Target-oriented Few-shot Transferring via Measuring Task Similarity

Zhipeng Zhou
University of Science and Technology of China
Hefei, China
zzp1994@mail.ustc.edu.cn

Wei Gong*
University of Science and Technology of China
Hefei, China
weigong@ustc.edu.cn

Haoquan Zhou
The First Affiliated Hospital of USTC
Hefei, China
zhouhq2005@qq.com

ABSTRACT

Despite significant progress in recent years, few-shot learning (FSL) still faces two critical challenges. Firstly, most FSL solutions in the training phase rely on exploiting auxiliary tasks, while target tasks are underutilized. Secondly, current benchmarks sample numerous target tasks, each with only an *N-way C-shot* shot query set in the evaluation phase, which is not representative of real-world scenarios. To address these issues, we propose Guidepost, a target-oriented FSL method that can implicitly learn task similarities using a task-level learn-to-learn mechanism and then re-weight auxiliary tasks. Additionally, we introduce a new FSL benchmark that satisfies realistic needs and aligns with our target-oriented approach. Mainstream FSL methods struggle under this new experimental setting. Extensive experiments demonstrate that Guidepost outperforms two classical few-shot learners, i.e., MAML and ProtoNet, and one state-of-the-art few-shot learner, i.e., RENet, on several FSL image datasets. Furthermore, we implement Guidepost as a domain adaptor to achieve high accuracy wireless sensing on our collected WiFi-based human activity recognition dataset.

CCS CONCEPTS

• **Computing methodologies → Machine learning algorithms**; **Knowledge representation and reasoning**.

KEYWORDS

task similarity, learn to learn, few-shot learning

ACM Reference Format:
Zhipeng Zhou, Wei Gong, and Haoquan Zhou. 2023. Target-oriented Few-shot Transferring via Measuring Task Similarity. In *Proceedings of 32nd ACM International Conference on Information and Knowledge Management (CIKM'23).* ACM, New York, NY, USA, 5 pages. https://doi.org/XXXXXXX.XXXXXXX

*Wei Gong is the corresponding author.

CIKM'23, October 21–25, 2023, Birmingham, UK

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

In the deep learning era, training deep models with big data has become a prevailing mode, and we have witnessed lots of successful applications [18]. Such a pattern works well in cases where massively labeled data is available. However, when it comes to such realistic scenarios as medical lesion recognition, drug discovery, etc., where a huge labeled dataset is impossible or very expensive to acquire [21], over-parameterized deep models would easily fall into overfitting. Therefore, few-shot learning (FSL) [13] has attracted many efforts as a promising solution towards tackling this issue and has been developed with many frameworks. Unlike traditional machine learning, it targets learning a model that can fast adapt to a new task with the knowledge extracted from the past tasks, instead of focusing on the optimization of a single task. To mitigate the data scarcity issue, the paradigm of meta learning (most popular FSL framework) usually assumes that there exists an auxiliary dataset that is full of labeled data but has no overlapped categories with the target tasks. As a comparison, traditional machine learning usually follows a simple but strong assumption: the training and test dataset are drawn i.i.d. As such, the trained model can be directly deployed during the test phase since it is supposed to have no out-of-distribution (OOD) samples. Therefore, how to transfer the learned ability on the auxiliary dataset to the target task is the core challenge of FSL, especially when we cannot guarantee the trained model is applicable for new coming target tasks.

Meta learning, the most widely used FSL solution, achieves such rapid adaptation by learning a robust metric space [13] or task-invariant representation [4], which requires only a few target-relevant samples. Taking a task as the smallest unit, these works adopt a so-called episodic training strategy to simulate the test phase. When we look at the design natures of these works, we can see that they strive for no preference on certain tasks so that the trained few-shot learners show equal generalization ability to all tasks. Unfortunately, such an intuition is infeasible when we care more about the performance of the target tasks at hand than the average performance of massive tasks. Purchasing overall performance implies the sacrifice of certain tasks that we care about, which goes against the needs of the most realistic scenario.

