{t^2 C_{\min}}{2\sigma^2}\right).$$

Choose

$$t = \frac{\lambda\eta}{\sqrt{C_{\min}} \max_{(i,j) \in S^c} |U_{i,j}|}, \tag{33}$$

then there holds

$$\mathbb{P}\left\{ \| \left(U_S^\top U_S\right)^{-1} U_S^\top U\varepsilon \|_{\max} \ge \frac{\lambda\eta}{\sqrt{C_{\min}} \max_{(i,j) \in S^c} |U_{i,j}|} \right\}$$
$$\le 2cm \exp\left(-\frac{\lambda^2 \eta^2}{2\sigma^2 \max_{(i,j) \in S^c} U_{i,j}^2}\right).$$

$\square$

Proof of Theorem 1 (Sufficiency). The results follow by applying Lemma 4 to Proposition 3. Inequality (22) holds if condition C2 and the first bound (31) hold, which proves the first part of the theorem. The sign consistency (23) holds if condition C3 and the second bound (32) hold, which gives the second part of the theorem. $\square$

Proof of Theorem 2 (Necessity). In both cases, the necessary conditions are due to the fact that random vectors in the left hand side of (31) and (32) are all Multivariate Gaussian of zero mean and thus zero median, whence if C2 or C3 fails there is an element with probability at least $1/2$ which violates the inequality (22) or (23), respectively. $\square$