

---

# The Curse of Low Task Diversity: On the Failure of Transfer Learning to Outperform MAML and their Empirical Equivalence

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Recently, it has been observed that a transfer learning solution might be all we need  
2 to solve many few-shot learning benchmarks – thus raising important questions  
3 about when and how meta-learning algorithms should be deployed. In this paper,  
4 we seek to clarify these questions by 1. proposing a novel metric – the *diversity*  
5 *coefficient* – to measure the diversity of tasks in a few-shot learning benchmark  
6 and 2. by comparing MAML and transfer learning under fair conditions (same  
7 architecture, same optimizer and all models trained to convergence). Using the  
8 diversity coefficient, we show that the popular MiniImagenet and Cifar-fs few-shot  
9 learning benchmarks have low diversity. This novel insight contextualizes claims  
10 that transfer learning solutions are better than meta-learned solutions in the regime  
11 of low diversity under a fair comparison. Specifically, we empirically find that a low  
12 diversity coefficient correlates with a high similarity between transfer learning and  
13 Model-Agnostic Meta-Learning (MAML) learned solutions in terms of accuracy  
14 at meta-test time and classification layer similarity (using feature based distance  
15 metrics like SVCCA, PWCCA, CKA, and OPD). To further support our claim,  
16 we find this meta-test accuracy holds even as the model size changes. Therefore,  
17 we conclude that in the low diversity regime, MAML and transfer learning have  
18 equivalent meta-test performance when both are compared fairly. We also hope  
19 our work inspires more thoughtful constructions and quantitative evaluations of  
20 meta-learning benchmarks in the future.

## 21 1 Introduction

22 The success of deep learning in computer vision (1; 2), natural language processing (3; 4), game  
23 playing (5; 6; 7) and more, keeps motivating a growing body of applications of deep learning on  
24 an increasingly wide variety of domains. In particular, deep learning is now routinely applied to  
25 few-shot learning – a research challenge that assesses a model’s ability to learn to adapt to new tasks,  
26 new distributions, or new environments. This has been the main research area where meta-learning  
27 algorithms have been applied – since such a strategy seems promising in a small data regime due to  
28 its potential to *learn to learn* or *learn to adapt*. However, it was recently shown (8) that a transfer  
29 learning model with a fixed embedding can match and outperform many modern sophisticated meta-  
30 learning algorithms on numerous few-shot learning benchmarks (9; 10; 11; 12). This growing body of  
31 evidence – coupled with these surprising results in meta-learning – raise the question if researchers are  
32 applying meta-learning with the right inductive biases (13; 14) and designing appropriate benchmarks  
33 for meta-learning. Our evidence suggests this is not the case.

34 In this work, we show that when the task diversity – a novel measure of variability across tasks – is  
35 low, then MAML (Model Agnostic Meta-Learning) (15) learned solutions have the same accuracy

36 as transfer learning (i.e., a supervised learned model with a fine-tuned final linear layer). We want  
37 to emphasize the importance of doing such an analysis fairly: with the same architecture, same  
38 optimizer and all models trained to convergence. This empirical equivalence remained true even as  
39 the model size changed – thus further suggesting this equivalence is more a property of the data than  
40 of the model. Therefore, we suggest taking a problem-centric approach to meta-learning and suggest  
41 applying Marr’s level of analysis (16; 17) to few-shot learning – to identify the family of problems  
42 suitable for meta-learning. Marr emphasized the importance of understanding the computational  
43 problem being solved and not only analyzing the algorithms or hardware that attempts to solve  
44 them. An example given by Marr is marveling at the rich structure of bird feathers without also  
45 understanding the problem they solve is flight. Similarly, there has been analysis of MAML solutions  
46 and transfer learning without putting the problem such solutions should solve into perspective (18; 19).  
47 Therefore, in this work, we hope to clarify some of these results by partially placing the current  
48 state of affairs in meta-learning from a problem-centric view. In addition, the novelty of our analysis  
49 compared to previous work is that we make analysis intrinsic of the data as a first class citizen.