To cope with this problem, researchers have started to develop task-adaptive or task-relevant FSL methods [14, 17], which resort to encode more task-level features into the model to help with task-specific fast adaptation. Unfortunately, none of these works leverage the target tasks to guide the meta training process, which cannot actually solve the problem mentioned above. Intuitively, more target task interventions are required during meta training. With this in mind, the source tasks are expected to be re-weighted by target tasks before feeding into the few-shot learner during meta training. However, the current benchmark prevents target tasks from participating in the meta training as it requires an evaluation on the thousands of sampled target tasks to show the generalization. It is meaningful for meta learning since it targets to learn a model that can fast adapt to all tasks. Yet it fails to well evaluate the few-shot learners due to two reasons: (1) A single target task in the current benchmark only contains ~15 shots, which can hardly tell the generalization on this single task. (2) In most realistic few-shot occasions, there is only one target task to be solved, e.g., wireless sensing, defect detection, drug discovery, etc. All we need is to train a model that performs well on this target task with the auxiliary dataset.

To achieve this goal, we first propose a novel task descriptor to represent a single task. Then a task-level learn-to-learn mechanism is employed to implicitly learn the task similarities. After such a module is well-trained, we re-weight the source tasks by measuring their similarities with the target tasks. Besides, to further evaluate the generalization of few-shot learners on the target tasks, we proposed a new FSL benchmark, which engages more target samples into the query set. At last, a target-oriented few-shot learner is trained by these re-weighted source tasks. In a nutshell, our contributions are three-fold: (1) We propose a task-level learn-to-learn mechanism to implicitly learn the task similarities with the proposed task descriptor and develop a target-oriented few-shot learner named Guidepost. (2) A new FSL benchmark is proposed to satisfy realistic needs, and we empirically observe that representative FSL models fail to perform well. (3) We evaluate Guidepost according to two aspects: domain adaptation and few-shot learning. Extensive experiments demonstrate that Guidepost outperforms its baselines on multiple benchmarks, which verifies the superiority of the target-oriented property.

2 THE PRINCIPLE OF GUIDEPOST

2.1 Problem Definition and Background

Formally, FSL adopts episodic training to train a few-shot learner. Assume there exist a training dataset (i.e., auxiliary dataset) $\mathcal{D}^{tr}$, a validation dataset $\mathcal{D}^{val}$, and a test dataset $\mathcal{D}^{ts}$. $\mathcal{D}^{tr}$ is utilized to train the model and D^{val} is utilized to select the best model to test on D^{ts}. All these three datasets have disjoint classes. Suppose we sample a single task $\mathcal{T}_i$ from the task distribution $p(\mathcal{T})$ of a certain dataset (D^{tr}, D^{val} or D^{ts}). During training/validation/test, $\mathcal{T}_i$ is further divided into support set $\mathcal{S}_i$ and query set Q_i. $\mathcal{S}_i$ contains N classes with K shots each class, while Q_i contains N classes with C shots each class. The labels of the query set Q_i are used for training only in the training phase and for evaluation in the reminder phases.

MAML [4] is regarded as a landmark work of FSL, and we take it as an example to elaborate Guidepost in this paper. Denote the initial model function as f and its weights as ϕ_i in the i_{th} episode. A task $\mathcal{T}_i$ is sampled from the task distribution $p(\mathcal{T})$ of the training dataset D^{train}. And $\mathcal{T}_i$ is divided into support set $\mathcal{S}_i$ and query set Q_i. MAML first feeds $\mathcal{S}_i$ into the model and updates its weights to temporary weights ϕ_{tmp}. Then Q_i passes through ϕ_{tmp} and the computed loss is derived from ϕ_i to get the gradients. At last, MAML takes the gradients to update ϕ_i to ϕ_{i+1}. The former update is usually termed inner loop optimization, while the latter update is termed outer loop optimization. Such a bi-level optimization can be denoted as below:

$$\textbf{Inner loop:} \quad \phi_{tmp} \leftarrow \phi_i - \alpha \bigtriangledown_{\phi_i} \mathcal{L}_{\mathcal{T}_i}(S_i, f_{\phi_i}) \tag{1}$$

$$\textbf{Outer loop:} \quad \phi_{i+1} \leftarrow \phi_i - \beta \bigtriangledown_{\phi_i} \mathcal{L}_{\mathcal{T}_i}(Q_i, f_{\phi_{tmp}}) \tag{2}$$

where α and β are the learning rates. From this bi-level optimization, we can observe that MAML does not care about the performance on the support set. Instead, it tries to minimize the error on the query set, which endows MAML with the good capability to generalize via multiple-step optimization in both inner and outer loops. In a more general view, MAML learns task-invariant representation, enabling it quickly adapt to those target tasks that are similar to the trained tasks. However, when a new coming task is quite dissimilar to the trained tasks, MAML fails to adapt quickly and may suffers from significant performance degradation. This is the major challenge that will be tackled in this paper. Our solution is to explicitly measure the similarities between the source tasks and target tasks before regular training, and re-weight the source tasks to help learn a target-oriented few-shot learner. Such an idea provides the potential to bridge the gap between the source domain and the target domain. The detailed design is elaborated as follows.

