
Supplementary Material of

Understanding Few-Shot Learning: Measuring Task Relatedness and Adaptation Difficulty via Attributes

1 Theoretical Results and Proofs

Recall that we have defined the distance between two categories y_k, y_t in Section 4.1 of the main paper, expressed as $d(y_k, y_t) = \sum_{a_i \in \mathcal{A}} \frac{1}{2L} \sum_{l=1}^L |p(a_i^l|y_k) - p(a_i^l|y_t)|$ if attribute space \mathcal{A} is countable. Based on the distance definition, we introduce the following theoretical results and proofs.

Lemma 1. *Let \mathcal{A} be the attribute space, L be the number of attributes. Assume all attributes are independent of each other given the class label, i.e. $p(a|y) = \prod_{l=1}^L p(a^l|y)$. For all $a_i \in \mathcal{A}$ and any two categories y_k, y_t , the following inequality holds:*

$$\sum_{a_i \in \mathcal{A}} |p(a_i|y_k) - p(a_i|y_t)| \leq d(y_k, y_t) + \Delta, \quad (1)$$

where $\Delta = \sum_{a_i \in \mathcal{A}} \frac{1}{2L} \sum_{l=1}^L (p(a_i^l|y_k) + p(a_i^l|y_t))$.

Proof. Firstly, the following inequality holds,

$$\frac{1}{L} \sum_{i=1}^L x_i \geq \left(\prod_{i=1}^L x_i \right)^{\frac{1}{L}} \geq \prod_{i=1}^L x_i, \quad (2)$$

where x_i is a non-negative real number and ranges from 0 to 1. The proof of Eq. (2) is straightforward: the first inequality is an application of AM-GM inequality (or the inequality of arithmetic and geometric means) on a list of L non-negative real numbers $\{x_1, \dots, x_L\}$, and the second inequality holds because $\prod_{i=1}^L x_i$ ranges from 0 to 1.

Next, we try to prove Eq. (1) based on the above inequality. If all attributes are independent of each other given the class label, for any conditional probabilities $p(a_i|y_k)$ and $p(a_i|y_t)$, we have

$$\begin{aligned} |p(a_i|y_k) - p(a_i|y_t)| &\leq \max(p(a_i|y_k), p(a_i|y_t)) \\ &= \max\left(\prod_{l=1}^L p(a_i^l|y_k), \prod_{l=1}^L p(a_i^l|y_t)\right). \end{aligned} \quad (3)$$

Note that $p(a_i^l|y_k)$ and $p(a_i^l|y_t)$ are both real numbers between 0 and 1. Thus, combining with Eq. (2), we have

$$\begin{aligned} |p(a_i|y_k) - p(a_i|y_t)| &\leq \max\left(\frac{1}{L} \sum_{l=1}^L p(a_i^l|y_k), \frac{1}{L} \sum_{l=1}^L p(a_i^l|y_t)\right) \\ &\leq \frac{1}{L} \sum_{l=1}^L \max(p(a_i^l|y_k), p(a_i^l|y_t)) \\ &\leq \frac{1}{2L} \sum_{l=1}^L (|p(a_i^l|y_k) - p(a_i^l|y_t)|) + \frac{1}{2L} \sum_{l=1}^L (p(a_i^l|y_k) + p(a_i^l|y_t)). \end{aligned} \quad (4)$$

For all $a_i \in \mathcal{A}$, denote $\Delta = \sum_{a_i \in \mathcal{A}} \frac{1}{2L} \sum_{l=1}^L (p(a_i^l|y_k) + p(a_i^l|y_t))$, we have

$$\sum_{a_i \in \mathcal{A}} |p(a_i|y_k) - p(a_i|y_t)| \leq d(y_k, y_t) + \Delta. \quad (5)$$

□

20 **Theorem 1.** *With the same notation and assumptions as Lemma 1, let \mathcal{H} be the hypothesis space with*
 21 *VC-dimension d , f_θ and g_ϕ be the meta-learner and the base-learner as introduced in Section 4.2*
 22 *respectively. Denote g_{ϕ^*} as the best base-learner on some specific task given a fixed meta-learner f_θ .*
 23 *For any training task $\tau_i = (\mathcal{D}_i, S_i)$ and novel task $\tau'_j = (\mathcal{D}'_j, S'_j)$, suppose the number of categories*
 24 *in the two tasks is the same, then with probability at least $1 - \delta$, $\forall g_\phi \circ f_\theta \in \mathcal{H}$, we have*

$$\epsilon(f_\theta, \tau'_j) \leq \hat{\epsilon}(f_\theta, \tau_i) + \sqrt{\frac{4}{m_i} (d \log \frac{2em_i}{d} + \log \frac{4}{\delta})} + d_\theta(\tau_i, \tau'_j) + \Delta' + \lambda, \quad (6)$$