50 **Our contributions** summarized as follows:

- 51 1. *We propose a novel metric that quantifies the **intrinsic diversity** of the data of a few-shot*  
52 *learning benchmark. We call it the *diversity coefficient*. It enables analysis of meta-learning*  
53 *algorithms through a *problem-centric framework*. It also goes beyond counting the number of*  
54 *classes or number of data points or counting the number concatenated data sets – and instead*  
55 *quantifies the expected diversity/variability of tasks in a few-shot learning benchmark.*
- 56 2. *We analyze the two most prominent few-shot learning benchmarks – MiniImagenet and*  
57 *Cifar-fs – and show that their diversity is low. These results are robust across different ways*  
58 *to measure the diversity coefficient, suggesting that our approach is robust.*
- 59 3. *With this context, we partially clarify the surprising results from (19) by comparing their*  
60 *transfer learning method against models trained with MAML (15). In particular, when*  
61 *making a fair comparison, transfer learning method with a fixed feature extractor fails to*  
62 *outperform MAML. We define a fair comparison when the two methods are compared using*  
63 *the same architecture (backbone), same optimizer and all models trained to convergence.*  
64 *We also show that their final layer makes similar predictions according to neural network*  
65 *distance techniques like distance based Singular Value Canonical Correlation Analysis*  
66 *(SVCCA), Projection Weighted (PWCCA), Linear Centered Kernel Analysis (LINCKA)*  
67 *and Orthogonal Procrustes Distance (OPD). This equivalence holds even as the model size*  
68 *increases.*
- 69 4. *Interestingly, we also find that even in the regime where task diversity is low (in MiniIma-*  
70 *genet and Cifar-fs), the features extracted by supervised learning and MAML are different –*  
71 *implying that the mechanism by which they function is different despite the similarity of*  
72 *their final predictions.*
- 73 5. *As an actionable conclusion, we provide a metric that can be used to analyze the intrinsic*  
74 *diversity of the data in a few-shot learning benchmarks and therefore build more thoughtful*  
75 *environments to drive research in meta-learning. In addition, our evidence suggests the*  
76 *following test to predict the empirical equivalence of MAML and transfer learning: if the*  
77 *task diversity is low, then transfer learned solutions might fail to outperform meta-learned*  
78 *solutions. This test is easy to run because our diversity coefficient can be done using*  
79 *the Task2Vec method (20) using pre-trained neural network. We also found that random*  
80 *networks were consistent with the results of pre-trained networks on Imagenet.*

81 We hope that this line of work inspires a problem-centric first approach to meta-learning – which  
82 appears to be especially sensitive to the properties of the problem in question. Therefore, we hope  
83 future work takes a more thoughtful and **quantitative** approach to benchmark creation – instead of  
84 focusing only on making huge data sets.

## 85 2 Background

86 In this section, we provide a summary of the background needed to understand our main results.

87 **Model-Agnostic Meta-Learning (MAML):** The MAML algorithm (15) attempts to meta-learn  
88 an initialization of parameters for a neural network so that it is primed for fast gradient descent

89 adaptation. It consists of two main optimization loops: 1) an outer loop used to prime the parameters  
 90 for fast adaptation, and 2) an inner loop that does the fast adaptation. During meta-testing, only the  
 91 inner loop is used to adapt the representation learned by the outer loop.

92 **Transfer Learning with Union Supervised Learning (USL):** Previous work (19) shows that  
 93 an initialization trained with supervised learning, on a union of all tasks, can outperform many  
 94 sophisticated methods in meta-learning. In particular, their method consists of two stages: 1) first  
 95 they use a union of all the labels in the few-shot learning benchmark during meta-training and train  
 96 with standard supervised learning (SL), then 2) during the meta-testing, they use an inference method  
 97 common in transfer learning: extract a fixed feature from the neural network and fully fine-tune the  
 98 final classification layer (i.e., the head). Note that our experiments only consider when the final layer  
 99 is regularized Logistic Regression trained with LBGFS.

100 **Distances for Deep Neural Network Feature Analysis:** To compute the distance between neural  
 101 networks we use the distance versions of Singular Value Canonical Correlation Analysis (SVCCA)  
 102 (21), Projection Weighted Canonical Correlation (PWCCA) (22), Linear Centered Kernel Analysis  
 103 (LINCKA) (23) and Orthogonal Procrustes Distance (OPD) (24). These distances are in the interval  
 104  $[0, 1]$  and are not necessarily a formal distance metric but are guaranteed to be zero when their  
 105 inputs are equal and nonzero otherwise. This is true because SVCCA, PWCCA, LINCKA are based  
 106 on similarity metrics and OPD is already a distance. Note that we use the formula  $d(X, Y) =$   
 107  $1 - sim(X, Y)$  for our distance metrics where  $sim$  is one either SVCCA, PWCCA, LINCKA  
 108 similarity metric and  $X, Y$  are matrices of activations (called layer matrices). The distance between  
 109 two models is computed by choosing a layer and then comparing the features/activations after  
 110 adaptation for that layer given a batch of tasks represented as a support and query set. A more  
 111 thorough overview of these metrics for the analysis of internal representations for convolutional  
 112 neural networks (CNNs) can be found in the appendix, section G.