2.2 Task Descriptor

Before measuring the task similarities, we need to figure out the descriptor of the task. Usually, most FSL methods naively sum up all prototypes in a single task to form a "task prototype" as the task descriptor. Such an approach seems simple but cannot represent the task since it considers little the correlations among prototypes. And the model measures the task similarity by computing the distance among "task prototypes", which is not stable and hard to generalize.

To generate a task descriptor that conserves both discriminative features and correlations among prototypes, we propose a new method as shown in Fig. 2. As for an *N-way K-shot* task $\mathcal{T}$, we first feed all prototypes into a shallow convolutional network C_φ to get their representations denoted as a matrix $P_{N\times D}$, where D is the dimension of the representation. After that, we adopt SVD to do decomposition and extract the components as follows:

$$T_{dp} = U\bar{P}_{N\times D}, \quad U, S, V = \text{SVD}(\bar{P}^T_{N\times D}) \tag{3}$$

$$\bar{P}_{N\times D} = P_{N\times D} - P^{mean}_{1\times D}, \quad P_{N\times D} = C_\varphi(\mathcal{T}) \tag{4}$$

where C_φ is the function of the shallow CNN; $P^{mean}_{1\times D}$ is the mean of $P_{N\times D}$ in the first dimension; U, S, V are the corresponding decomposition results of $\bar{P}^T_{N\times D}$; T_{dp} is the final task descriptor. After representing the task with the descriptor, we can measure task similarity and optimize it. Next, we propose a task-level learn-to-learn mechanism to optimize the task similarity metric module.

2.3 Task-level Learn-to-learn Mechanism

The goal of our task similarity metric module is to assign the related tasks with higher weights so that they can contribute more during few-shot learning. As depicted in Fig. 1(a), with generated task descriptors, we calculate their cosine similarity as the similarity metric. Denoted the parameters of the auxiliary model as θ_0, then we feed task 1 into the model and update it to θ_{tmp}. During this update process, we leverage the similarity of task 1 and task 2 to

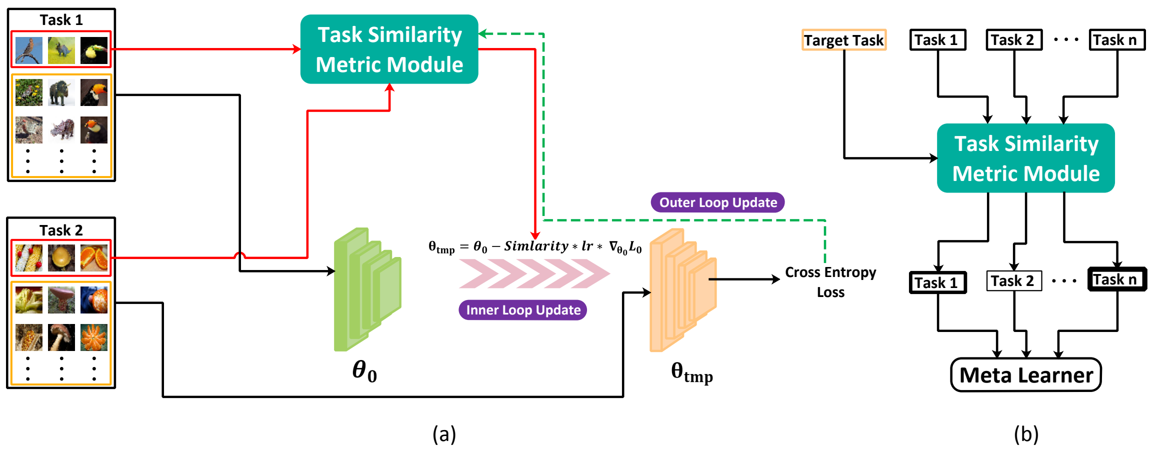

Figure 1: The overview of Guidepost. Illustrated with 3-way 1-shot tasks. (a) is the framework of learning task similarity, and only the task similarity metric module will be updated during training. (b) is the framework of training by re-weighted source tasks. The weight of each source task is assigned by the task similarity metric module trained by step (a).