25 *where $\lambda = \lambda_i + \lambda'_j$ is the generalization error of g_{ϕ^*} and f_θ on the two tasks, i.e., $\lambda_i =$*
 26 *$\mathbb{E}_{(x,y) \sim \mathcal{D}_i} [\mathbb{I}(g_{\phi_i^*}(f_\theta(x)) \neq y)]$, $\lambda'_j = \mathbb{E}_{(x,y) \sim \mathcal{D}'_j} [\mathbb{I}(g_{\phi_j^*}(f_\theta(x)) \neq y)]$. Δ' is a term depending*
 27 *on learned base-learners $g_{\phi_i}, g_{\phi'_j}$ and the best base-learners $g_{\phi_i^*}, g_{\phi_j^*}$.*

28 *Proof.* Note that the error $\epsilon(f_\theta, \tau'_j) = \mathbb{E}_{(x,y) \sim \mathcal{D}'_j} [\mathbb{I}(g_{\phi'_j}(f_\theta(x)) \neq y)]$ can be decomposed into two
 29 parts: with the same f_θ , (1) the probability that the learned base-learner $g_{\phi'_j}$ agrees with the best
 30 base-learner $g_{\phi_j^*}$, but they both output the wrong prediction; and (2) the probability that the learned
 31 base-learner $g_{\phi'_j}$ disagrees with the best base-learner $g_{\phi_j^*}$, while $g_{\phi'_j}$ outputs the wrong prediction. The
 32 first part can be bounded by the error of $g_{\phi_j^*}$, and the second part can be bounded by the probability
 33 that $g_{\phi'_j}$ disagrees with $g_{\phi_j^*}$. Denote $Z^{j'} = \{(x, y) | g_{\phi'_j}(f_\theta(x)) \neq g_{\phi_j^*}(f_\theta(x)), (x, y) \sim \mathcal{D}'_j\}$, then
 34 $P_{\mathcal{D}'_j}[Z^{j'}]$ represents the probability that $g_{\phi'_j}$ disagrees with $g_{\phi_j^*}$ based on the same f_θ on distribution
 35 \mathcal{D}'_j . We have

$$\begin{aligned} \epsilon(f_\theta, \tau'_j) &\leq \lambda'_j + P_{\mathcal{D}'_j}[Z^{j'}] \\ &= \lambda'_j + P_{\mathcal{D}_i}[Z^i] + P_{\mathcal{D}'_j}[Z^{j'}] - P_{\mathcal{D}_i}[Z^i] \\ &\leq \lambda'_j + \lambda_i + \epsilon(f_\theta, \tau_i) + P_{\mathcal{D}'_j}[Z^{j'}] - P_{\mathcal{D}_i}[Z^i] \\ &= \lambda + \epsilon(f_\theta, \tau_i) + P_{\mathcal{D}'_j}[Z^{j'}] - P_{\mathcal{D}_i}[Z^i]. \end{aligned} \quad (7)$$

36 Assume that the two tasks both have C categories and $p(y)$ is uniform, we can decompose the proba-
 37 bility $P_{\mathcal{D}'_j}[Z^{j'}]$ as $P_{\mathcal{D}'_j}[Z^{j'}] = \frac{1}{C} \sum_{t=1}^C P_{x|y_t}[Z_t^{j'}]$, where $Z_t^{j'} = \{x | g_{\phi'_j}(f_\theta(x)) \neq g_{\phi_j^*}(f_\theta(x)), x \sim$
 38 $p(x|y_t)\}$. Thus, we have

$$\begin{aligned} \epsilon(f_\theta, \tau'_j) &\leq \lambda + \epsilon(f_\theta, \tau_i) + \frac{1}{C} \left(\sum_{t=1}^C P_{x|y_t}[Z_t^{j'}] - \sum_{k=1}^C P_{x|y_k}[Z_k^i] \right) \\ &= \lambda + \epsilon(f_\theta, \tau_i) + \frac{1}{C} \sum_{e_{tk} \in M} (P_{x|y_t}[Z_t^{j'}] - P_{x|y_k}[Z_k^i]), \end{aligned} \quad (8)$$

39 where M is a maximum matching, which contains C edges and each edge $e_{tk} \in M$ links two
 40 categories y_t, y_k in task τ'_j and τ_i respectively.