113 **Task2Vec Embeddings for Distances between Tasks:** The diversity coefficient we propose is  
 114 the expectation of distance between tasks (explain in more detail in section 3). Therefore, it is  
 115 essential to define the distance between different pairs of tasks. We choose the cosine distance  
 116 between Task2Vec (vectorial) embeddings as in (20). Therefore, we provide a summary of the  
 117 Task2Vec method to compute task embeddings. The vectorial representation of tasks provided by  
 118 Task2Vec (20) is the vector of diagonal entries of the Fisher Information Matrix (FIM) given a fix  
 119 neural network as a feature extractor – also called a **probe network** – after fine-tuning the final  
 120 classification layer to the task. The authors explain this is a good vectorial representation of tasks  
 121 because 1. It approximately indicates the most informative weights for solving the current task  
 122 (up to a second order approximation) 2. For rich probe networks like CNNs, the diagonal is more  
 123 computationally tractable. We choose Task2Vec because the original authors provide extensive  
 124 evidence that their embeddings correlate with semantic and taxonomic relations between different  
 125 visual classes – making it a convincing embedding for tasks (20). The Task2Vec embedding of task  $\tau$   
 126 is the diagonal of the following matrix:

$$\hat{F}_{D_\tau, f_w} = \hat{F}(D_\tau, f_w) = \mathbb{E}_{x, y \sim \hat{p}(x|\tau)p(y|x, f_w)} [\nabla_w \log p(y | x, f_w) \nabla_w p(y | x, f_w)^\top] \quad (1)$$

127 where  $f_w$  is the neural networks used as a feature extractor with architecture  $f$  and weights  $w$ ,  
 128  $\hat{p}(x | \tau)$  is the empirical distribution defined by the training data  $D_\tau = \{(x_i, y_i)\}_{i=1}^n$  for task  $\tau$ , and  
 129  $p(y | x, f_w)$  is a deep neural network trained to approximate the (empirical) posterior  $\hat{p}(y | x, \tau)$ .  
 130 We'd like to emphasize that there is a dependence on target label since Task2Vec fixes the  
 131 feature extractor (using  $f_w$ ) and then fits the final layer (or "head") to approximate the task posterior  
 132 distribution  $\hat{p}(y | x, \tau)$ .

### 133 3 Definition of the Diversity Coefficient

134 The diversity coefficient aims to measure the intrinsic diversity (or variability) of tasks in a few-shot  
 135 learning benchmark. At a high level, the diversity coefficient is the expected distance between a  
 136 pair of different tasks **given a fixed probe network**. In this work, we choose the distance to be the  
 137 cosine distance between vectorial representations (i.e. embeddings) of tasks according to Task2Vec  
 138 as described in section 2. Using a fixed probe networks is essential because: 1. Using a fixed probe  
 139 network means that the distances between different tasks are comparable, as discussed in the original  
 140 Task2Vec (20) and 2. Since we are computing the distance between different tasks, we need to make  
 141 sure the difference comes from intrinsic properties of the data and not from a different source, e.g. if

142 one uses different models then this might confound the source of variability in our metric. We define  
 143 the **diversity coefficient** of a few-shot learning benchmark  $B$  as follows:

$$\hat{div}(B) = \mathbb{E}_{\tau_1 \sim \hat{p}(\tau|B), \tau_2 \sim \hat{p}(\tau|B)} \mathbb{E}_{D_1 \sim \hat{p}(x_1, y_1 | \tau_1), D_2 \sim \hat{p}(x_2, y_2 | \tau_2)} \left[ d(\hat{F}_{D_1, f_w}, \hat{F}_{D_2, f_w}) \right] \quad (2)$$

144 where  $f_w$  is the neural networks used as a feature extractor with architecture  $f$  and weights  $w$ ,  
 145  $\hat{p}(x | \tau)$  is the empirical distribution defined by the training data  $D_\tau = \{(x_i, y_i)\}_{i=1}^n$  for task  $\tau$ ,  
 146  $\tau_1, \tau_2$  are tasks sampled from the empirical distribution of tasks  $\hat{p}(\tau | B)$  for the current benchmark  
 147  $B$  (i.e. a batch of tasks with their data sets  $\mathcal{D} = (\tau_i, D_{\tau_i})_{i=1}^N$ ), a task  $\tau_i$  is the probability distribution  
 148  $p(x, y | \tau)$  of the data,  $d$  is a distance metric (for us cosine),  $f_w$  is the neural networks used as  
 149 a feature extractor with architecture  $f$  and weights  $w$ , and  $\hat{p}(x | \tau)$  is the empirical distribution  
 150 defined by the training data  $D_\tau = \{(x_i, y_i)\}_{i=1}^n$  for task  $\tau$ . We’d also like to recall the reader that the  
 151 definition of a task in this setting is of a n-way, k-shot few-shot learning task. Therefore, each task has  
 152 n classes sampled with k examples used for the adaptation. We’d like to emphasize that the adaptation  
 153 here is only to fine-tune the final layer according to the Task2Vec method for the correct computation  
 154 of the FIM. Therefore, in this setting we combine the support and query set as the split is not relevant  
 155 for the computation of the task embedding using Task2Vec. Note that the above formulation can be  
 156 easily adapted to any distance function between tasks, and is not necessarily specific to using the  
 157 FIM or cosine distance. For example, given the true distributions for tasks one can use real distances  
 158 between probability distributions e.g. Hellinger distance. In addition, it is obvious one can use a  
 159 distance function besides the cosine distance – but choose it in accordance to the original work of  
 160 Task2Vec (20).