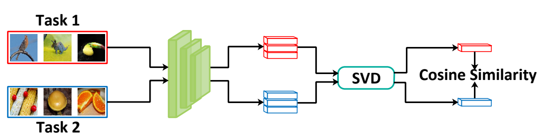

Figure 2: The detailed structure of the task similarity module $\mathcal{M}_\varphi$. It comprises a shallow CNN noted as C_φ and an SVD layer, while φ consists of 64-channel 4 stacked CNN layers.

modulate the inner loop learning rate lr_{in} so that the inner loop optimization is determined by the task correlation:

$$\theta_{tmp} = \theta_0 - Sim * lr_{in} * \nabla_{\theta_0} \mathcal{L}_0 \tag{5}$$

$$Sim = exp(cos\_sim(T^1_{dp}, T^2_{dp})) \tag{6}$$

$$cos\_sim(T^1_{dp}, T^2_{dp}) = \frac{T^1_{dp} \cdot T^2_{dp}}{\left\|T^1_{dp}\right\| \cdot \left\|T^2_{dp}\right\|} \tag{7}$$

Specially, to avoid the semantic confusion between the inner and outer loop, the classifier of the model (i.e., θ_0, θ_{tmp}) is non-parametric. Thus, according to the metric-based FSL, the objectiveness of the inner loop is:

$$\mathcal{L}_0 = -log(p_{\theta_0}(y = k|x)) \tag{8}$$

$$p_{\theta_0}(y = k|x) = \frac{exp(-d(f_{\theta_0}(x), \mathbf{c}_\mathbf{k}))}{\sum_{k'} exp(-d(f_{\theta_0}(x), \mathbf{c}_{\mathbf{k'}}))} \tag{9}$$

where c_k is the prototype of the class k and $d(\cdot)$ is the distance metric function. Task 2 is later fed into θ_{tmp} and the corresponding loss $\mathcal{L}_1$ is computed by the same formulas as Eq. (8) and Eq. (9). And the gradients of the outer loop optimization are derived from φ, since we only update the task similarity metric module with outer loop learning rate lr_{ou} during training:

$$\varphi \leftarrow \varphi - lr_{ou} * \nabla_\varphi \mathcal{L}_1 \tag{10}$$

2.4 Target-oriented Bi-level Optimization

Now that we have a module to measure the similarities among tasks, we can meta learn with source tasks guided by target tasks. As shown in Figure 1(b), each source task sampled from the auxiliary dataset is compared with the target task and then assigned to a weight. With amounts of weighted source tasks, the optimization of the few-shot learner (e.g., MAML) can be formulated as:

$$\phi_{tmp} \leftarrow \phi_i - \alpha\omega \nabla_{\phi_i} \mathcal{L}_{\mathcal{T}_i}(\mathcal{S}_i, f_{\phi_i}) \tag{11}$$

$$\phi_{i+1} \leftarrow \phi_i - \beta \nabla_{\phi_i} \mathcal{L}_{\mathcal{T}_i}(\mathcal{Q}_i, f_{\phi_{tmp}}) \tag{12}$$

where ω is the corresponding weight assigned by $\mathcal{M}_\varphi$. And $\mathcal{L}_{\mathcal{T}_i}$ here is the task-specific loss of the few-shot learner, which could be a parametric/non-parametric classifier or regressor.

3 PERFORMANCE EVALUATION

3.1 Evaluation Setup

We implement Guidepost as both the few-shot learner and domain adaptor to evaluate its rapid adaptation ability. Specifically, we evaluate Guidepost on both image and wireless-based benchmarks. As for the image-based tasks, to satisfy the realistic needs and fit our target-oriented property, we propose a new benchmark based on mainstream FSL datasets. The previous benchmark usually samples a number of target tasks with *N-way C-shot* query set from the test dataset, thus the performance of the few-shot learner on a certain task is regarded as the accuracy on its query set. However, with only *C-shot* in its query set, the few-shot learner cannot be well evaluated. In addition, we only face a single target task in most realistic scenarios. **To address it, we sample the target task, which comprises a few-shot support set and a query set that**

Table 1: Comprehensive 1 on 1 domain adaptation experiment results. 'FT' represents the model is fine-tuned by target 1-shot labeled sample. The bold one represents rank first. 'Guidepost-M' is the MAML version of Guidepost.