41 Next, we consider to replace the conditional distribution $p(x|y)$ with the attribute conditional
 42 distribution $p(a|y)$, because the former is usually unknown and difficult to estimate. For a con-
 43 ditional distribution $p(x|y)$ and a mapping $f_\theta : \mathcal{X} \rightarrow \mathcal{A}$, a new distribution can be induced
 44 over the space \mathcal{A} as $p_\theta(a|y) \triangleq p(f_\theta(x)|y)$. Based on the induced distribution $p_\theta(a|y)$, we have
 45 $P_{x|y_t}[Z_t^{j'}] = P_{a|y_t}[\{a | g_{\phi'_j}(a) \neq g_{\phi_j^*}(a), a \sim p_\theta(a|y_t)\}]$. For clarity, we define $A_t = \{a | g_{\phi'_j}(a) \neq$
 46 $g_{\phi_j^*}(a), a \sim p_\theta(a|y_t)\}$ and $A_k = \{a | g_{\phi_i}(a) \neq g_{\phi_i^*}(a), a \sim p_\theta(a|y_k)\}$. Thus, Eq. (8) can be rewritten
 47 as

$$\epsilon(f_\theta, \tau'_j) \leq \lambda + \epsilon(f_\theta, \tau_i) + \frac{1}{C} \sum_{e_{tk} \in M} (P_{a|y_t}[A_t] - P_{a|y_k}[A_k]). \quad (9)$$

48 Let $A_{t \cup k} = A_t \cup A_k$ be the union set of A_t and A_k , $A_{t \cap k} = A_t \cap A_k$ be the intersection set of A_t
 49 and A_k , then we have another inequality as

$$\begin{aligned} P_{a|y_t}[A_t] - P_{a|y_k}[A_k] &\leq (P_{a|y_t}[A_{t \cup k}] - P_{a|y_k}[A_{t \cup k}]) + |P_{a|y_t}[A_{t \cap k}] - P_{a|y_k}[A_{t \cap k}]| \\ &\quad + |P_{a|y_t}[A_k] - P_{a|y_k}[A_t]|. \end{aligned} \quad (10)$$

50 For clarity, we use two notions Δ_1 and Δ_2 to denote $\frac{1}{C} \sum_{e_{tk} \in M} |P_{a|y_t}[A_{t \cap k}] - P_{a|y_k}[A_{t \cap k}]|$ and
 51 $\frac{1}{C} \sum_{e_{tk} \in M} |P_{a|y_t}[A_k] - P_{a|y_k}[A_k]|$, respectively. Based on Lemma 1 and Eq. (10), we have

$$\begin{aligned} \frac{1}{C} \sum_{e_{tk} \in M} (P_{a|y_t}[A_t] - P_{a|y_k}[A_k]) &\leq \frac{1}{C} \sum_{e_{tk} \in M} (P_{a|y_t}[A_{t \cup k}] - P_{a|y_k}[A_{t \cup k}]) + \Delta_1 + \Delta_2 \\ &= \frac{1}{C} \sum_{e_{tk} \in M} \sum_{a_i \in A_{t \cup k}} (p_\theta(a_i|y_t) - p_\theta(a_i|y_k)) + \Delta_1 + \Delta_2 \\ &\leq \frac{1}{C} \sum_{e_{tk} \in M} d_\theta(y_t, y_k) + \Delta + \Delta_1 + \Delta_2 \\ &= d_\theta(\tau'_j, \tau_i) + \Delta + \Delta_1 + \Delta_2, \end{aligned} \quad (11)$$

52 where $\Delta = \frac{1}{C} \sum_{e_{tk} \in M} \sum_{a_i \in A_{t \cup k}} \frac{1}{2L} \sum_{l=1}^L (p_\theta(a_i^l|y_k) + p_\theta(a_i^l|y_t))$. Denoting $\Delta' = \Delta + \Delta_1 + \Delta_2$,
 53 and combining Eq. (9) and Eq. (11), we can get

$$\epsilon(f_\theta, \tau'_j) \leq \lambda + \epsilon(f_\theta, \tau_i) + d_\theta(\tau_i, \tau'_j) + \Delta'. \quad (12)$$

54 Finally, we apply Vanik-Chervonenkis theory [11] to bound the generalization error $\epsilon(f_\theta, \tau_i)$ in
 55 Eq. (12) by its empirical estimate $\hat{\epsilon}(f_\theta, \tau_i)$. Namely, if S_i is a m_i -size i.i.d sample set, then with
 56 probability at least $1 - \delta$,

$$\epsilon(f_\theta, \tau_i) \leq \hat{\epsilon}(f_\theta, \tau_i) + \sqrt{\frac{4}{m_i} (d \log \frac{2em_i}{d} + \log \frac{4}{\delta})}. \quad (13)$$

57 Combining with Eq. (12), with probability at least $1 - \delta$, we have

$$\epsilon(f_\theta, \tau'_j) \leq \hat{\epsilon}(f_\theta, \tau_i) + \sqrt{\frac{4}{m_i} (d \log \frac{2em_i}{d} + \log \frac{4}{\delta})} + d_\theta(\tau_i, \tau'_j) + \Delta' + \lambda. \quad (14)$$