## 161 4 Experiments

162 This section explains the experiments backing up our main results outlined in our list of contributions.  
 163 Experimental details are provided in the supplementary section A and the learning curves displaying  
 164 the convergence for a fair comparison are in supplementary section B.

### 165 4.1 The Diversity Coefficient of MiniImagenet and Cifar-fs

166 To put our analysis into a problem-centric framework, we first analyze the problem they are trying  
 167 to solve through the diversity coefficient. Recall that the diversity coefficient aims to quantify the  
 168 intrinsic variation of tasks in a few-shot learning benchmark. We show that the diversity coefficient  
 169 of the popular MiniImagenet and Cifar-fs benchmarks are low with good confidence intervals using  
 170 four different probe networks in table I.

Probe Network	Diversity on MI	Diversity on Cifar-fs
Resnet18 (pt)	0.117 ± 2.098e-5	0.100 ± 2.18e-5
Resnet18 (rand)	0.0955 ± 1.29e-5	0.103 ± 1.05e-5
Resnet34 (pt)	0.0999 ± 1.95e-5	0.0847 ± 3.06e-5
Resnet34 (rand)	0.0620 ± 8.12e-6	0.0643 ± 9.64e-6

Table 1: **The diversity coefficient of MiniImagenet (MI) and Cifar-fs is low.** The diversity coefficient was computed using the cosine distance between different standard n-way, k-shot classification tasks from the few-shot learning benchmark using the Task2Vec method described in section 3. We used n=5 (number of classes) and k=20 (number of examples per class) since we can use the whole task data to compute the diversity coefficient (no splitting of support and query set required). We used Resnet18 and Resnet34 networks as probe networks – both pre-trained on ImageNet (indicated as “pt” on table) and randomly initialized (indicated as “rand” on table). We observe that both type of networks and weights give similar diversity results. All confidence intervals were at 95%. To compute results, we used 500 few-shot learning tasks and only compare pairs of different tasks. This results in  $(500^2 - 500)/2 = 124,750$  pair-wise distances used to compute the diversity coefficient.

### 171 4.2 Low Diversity Correlates with Equivalence of MAML and Transfer Learning

172 Now that we have placed ourselves in a problem-centric framework and shown the diversity coefficient  
 173 of the popular MiniImagenet and Cifar-fs benchmarks are low – we proceed to show the failure of

174 transfer learning (with USL) to outperform MAML. Crucially, the analysis was done using a fair  
 175 comparison: using the same model architecture, optimizer, and training all models to convergence  
 176 – details in section A. We used the five-layer CNN used in (15; 25) and Resnet12 as in (19). We  
 177 provide evidence that in the setting of low diversity:

- 178 1. The accuracy of an adapted MAML meta-learner vs. an adapted USL pre-trained model are  
 179 similar and statistically significant, except for one result where transfer learning with USL  
 180 is worse. This is shown in table 2 and 1.
- 181 2. The distance for the classification layer decreases sharply according to four distance-based  
 182 metrics – SVCCA, PWCCA, LINCKA, and OPD – as shown in figure 2. This implies the  
 183 predictions of the two are similar.

184 For the first point, we emphasize that tables 1 and table 2 taken together support our central hypothesis:  
 185 that models trained with meta-learning are not inferior to transfer learning models (using USL) when  
 186 the diversity coefficient is low. Careful inspection reveals that the methods have the same meta-test  
 187 accuracy with intersecting confidence intervals – making the results statistically significant across  
 188 few-shot benchmarks and architectures. The one exception is the third set of bar plots, where transfer  
 189 learning with USL is in fact worse.