Accuracy (%)	**2→1**	**2→3**	**2→4**	**2→5**	**2→6**	**1→3**	**1→4**	**1→5**	**1→6**	**3→4**	**3→5**	**3→6**	**4→5**	**4→6**	**5→6**	**Ave.**
CNN	80.22	80.22	64.81	79.64	82.53	75.15	67.60	81.26	79.30	60.45	80.31	82.48	69.74	62.20	77.53	74.90
CNN_FT	79.34	75.82	64.63	79.75	83.16	74.10	67.78	81.82	**80.34**	60.45	80.31	82.38	67.11	63.56	81.02	74.77
EI	82.70	81.79	65.83	79.42	84.72	70.75	64.72	76.40	76.64	62.89	79.47	82.01	72.09	65.02	78.52	74.86
EI_FT	75.33	73.28	67.25	76.01	85.77	74.40	63.07	77.91	75.91	56.01	83.22	78.94	54.59	62.15	79.25	72.21
MatNet	78.54	75.52	**73.52**	75.17	85.74	76.49	68.32	80.86	76.70	63.02	84.15	83.32	68.81	67.82	**87.29**	76.35
HDA	84.66	83.71	69.97	82.77	84.70	**76.50**	70.23	85.09	78.56	54.86	86.00	83.33	78.91	81.18	72.38	78.17
PACL	75.40	79.31	72.00	79.65	82.21	72.31	**70.82**	**86.94**	80.11	65.82	**88.76**	**84.53**	**86.29**	66.95	80.53	78.11
RFNet	60.37	45.78	50.62	67.23	62.86	47.83	58.33	66.55	56.04	43.92	62.84	56.51	55.84	45.35	58.68	55.92
MetaSense	85.57	87.01	68.99	87.02	87.80	67.69	60.45	65.60	73.21	65.24	80.76	81.18	78.91	78.57	81.26	76.85
MetaSense_FT	85.57	87.16	68.81	87.02	87.75	70.97	60.54	68.46	74.94	65.24	80.76	82.01	78.91	78.73	81.33	77.27
Guidepost-M	**86.62**	**88.51**	69.86	**87.55**	**88.01**	71.64	60.28	72.37	74.54	**65.94**	81.82	81.88	80.98	**81.75**	82.38	**78.28**

Table 2: Few-shot classification results on CUB-200 on 5-way 1-shot. ProtoNet.X is the X-way version of ProtoNet during training.

Accuracy (%)	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Ave.**
MAML	21.38	26.10	41.69	46.44	47.80	48.81	45.42	45.42	39.31	44.75	40.71
Guidepost + MAML	45.86	40.00	46.44	46.78	47.80	52.20	54.24	51.53	47.93	48.14	**45.75**
ProtoNet.5	47.93	43.05	35.59	24.07	44.07	42.37	51.86	46.44	17.24	41.69	39.43
ProtoNet.30	22.41	22.03	31.19	41.69	47.46	26.10	54.58	44.07	32.41	23.73	34.57
Guidepost + ProtoNet.5	46.55	46.10	44.07	44.41	45.76	49.49	46.78	50.17	47.59	48.47	**46.94**
RENet	24.32	23.58	26.04	23.53	22.33	23.32	27.82	25.10	24.23	27.49	24.78
Guidepost + RENet	47.92	49.26	44.44	50.00	49.26	50.37	48.52	42.96	44.15	41.48	**46.84**

Table 3: Few-shot classification results on MiniImageNet on 5-way 1-shot.