58 \square

59 **Corollary 1.** *With the same notation and assumptions as Theorem 1, for n training tasks $\{\tau_i\}_{i=1}^n$
 60 and a novel task τ'_j , define $\hat{\epsilon}(f_\theta, \tau_{i=1}^n) = \frac{1}{n} \sum_{i=1}^n \hat{\epsilon}(f_\theta, \tau_i)$, then with probability at least $1 - \delta$,
 61 $\forall g_\phi \circ f_\theta \in \mathcal{H}$, we have*

$$\epsilon(f_\theta, \tau'_j) \leq \hat{\epsilon}(f_\theta, \tau_{i=1}^n) + \frac{1}{n} \sum_{i=1}^n \sqrt{\frac{4}{m_i} (d \log \frac{2em_i}{d} + \log \frac{4}{\delta})} + \frac{1}{n} \sum_{i=1}^n d_\theta(\tau_i, \tau'_j) + \Delta' + \lambda, \quad (15)$$

62 where $\lambda = \frac{1}{n} \sum_{i=1}^n \lambda_i + \lambda'_j$, and Δ' is a term depending on the learned base-learners $\{g_{\phi_i}\}_{i=1}^n, g_{\phi'_j}$
 63 and the best base-learners $\{g_{\phi_i^*}\}_{i=1}^n, g_{\phi_j'^*}$.

64 *Proof.* The proof of Corollary 1 is similar to the proof of Theorem 1. Denote $\lambda = \frac{1}{n} \sum_{i=1}^n \lambda_i + \lambda'_j$,
 65 we have

$$\begin{aligned} \epsilon(f_\theta, \tau'_j) &\leq \lambda'_j + P_{\mathcal{D}'_j}[Z^{j'}] \\ &= \lambda'_j + \frac{1}{n} \sum_{i=1}^n P_{\mathcal{D}_i}[Z^i] + P_{\mathcal{D}'_j}[Z^{j'}] - \frac{1}{n} \sum_{i=1}^n P_{\mathcal{D}_i}[Z^i] \\ &\leq \lambda'_j + \frac{1}{n} \sum_{i=1}^n \lambda_i + \epsilon(f_\theta, \tau_{i=1}^n) + P_{\mathcal{D}'_j}[Z^{j'}] - \frac{1}{n} \sum_{i=1}^n P_{\mathcal{D}_i}[Z^i] \\ &= \lambda + \epsilon(f_\theta, \tau_{i=1}^n) + \frac{1}{n} \sum_{i=1}^n (P_{\mathcal{D}'_j}[Z^{j'}] - P_{\mathcal{D}_i}[Z^i]). \end{aligned} \quad (16)$$

66 Now, we can follow the same procedure as the proof in Theorem 1 and have the following inequality

$$\epsilon(f_\theta, \tau'_j) \leq \lambda + \hat{\epsilon}(f_\theta, \tau_{i=1}^n) + \frac{1}{n} \sum_{i=1}^n \sqrt{\frac{4}{m_i} (d \log \frac{2em_i}{d} + \log \frac{4}{\delta})} + \frac{1}{n} \sum_{i=1}^n d_\theta(\tau_i, \tau'_j) + \Delta', \quad (17)$$

67 where $\Delta' = \frac{1}{n} \sum_{i=1}^n \Delta'_i$, and Δ'_i corresponds to the additional non-negative term as in Eq. (12),
 68 which is derived from $P_{\mathcal{D}'_j}[Z^{j'}] - P_{\mathcal{D}_i}[Z^i]$. \square

69 **Theorem 2.** *With the same notation and assumptions as in Corollary 1, assume the conditional*
70 *distribution $p(x|a^l)$ is task agnostic. If the number of labeled samples m_i in n training tasks and the*
71 *number of labeled samples m'_j in novel task τ'_j tend to be infinite, the following inequality holds:*

$$\frac{1}{n} \sum_{i=1}^n d_{\theta}(\tau_i, \tau'_j) \leq \frac{1}{n} \sum_{i=1}^n d(\tau_i, \tau'_j). \quad (18)$$