190 For the second point, refer to figure 2 and observe that as the depth of the network increases, the  
 191 distance between the activation layers of a model trained with MAML vs USL increases until it  
 192 reaches the final classification layer – where all four metrics display a noticeable dip. In particular,  
 193 PWCCA considers the two prediction layers identical (approximately zero distance). This final point  
 194 is particularly interesting because PWCCA is weighted according to the CCA weights that stabilize  
 195 with the final predictions of the network. This means that the PWCCA distance value is reflective of  
 196 what the networked actually learned and gives a more reliable distance metric (for details, refer to the  
 197 appendix section G.5). This is important because this supports our main hypothesis: that at prediction  
 198 time there is an equivalence between transfer learning and MAML when the diversity coefficient is  
 199 low.




Figure 1: **MAML trained models and union supervised trained (USL) models have statistically equivalent meta-test accuracy for MiniImagenet and Cifar-fs with Resnet12 and five layer CNNs.** This holds for both the Resnet12 architecture used in (19) and the 5 layer CNN (indicated as “5CNN”) in (25). Results used a (meta) batch-size of 100 tasks and 95% confidence intervals. All MAML models were trained with 5 inner steps during meta-training. “MAML5” and “MAML10” in the bar plot indicates the adaptation method used at test time i.e. we used 5 inner steps and 10 inner steps at test time. MiniImagenet is abbreviated as “MI” in the figure.

### 200 4.3 Is the Equivalence of MAML and Transfer Learning related to Model Size or Low 201 Diversity?

202 An alternative hypothesis to explain the equivalence of transfer learning (with USL) and MAML could  
 203 be due to the capabilities of large neural networks to be better meta-learners in general. Inspired by the  
 204 impressive ability of large language models to be few-shot (or even zero-shot) learners (4; 27; 28; 3) –  
 205 we hypothesized that perhaps the meta-learning capabilities of deep learning models is a function  
 206 of the model size. If this were true, then we expected to see the difference in meta-test accuracy

Meta-train Initialization	Adaptation at Inference	Meta-test Accuracy
Random	no adaptation	19.3 $\pm$ 0.80
MAML0	no adaptation	20.0 $\pm$ 0.00
USL	no adaptation	15.0 $\pm$ 0.26
Random	MAML5 adaptation	34.2 $\pm$ 1.16
<b>MAML5</b>	<b>MAML5 adaptation</b>	<b>62.4 <math>\pm</math> 1.64</b>
USL	MAML5 adaptation	25.1 $\pm$ 0.98
Random	MAML10 adaptation	34.1 $\pm$ 1.23
<b>MAML5</b>	<b>MAML10 adaptation</b>	<b>62.3 <math>\pm</math> 1.50</b>
USL	MAML10 adaptation	25.1 $\pm$ 0.97
Random	Adapt Head only (with LR)	40.2 $\pm$ 1.30
MAML5	Adapt Head only (with LR)	59.7 $\pm$ 1.37
<b>USL</b>	<b>Adapt Head only (with LR)</b>	<b>60.1 <math>\pm</math> 1.37</b>

Table 2: **MAML trained representations and supervised trained representation have statistically equivalent meta-test accuracy on MiniImagenet – which has low diversity.** The transfer model’s adaptation is labeled as “Adapted Head only (with LR)” – which stands for “Logistic Regression (LR)” used in (19). More precisely, we used Logistic Regression (LR) with LBFGS with the default value for the l2 regularization parameter given by Python’s Sklearn. Note that an increase in inner steps from 5 to 10 with the MAML5 trained model does not provide an additional meta-test accuracy boost, consistent with previous work (26). Note that the fact that the MAML5 representation matches the USL representation when both use the same adaptation method is not surprising – given that: 1) previous work has shown that the distance between the body of an adapted MAML model is minimal compared to the unadapted MAML (which we reproduce in 5 in the green line) and 2) the fact that a MAML5 adaptation is only 5 steps of MAML while LR fully converges the prediction layer. We want to highlight that only the MAML5 model achieved the maximum meta-test performance of 0.6 with the MAML5 adaptation – suggesting that the USL and MAML5 meta-learning algorithms might learn different representations. For USL to have a fair comparison during meta-test time when using the MAML adaptation, we provide the MAML final layer learned initialization parameters to the USL model (but any is fine due to convexity when using a fixed feature extractor). This is needed since during meta-training USL is trained with a union of all the labels (64) – so it does not even have the right output size of 5 for few-shot prediction. Meta-testing was done in the standard 5-way, 5-shot regime.