Accuracy (%)	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Ave.**
MAML	26.24	26.21	30.88	26.68	26.51	29.28	28.81	26.88	26.54	28.41	27.64
Guidepost + MAML	26.30	26.91	31.68	28.14	27.95	29.62	29.24	28.35	26.58	28.54	**28.33**
ProtoNet.5	21.53	23.48	35.59	12.71	23.38	20.07	24.57	17.29	21.58	23.55	22.38
ProtoNet.30	15.98	21.08	33.27	11.64	23.23	19.72	28.83	15.46	22.39	21.52	21.31
Guidepost + ProtoNet.5	24.37	26.91	23.54	22.80	27.05	23.37	23.81	23.81	21.80	25.11	**24.26**
RENet	23.56	23.29	25.56	24.59	21.95	24.25	26.64	23.84	23.82	25.58	24.31
Guidepost + RENet	24.05	24.53	26.38	25.52	22.03	24.03	27.37	23.16	23.65	26.35	**24.68**

contains all samples per category except for those in the support set and validation set. We also develop Guidepost as a domain adaptor and evaluate it on a wireless sensing task.

3.2 Guidepost for Image

Following the aforementioned evaluation setup, we mainly conduct experiments on two FSL datasets: MiniImageNet and CUB-200. To show the effectiveness of Guidepost, we take three representative FSL methods (i.e., MAML, ProtoNet, and RENet [10]) as the baselines for evaluation. We randomly sample 2000 source tasks for meta training and 10 target tasks for evaluation, and the experimental results on both datasets are presented in Table 2 and Table 3.

CUB-200: We implement MAML and RENet according to their original researches. As for ProtoNet, popular metric-based FSL methods usually train with sampled 30-way 1/5-shot source tasks. Therefore, we provide the experimental results of the ProtoNet trained with 30-way 1-shot source tasks (denoted as ProtoNet.30). We also conduct the experiments of training with 5-way 1-shot source tasks (denoted as ProtoNet.5) for the convenience of deploying Guidepost. Surprisingly, different from most researches reported, ProtoNet.5 largely outperforms ProtoNet.30 under our evaluation setup. According to Table 2, MAML, ProtoNet and RENet are significantly improved with Guidepost, which shows the superiority of the target-oriented property. And although RENet serves as a state-of-the-art method under the traditional FSL setting, it cannot generalize well when engaging more evaluation samples in a single task.

MiniImageNet: According to Table 3, MAML, ProtoNet, and RENet cannot achieve fast adaptation on the larger query set. With Guidepost augmented, ProtoNet still can be benefited with a ~2% performance gain. Differently, MAML and RENet are less augmented by Guidepost though it's still benefited. It might owe to that MiniImageNet has a larger domain shift between its training and test datasets compared to CUB-200 and learning an task-invariant representation is harder than learning a "common" metric space.

3.3 Guidepost for Wireless Sensing

Wireless sensing [3, 6–8, 15, 16, 19, 20, 22, 23] has been a promising solution for smart homes, healthcare, and VR/AR, etc. Here we collect a WiFi-based human activity recognition (HAR) dataset, which employs 9 volunteers to perform 4 kinds of activities (i.e., walking, standing up/sitting down, jumping, and turning around) across 6 different environments (domains, numbered 1 to 6). It totally contains 9156 samples while each domain has a roughly equal number of samples. In this evaluation, we conduct few-shot domain adaptation experiments by sampling 1-shot labeled sample in the target domain. And we compare Guidepost with the following seven mainstream domain adaptation for WiFi-based HAR methods: CNN, EI [9], MatNet [12], HDA [1], PACL [11], MetaSense [5], and RFNet [2]. The main purpose of this experiment is to verify the effectiveness of the Guidepost framework rather than simply achieve performance gain, so we apply the same pre-processing operations to all methods. And all methods have the same four stacked convolutional layers as the backbone network.

According to Table 1, MAML-based Guidepost is an excellent domain adaptor for WiFi-based HAR. Although sometimes may not correctly re-weight the source tasks, it still ranks top in many 1 on 1 domain adaptation cases and achieves the best performance averagely compared to state-of-the-art WiFi-based HAR domain adaptation methods. It indicates the generalization of Guidepost on bridging the gap between the source domain and target domain.

REFERENCES

[1] Shuhao Cui, Xuan Jin, et al. 2020. Heuristic domain adaptation. *Advances in Neural Information Processing Systems* 33 (2020), 7571–7583.

[2] Shuya Ding, Zhe Chen, Tianyue Zheng, and Jun Luo. 2020. RF-net: A unified meta-learning framework for RF-enabled one-shot human activity recognition. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 517–530.