72 *Proof.* Without loss of generality, we first consider a single training task τ_i and prove $d_{\theta}(\tau_i, \tau'_j) \leq$
73 $d(\tau_i, \tau'_j)$. Assume the number of training samples is infinite. Thus, for any category y_k in train-
74 ing task τ_i , the induced distribution $p_{\theta}^i(a|y_k)$ is equal to the ground-truth distribution $p^i(a|y_k)$.
75 However, for any category y_t in novel task τ'_j , even with the infinite novel samples, the induced
76 distribution $p_{\theta}^j(a|y_t)$ does not equal the ground-truth distribution $p^j(a|y_t)$. This is because we
77 fix f_{θ} and train a new base-learner $g_{\phi'_j}$ to adapt to novel task τ'_j , as introduced in Section 4.2.
78 Thus, we have $d_{\theta}(\tau_i, \tau'_j) = \frac{1}{C} \sum_{e_{kt} \in M} \frac{1}{L} \sum_{l=1}^L d_{L_1}(p^i(a^l|y_k), p_{\theta}^j(a^l|y_t))$, which measures the
79 distance between the ground-truth distribution $p^i(a^l|y_k)$ and the induced distribution $p_{\theta}^j(a^l|y_t)$.
80 Next, for any attribute l , we consider three cases: (1) the values of attribute l in training and
81 novel tasks are disjoint, which means it is a new attribute or new values of observed attribute
82 for novel task τ'_j . In this case, for any two categories y_k and y_t , the model-related distance
83 $d_{L_1}(p^i(a^l|y_k), p_{\theta}^j(a^l|y_t)) \leq d_{L_1}(p^i(a^l|y_k), p^j(a^l|y_t)) = 1$; (2) the values of attribute l in train-
84 ing and novel tasks are completely overlapped. As the conditional distribution $p(x|a^l)$ is task-
85 agnostic, the attribute classifier f_{θ} can also identify attribute l in novel task τ'_j . In this case,
86 $d_{L_1}(p^i(a^l|y_k), p_{\theta}^j(a^l|y_t)) = d_{L_1}(p^i(a^l|y_k), p^j(a^l|y_t))$; (3) the values of attribute l in training and
87 novel tasks are overlapped but not the same. We can divide the values of attribute l into two parts:
88 the completely overlapped values and the disjoint values, then we can follow the same analysis
89 procedures as in the case (1) and case (2).

90 In summary, we can arrive that the model-related distance $d_{\theta}(\tau_i, \tau'_j)$ is no more than the model-
91 agnostic distance $d(\tau_i, \tau'_j)$, thus we have $\frac{1}{n} \sum_{i=1}^n d_{\theta}(\tau_i, \tau'_j) \leq \frac{1}{n} \sum_{i=1}^n d(\tau_i, \tau'_j)$.

92 □

93 2 Attribute Prototypical Network

94 Our theoretical analysis is based on a specific meta-learning framework with attribute learning. Thus,
95 we instantiate a simple model under that framework as an example, we call this model Attribute
96 Prototypical Network (APNet). A sketch of APNet is presented in Fig. 1.

97 Let $S = \{(x_k, y_k)\}_{k=1}^m$ include all labeled samples in n training tasks. For each sample $(x_k, y_k) \in S$,
98 assume we have L binary attribute labels $\{a_k^l\}_{l=1}^L$. As Corollary 1 reveals, we can reduce the
99 generalization error on novel tasks by maximizing the attribute discrimination ability of meta-learner
100 f_{θ} and the classification ability of base-learner g_{ϕ} . Specifically, we adopt a convolutional network
101 with an additional MLP as f_{θ} . The convolutional network extracts feature representations from
102 images, then the MLP takes features as input and output attribute labels. The attribute classification
103 loss is defined as

$$\mathcal{L}(f_{\theta}) = -\frac{1}{m} \sum_{k=1}^m \frac{1}{L} \sum_{l=1}^L [a_k^l \log z_k^l + (1 - a_k^l) \log(1 - z_k^l)], \quad (19)$$

104 where z_k^l is the l -th dimension of $f_{\theta}(x_k)$ after a sigmoid function.

105 For base-learner, we simply choose a non-parametric base-learner like ProtoNet [9]¹, which takes
106 the attributes generated by the meta-learner f_{θ} as input to calculate cosine distance between test
107 samples and attribute prototypes, then predicts the target label. The few-shot classification loss is

¹Any other models that map $a \in \mathcal{A}$ to $y \in \mathcal{Y}$ are also feasible, such as a MLP in Relation Network [10] and a parametric cosine classifier in Baseline++ [3].

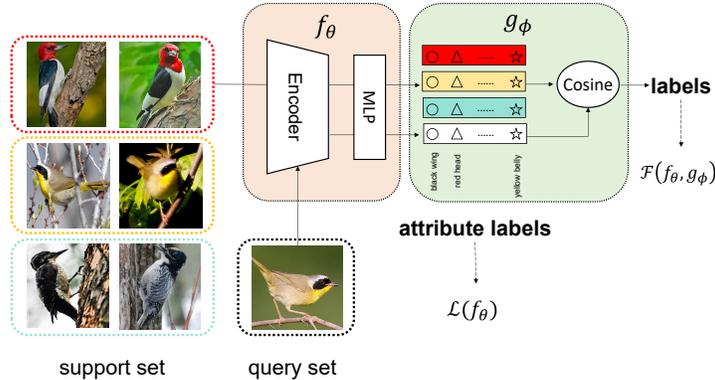


Figure 1: A sketch of APNet.