207 of MAML and USL to be larger for smaller models and the difference to decrease as the model  
208 size increased. Once the two models were, of the same size but large enough, we hypothesized that  
209 the meta-test accuracy would be the same. We tested this to rule out that our observations were a  
210 consequence of the model size. The results were negative and surprisingly the equivalence between  
211 MAML and USL seems to hold even as the model increased – strengthening our hypothesis that the  
212 low task diversity might be a bigger factor explaining our observations. We show this in figure 3  
213 and we want to draw attention to the fact this statistical equivalence holds even when using only four  
214 filters – the case where we expected the biggest difference.

#### 215 4.4 MAML learns a different base model compared to Union Supervised Learned models – 216 even in the presence of low task diversity

217 The first four layers of figure 2 shows how large the distance is of a MAML representation compared  
218 to a SL representation. In particular, it is much larger than the distance value in the range  $[0, 0.1]$   
219 from previous work that compared MAML vs. adapted MAML (18). We reproduced that and indeed  
220 MAML vs. adapted MAML has a small difference (smaller for us) – supporting our observations that  
221 a MAML vs. a USL learned representations are different at the feature extractor layer even when the  
222 diversity is low. Results are statistically significant.

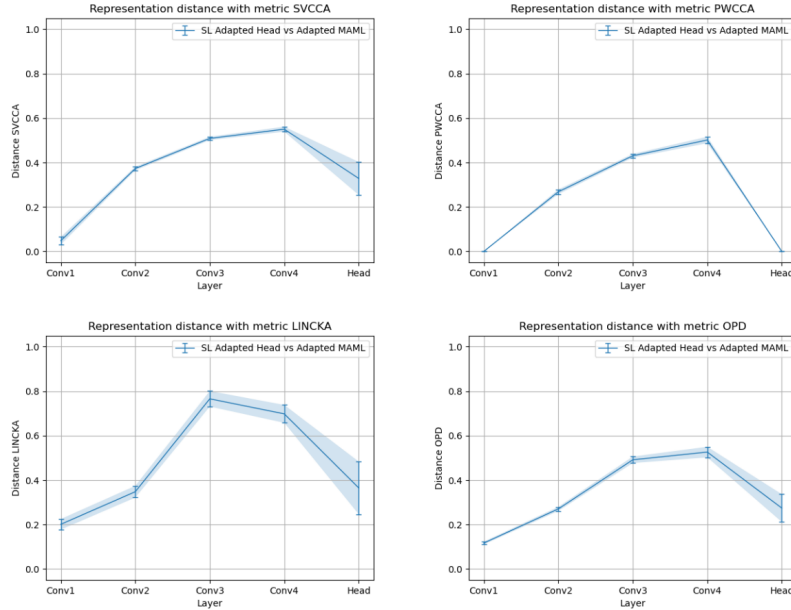


Figure 2: **The classification layer of transfer learning and a MAML5 model decrease in distance – implying similar predictions.** More precisely, an initialization trained with 5 inner steps (MAML5) has an increasingly similar head (classifier) after adaptation with MAML5 compared to the classifier layer of the union supervise learned (USL) model that has been adapted only at the final layer. In particular, the USL model has been adapted with Logistic Regression (LR) with LBFGS with the default value for the  $l_2$  regularization parameter given by Python’s Sklearn (as in (19)). We showed this trend with four different distance metrics SVCCA, PWCCA, LICKA, and OPD referenced in section 2. Observe that according to PWCCA the distance between the predictions is zero. This is true because the distance of classification layer (indicated as “head” in the figure) is zero. The architecture used here is a five layer CNN as in (15; 25) with their same setup. The benchmark used for this analysis is MiniImagenet.




Figure 3: **The meta-test accuracy of MAML and transfer learning using USL is similar in a statistically significant way – regardless of the model size.** In this experiment, we used the MiniImagenet benchmark, the five layer CNN used in (15; 25), and only increased the filter size using sizes 4, 8, 16, and 32. We made sure the comparison was fair by using the same architecture, optimizer, and trained all models to convergence. During meta-training, the MAML model was trained using 5 inner steps. The legends indicating MAML5 and MAML10 refer to the number of inner steps used at test time. We used a (meta) batch size of 100 tasks.

223 **4.5 Synthetic Experiments showing closeness of MAML and Transfer Learning as Diversity**  
 224 **Changes**

225 In this section, we show the closeness of MAML and transfer learning (with USL) for synthetic  
 226 experiments for low and high diversity regimes in Figure 4. In the low regime, the two methods are  
 227 equivalent in a statistically significant way – which supports the main claims of our paper. As the  
 228 diversity increases, however, the difference between USL and MAML increases (in favor of USL).  
 229 This will be explored further in future work.