[3] Xiaoyi Fan, Fangxin Wang, Feng Wang, Wei Gong, and Jiangchuan Liu. 2019. When RFID meets deep learning: Exploring cognitive intelligence for activity identification. *IEEE wireless Communications* 26, 3 (2019), 19–25.

[4] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*. PMLR, 1126–1135.

[5] Taesik Gong, Yeonsu Kim, et al. 2019. Metasense: few-shot adaptation to untrained conditions in deep mobile sensing. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. 110–123.

[6] Wei Gong and Jiangchuan Liu. 2017. Robust indoor wireless localization using sparse recovery. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 847–856.

[7] Wei Gong, Kebin Liu, and Yunhao Liu. 2014. Directional diagnosis for wireless sensor networks. *IEEE Transactions on Parallel and Distributed Systems* 26, 5 (2014), 1290–1300.

[8] Wei Gong, Longzhi Yuan, Qiwei Wang, and Jia Zhao. 2020. Multiprotocol backscatter for personal IoT sensors. In *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies*. 261–273.

[9] Wenjun Jiang et al. 2018. Towards environment independent device free human activity recognition. In *Proceedings of the 24th annual international conference on mobile computing and networking*. 289–304.

[10] Dahyun Kang, Heeseung Kwon, et al. 2021. Relational embedding for few-shot classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 8822–8833.

[11] Kai Li, Chang Liu, Handong Zhao, Yulun Zhang, and Yun Fu. 2020. Semi-supervised domain adaptation with prototypical alignment and consistency learning. (2020).

[12] Zhenguo Shi et al. 2020. Environment-robust device-free human activity recognition with channel-state-information enhancement and one-shot learning. *IEEE Transactions on Mobile Computing* (2020).

[13] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems* 30 (2017).

[14] Qiuling Suo, Jingyuan Chou, et al. 2020. Tadanet: Task-adaptive network for graph-enriched meta-learning. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1789–1799.

[15] Fangxin Wang, Jiangchuan Liu, and Wei Gong. 2019. WiCAR: WiFi-based in-car activity recognition with multi-adversarial domain adaptation. In *2019 IEEE/ACM 27th International Symposium on Quality of Service (IWQoS)*. IEEE, 1–10.

[16] Fangxin Wang, Jiangchuan Liu, and Wei Gong. 2020. Multi-adversarial in-car activity recognition using RFIDs. *IEEE Transactions on Mobile Computing* 20, 6 (2020), 2224–2237.

[17] Huaxiu Yao, Yu Wang, Ying Wei, Peilin Zhao, Mehrdad Mahdavi, Defu Lian, and Chelsea Finn. 2021. Meta-learning with an adaptive task scheduler. *Advances in Neural Information Processing Systems* 34 (2021), 7497–7509.

[18] Mingyue Yuan, Dong Yin, Jinwen Ding, Zhipeng Zhou, Chengfeng Zhu, Rui Zhang, and An Wang. 2019. A multi-image Joint Re-ranking framework with updateable Image Pool for person re-identification. *Journal of Visual Communication and Image Representation* 59 (2019), 527–536.

[19] Haoqiang Zhang, Zhipeng Zhou, and Wei Gong. 2021. Wi-adaptor: Fine-grained domain adaptation in wifi-based activity recognition. In *2021 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 1–6.

[20] Jia Zhao, Wei Gong, and Jiangchuan Liu. 2020. Towards scalable backscatter sensor mesh with decodable relay and distributed excitation. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*. 67–79.

[21] Zhipeng Zhou, Lanqing Li, Peilin Zhao, Pheng-Ann Heng, and Wei Gong. 2023. Class-Conditional Sharpness-Aware Minimization for Deep Long-Tailed Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3499–3509.

[22] Zhipeng Zhou, Feng Wang, Jihong Yu, Ju Ren, Zhi Wang, and Wei Gong. 2022. Target-oriented Semi-supervised Domain Adaptation for WiFi-based HAR. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 420–429.

[23] Zhipeng Zhou, Jihong Yu, Zheng Yang, and Wei Gong. 2020. Mobifi: Fast deep-learning based localization using mobile wifi. In *Globecom 2020-2020 IEEE Global Communications Conference*. IEEE, 1–6.