108 defined as

$$\mathcal{F}(f_\theta, g_{\phi_i}) = -\frac{1}{m_i} \sum_{k=1}^{m_i} y_k \log \frac{\exp(d_k/t)}{\sum_{y_k} \exp(d_k/t)}, \quad (20)$$

109 where $d_k \triangleq \cos(f_\theta(x_k), c_{y_k})$ denotes cosine similarity and c_{y_k} denotes the attribute prototype of
 110 category y_k . t is a scalar temperature factor. In practice, we use a hyperparameter β to balance the
 111 two losses, so that the final training objective is

$$\mathcal{L} = \beta * \mathcal{L}(f_\theta) + \frac{1}{n} \sum_{i=1}^n \mathcal{F}(f_\theta, g_{\phi_i}). \quad (21)$$

112 During the inference phrase, we fix f_θ then calculate the cosine distance between each query sample
 113 and attribute prototypes to predict the target label.

114 3 Experiment Details

115 3.1 Implementation Details

116 We run experiments with APNet and five classical FSL methods (MatchingNet [12], ProtoNet [9],
 117 RelationNet [10], MAML [4], Baseline++ [3]). Here we explain more implementation details about
 118 these methods. As existing work [3] has provided a unified testbed for several different FSL methods,
 119 we use the codebase and run the experiments for the above methods. For a fair comparison, we use
 120 the four-layer convolution network (Conv4) as backbone model for all methods. On the CUB dataset,
 121 we perform standard data augmentation, including random crop, rotation, horizontal flipping and
 122 color jittering, as in [2]. On the SUN dataset, we simply use two augmentation operations, including
 123 image scaling and horizontal flipping. For APNet, we use all provided attribute information (attribute
 124 locations and labels) to calculate the attribute classification loss $\mathcal{L}(f_\theta)$. Because the SUN dataset
 125 does not provide attribute locations, we only use attribute labels to calculate $\mathcal{L}(f_\theta)$. We use the Adam
 126 optimizer [5] with an initial learning rate of 10^{-3} and weight decay of 0. We train models on 5-shot
 127 tasks for 40,000 episodes and on 1-shot tasks for 60,000 episodes. The hyperparameter β is tuned on
 128 the validation set. We set β to 0.6 and 1.0 for 1-shot and 5-shot setting respectively on CUB dataset,
 129 and 0.6 for both settings on SUN dataset.

130 3.2 Complete Results

131 Here we show complete experimental results which have been partially shown in the main paper.
 132 Tab. 1 shows the results on CUB and SUN dataset. Tab. 2 shows the results on *miniImageNet* and the
 133 cross-dataset scenario (*miniImageNet* \rightarrow CUB).

Table 1: 5-way 1-shot and 5-shot performance of different FSL methods on CUB and SUN datasets. Conv4 is used as the backbone model. We report the average accuracy on 600 novel tasks with 95% confidence interval.

Method	Backbone	CUB		SUN	
		5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
MatchingNet	Conv4	61.02 (0.88)	79.99 (0.75)	57.87 (0.95)	76.80 (0.68)
ProtoNet		57.12 (0.94)	76.67 (0.65)	60.20 (0.90)	76.75 (0.65)
RelationNet		61.86 (0.98)	76.63 (0.71)	60.52 (0.91)	76.49 (0.65)
MAML		58.73 (0.97)	76.20 (0.69)	59.65 (0.94)	76.82 (0.68)
Baseline++		60.57 (0.80)	80.17 (0.61)	49.78 (0.82)	74.09 (1.11)
APNet	Conv4	72.96 (0.89)	85.48 (0.55)	60.53 (0.86)	76.35 (0.63)

Table 2: 5-way 1-shot and 5-shot performance of different FSL methods on *miniImageNet* and cross-dataset scenario (*miniImageNet*→CUB). Conv4 is used as the backbone model. We report the average accuracy on 600 novel tasks with 95% confidence interval.