230 The task is the usual n-way, k-shot tasks, but the data comes from a Gaussian and the meta-learners  
 231 are tasked with classifying from which Gaussian the data points came from in a few-shot learning  
 232 manner. Benchmarks are created by sampling a Gaussian distribution with means moving away from  
 233 the origin as the benchmark changes. Therefore, the Gaussian benchmark with the highest diversity  
 234 coefficient has Gaussians that are the furthest from the origin. We computed the diversity coefficient  
 235 using a proper distance between distributions using the Hellinger distance eluded in section 3 instead  
 236 of the FIM distance. We can do this because we know the ground truth distribution in our synthetic  
 237 experiments, and Gaussians have a closed form Hellinger distance. Details on the n-way Gaussian  
 238 benchmark and diversity coefficient using the Hellinger distance can be found in supplementary  
 239 section E and F.




Figure 4: **The meta-test accuracy of MAML and transfer learning using USL is similar in a statistically equivalent way in the low diversity regime in the 5-way, 10-shot Gaussian Benchmarks.** MAML models were trained with 5 inner steps. MAML5 and MAML10 indicate the adaptation procedure at test time. Results used a (meta) batch-size of 500 tasks and 95% confidence intervals. As the diversity of the benchmark increases, the Gaussian tasks are sampled further away from the origin. Note, as the diversity increases, the difference between USL and MAML increases (in favor of USL).

240 **5 Related Work**

241 Our work proposes a problem-centric framework for the analysis of meta-learning algorithms inspired  
 242 from previous puzzling results (19). We propose to use a pair-wise distance between tasks and  
 243 analyze how this metric might correlate with meta-learning. The closest line of work for this is the  
 244 long line of work by (20) where they suggest methods to analyze the complexity of a task, propose  
 245 unsymmetrical distance metrics for data sets, reachability of tasks with SGD, ways to embed entire  
 246 data sets and more (20; 29; 30; 31). We believe this line of work to be very fruitful and hope that  
 247 more people adopt tools like the ones they suggest and we propose in this paper before researching or  
 248 deploying meta-learning algorithms. We hope this helps meta-learning methods succeed in practice –  
 249 since cognitive science suggests meta-learning is a powerful method humans use to learn (32). In  
 250 the future, we hope to compare (20)’s distance metrics between tasks with ours to provide a further  
 251 unified understanding of meta-learning and transfer learning. A contrast between their work and ours  
 252 is that we focus our analysis from a meta-learning perspective applied to few-shot learning – while  
 253 their focus is understanding transfer learning methods between data sets.

254 The use of a distance metric in our definition of the diversity coefficient is inspired by the analysis  
 255 done by (18). They showed that MAML functions mainly via feature re-use than by rapid learning i.e.,  
 256 that a model trained with MAML changes very little after the MAML adaptation. The main difference



257 of their work with our is: 1) that we compare MAML trained models against union supervised learned  
258 models (USL) instead of only comparing MAML against adapted MAML, and 2) that we explicitly  
259 analyzed properties of the data sets. In addition, we use a large set of distance metrics for our analysis  
260 including: SVCCA, PWCCA, LINCKA and OPD as proposed by (21; 22; 23; 24).

261 Our work is most influenced by previous work suggesting modern meta-learning requires rethinking  
262 (19). The main difference of our work with theirs is that we analyzed the internal representation  
263 of the meta-learning algorithms and contextualize these with quantifiable metrics of the problem  
264 being solved. Unlike their work, we focused on a fair comparison between meta-learning methods by  
265 ensuring the same neural network backbone was used. Another difference is that they gained further  
266 accuracy gains by using distillation – a method we did not analyze and leave for future work.

267 A related line of work (33; 26) first showed that there exist synthetic data sets that are capable of  
268 exhibiting higher degrees of adaptation as compared to the original work by (18). The difference is  
269 that they did not compare MAML models against transfer learning methods like we did here. Instead,  
270 they focused on comparing adapted MAML models vs. unadapted MAML models.

271 Another related line of work is the predictability of adversarial transferability and transfer learning.  
272 They show this both theoretically and with extensive experiments (34). The main difference between  
273 their work and ours is that they focus their analysis mainly on transfer learning, while we concentrated  
274 on meta-learning for few-shot learning. In addition, we did not consider adversarial transferability –  
275 while that was a central piece of their analysis. Further, related work is outlined in the supplementary  
276 section II.