Method	Backbone	<i>miniImageNet</i>		<i>miniImageNet</i> →CUB	
		5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
MatchingNet	Conv4	49.36 (0.79)	62.77 (0.69)	37.48 (0.68)	49.98 (0.66)
ProtoNet		42.53 (0.84)	62.89 (0.72)	33.91 (0.67)	53.74 (0.72)
RelationNet		48.38 (0.80)	64.37 (0.72)	38.19 (0.69)	52.57 (0.66)
MAML		45.70 (0.85)	62.64 (0.72)	36.97 (0.69)	51.60 (0.70)
Baseline++		47.01 (0.71)	66.72 (0.62)	37.11 (0.66)	52.42 (0.67)

134 4 Additional Experiments

135 4.1 Deeper Backbone

136 As mentioned in the main paper, we have shown that TAD can serve as a metric to measure the
 137 adaptation difficulty on novel tasks for different FSL methods. Here we consider how a deeper
 138 backbone affects this conclusion. Following [3], we use ResNet18 as backbone model and train the
 139 five FSL models. The experimental results are shown in Tab. 3. Fig. 2 shows the task distance and the
 140 corresponding accuracy of 2,400 novel tasks. As shown in Fig. 2, we observe similar phenomenon
 141 that with the increase of task distance, the accuracy of these models tends to decrease. This indicates
 142 that the proposed TAD metric works for different FSL methods with a deeper backbone model.

143 4.2 Comparison with Other Metrics

144 Here we present comparisons between proposed TAD and other metrics to show the effectiveness of
 145 it. For comparing different metrics, we design a task selection experiment. More specifically, we
 146 select top 5% novel tasks with the highest distances computed by different metrics, and then evaluate
 147 the accuracy of FSL models on these chosen tasks. The central hypothesis behind this experiment is

Table 3: 5-way 1-shot and 5-shot performance of different FSL methods on CUB and SUN. ResNet18 is used as the backbone model. We report the average accuracy on 600 novel tasks with 95% confidence interval.

Method	Backbone	CUB		SUN	
		5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
MatchingNet	ResNet18	74.62 (0.87)	85.02 (0.54)	67.42 (0.90)	78.30 (0.67)
ProtoNet		74.04 (0.88)	87.30 (0.49)	67.71 (0.87)	81.68 (0.60)
RelationNet		70.37 (0.98)	84.41 (0.57)	63.85 (0.93)	79.60 (0.67)
MAML		70.76 (1.04)	81.31 (0.70)	61.47 (0.98)	75.24 (0.73)
Baseline++		70.20 (0.93)	84.11 (0.57)	53.06 (0.77)	74.21 (0.68)

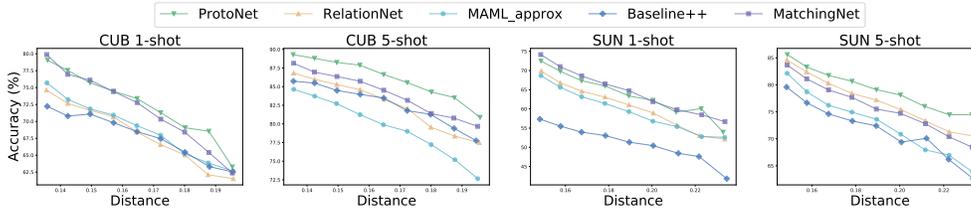


Figure 2: Accuracy of different methods in terms of the average task distance. From left to right, 5-way 1-shot and 5-shot on CUB/SUN. ResNet18 is used as the backbone model.

Table 4: The difference in accuracy between the chosen tasks and all novel tasks for various FSL models. In the last column, we report the computation time (seconds) of different metrics. We run the experiment on the CUB dataset with 5-way 1-shot setting.

Metrics	MatchingNet	ProtoNet	RelationNet	Baseline++	APNet	Time
FID	-3.71	-1.78	-0.57	-0.33	-3.63	480
EMD	-0.24	-1.33	1.60	-0.75	1.17	22
Task2Vec	-0.54	-0.30	-3.64	-0.54	-2.45	6000
TAD	-8.23	-6.64	-7.39	-3.52	-7.24	3

148 that if a distance metric can better reflect task difficulty, then novel tasks with the highest distances
 149 should be more challenging. We choose three methods for comparison, which have been proposed in
 150 the few-shot learning or related area: (1) Fréchet Inception Distance (**FID**) [6], FID is a metric to
 151 measure the distance between two image distributions by comparing their mean and covariance. (2)
 152 Earth Mover’s Distance (**EMD**) [7], EMD is a measure of dissimilarity between two distributions
 153 by considering the distance as the cost of moving images from one distribution to the other. (3)
 154 **Task2Vec** [1], TaskVec is a task embedding method which represents each task as an embedding
 155 with Fisher Information Matrix, and the norm of embedding reflects task difficulty. Note that the
 156 above three methods rely on a pretrained model. Following [7], we use ResNet-101 pre-trained on
 157 ImageNet for them. Tab. 4 illustrates the results of different metrics on the CUB dataset. We find that,
 158 with human-annotated attributes, TAD significantly outperforms other three methods in identifying
 159 more challenging novel tasks across all FSL models, demonstrating the effectiveness of TAD metric.
 160 Furthermore, the computational efficiency of TAD greatly surpasses other methods, as illustrated in
 161 the last column of table. Notably, TAD requires only 3 seconds to compute across 2400 novel tasks,
 162 underscoring its advantage of ease of computation.

163 5 Analysis of Auto-Annotated Attributes

164 We try to evaluate the quality of the auto-annotated attributes generated by pretrained CLIP and
 165 then give some examples for qualitative analysis. Due to the absence of attribute annotations in
 166 the miniImagenet dataset, we collect the annotations ourselves. Initially, we predefine 25 attribute
 167 labels (as shown in Tab. 5) following [8], and then randomly select 50 images for annotation. By
 168 comparing the ground-truth annotations with the results produced by the CLIP model, we discern
 169 that the average accuracy across the 25 attributes approaches 0.65. Notably, we observe that CLIP
 170 achieves good performance across the majority of attributes, with accuracies ranging from 0.7 to
 171 0.9. However, it fails on some attributes such as "white," "pink," "smooth," and "shiny," where the
 172 accuracy decreases to approximately 0.2. Fig. 3 shows qualitative examples of the obtained attributes.
 173 We find that the common wrong cases are color attributes, as the CLIP model always predicts more
 174 colors than manual annotations. Thus, in all other experiments under the cross-dataset scenario, we
 175 remove 11 color attributes and only consider the remaining 14 attributes when calculating the task
 176 distance for the TAD metric.

177 References

178 [1] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhansu Maji,
 179 Charless C Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-

Table 5: Details of pre-defined 25 attribute labels, including color, pattern, shape and texture.

Attributes	
Color	black, blue, brown, gray, green, orange, pink, red, violet, white, yellow
Pattern	spotted, striped
Shape	long, round, rectangular, square
Texture	furry, smooth, rough, shiny, metallic, vegetation, wooden, wet

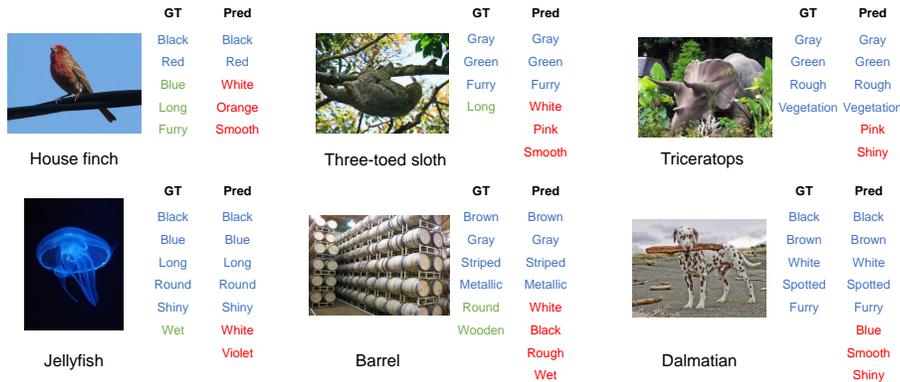


Figure 3: Qualitative examples of auto-annotated attributes. Attributes predicted by the CLIP model and overlapped with the ground truth are highlighted in blue, while the ground truth without match is marked in green. False positive predictions are represented in red.

- 180 learning. In *ICCV*, 2019.
- 181 [2] Kaidi Cao, Maria Brbic, and Jure Leskovec. Concept learners for few-shot learning. In *ICLR*,
182 2021.
- 183 [3] Weiyu Chen, Yencheng Liu, Zsolt Kira, Yuchiang Frank Wang, and Jiabin Huang. A closer
184 look at few-shot classification. In *ICLR*, 2019.
- 185 [4] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adapta-
186 tion of deep networks. In *ICML*, 2017.
- 187 [5] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*,
188 2015.
- 189 [6] Timo Milbich, Karsten Roth, Samarth Sinha, Ludwig Schmidt, Marzyeh Ghassemi, and Bjorn
190 Ommer. Characterizing generalization under out-of-distribution shifts in deep metric learning.
191 In *NeurIPS*, 2021.
- 192 [7] Jaehoon Oh, Sungnyun Kim, Namgyu Ho, Jin-Hwa Kim, Hwanjun Song, and Se-Young
193 Yun. Understanding cross-domain few-shot learning based on domain similarity and few-shot
194 difficulty. In *NeurIPS*, 2022.
- 195 [8] Olga Russakovsky and Li Fei-Fei. Attribute learning in large-scale datasets. In *ECCV*, 2010.
- 196 [9] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In
197 *NeurIPS*, 2017.
- 198 [10] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales.
199 Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018.
- 200 [11] Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media,
201 1999.
- 202 [12] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks
203 for one shot learning. In *NeurIPS*, 2016.