## 277 6 Discussion and Future Work

278 In this work, we presented a problem-centric framework when comparing transfer learning methods  
279 with meta-learning algorithms – using USL and MAML as the canonical representatives of transfer  
280 and meta-learning methods respectively. We showed the diversity coefficient of the popular MiniIma-  
281 genet and Cifar-fs benchmark is low and that under a fair comparison – MAML is very similar to  
282 transfer learning with USL. This was also true even when decreasing the model size – removing the  
283 alternative hypothesis that the equivalence of MAML and transfer learning with USL held due to  
284 large models. Instead, this suggests strengthens our hypothesis that the diversity of the data might be  
285 the driving factor. The equivalence of MAML and USL also replicated in our synthetic experiments.  
286 Therefore, we challenge the suggestions from previous work (19) that only a good embedding can  
287 beat more effective than sophisticated meta-learning – especially in the low diversity regime. In  
288 addition, our synthetic experiments show a promising scenario where we can systematically differ-  
289 entiate meta-learning algorithms from transfer learning algorithms – which supports our actionable  
290 suggestion to use the diversity coefficient to effectively study meta-learning and transfer learning  
291 algorithms. We hope to study this in more depth in the future with real and synthetic data.

292 We also have theoretical results from a statistical decision perspective in the supplementary section ??  
293 that inspired this work and suggest that when the distance between tasks is zero – then the predictions  
294 of transfer learning, meta-learning and even a fixed model with no adaptation are all equivalent (with  
295 the l2 loss). The results are theoretically limited because we can only reason when the diversity is  
296 exactly zero, but regardless provided an interesting perspective to study and inspire empirical work.

297 We hope this work inspires the community in meta-learning and machine learning to construct  
298 benchmarks from a problem-centric perspective – that go beyond large scale data sets – using have  
299 quantitative metrics.

## 300 References

- 301 [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet Classification with Deep Convolu-  
302 tional Neural Networks,”
- 303 [2] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” *Proceed-*  
304 *ings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*,  
305 vol. 2016-December, pp. 770–778, dec 2015.
- 306 [3] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional  
307 Transformers for Language Understanding,” *NAACL HLT 2019 - 2019 Conference of the*

- 308 *North American Chapter of the Association for Computational Linguistics: Human Language*  
309 *Technologies - Proceedings of the Conference*, vol. 1, pp. 4171–4186, oct 2018.
- 310 [4] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan,  
311 P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child,  
312 A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray,  
313 B. Chess, J. Clark, C. Berner, S. Mccandlish, A. Radford, I. Sutskever, and D. A. Openai,  
314 “Language Models are Few-Shot Learners,” tech. rep., 2020.
- 315 [5] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser,  
316 I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalch-  
317 brenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis,  
318 “Mastering the game of Go with deep neural networks and tree search,” *Nature 2016 529:7587*,  
319 vol. 529, pp. 484–489, jan 2016.
- 320 [6] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller,  
321 “Playing Atari with Deep Reinforcement Learning,”
- 322 [7] W. Ye, S. Liu, T. Kurutach, P. Abbeel, Y. Gao, T. University, U. C. Berkeley, S. Qi, and  
323 Z. Institute, “Mastering Atari Games with Limited Data,” oct 2021.
- 324 [8] Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, and P. Isola, “Rethinking Few-Shot Image  
325 Classification: a Good Embedding Is All You Need?,” 2020.
- 326 [9] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang, “A Closer Look at Few-shot  
327 Classification,” *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- 328 [10] Y. Chen, X. Wang, Z. Liu, H. Xu, and T. Darrell, “A New Meta-Baseline for Few-Shot Learning,”  
329 tech. rep.
- 330 [11] G. S. Dhillon, P. Chaudhari, A. Ravichandran, and S. Soatto, “A Baseline for Few-Shot Image  
331 Classification,” 2019.
- 332 [12] S. Huang and D. Tao, “All you need is a good representation: A multi-level and classifier-centric  
333 representation for few-shot learning,” 2019.
- 334 [13] T. M. Mitchell, “The Need for Biases in Learning Generalizations by The Need for Biases in  
335 Learning Generalizations,” 1980.
- 336 [14] S. B.-D. Shai Shalev-Shwartz, “Understanding Machine Learning: From Theory to Algorithms,”  
337 *Cambridge University Press*, 2014.
- 338 [15] C. Finn, P. Abbeel, and S. Levine, “Model-Agnostic Meta-Learning for Fast Adaptation of Deep  
339 Networks,” 2017.
- 340 [16] J. B. Hamrick Deepmind and S. M. Deepmind, “LEVELS OF ANALYSIS FOR MACHINE  
341 LEARNING,” 2020.
- 342 [17] D. Marr, “Vision: A Computational Investigation into the Human Representation and Processing  
343 of Visual Information,” *Phenomenology and the Cognitive Sciences*, vol. 8, no. 4, p. 397, 1982.
- 344 [18] A. Raghu, M. Raghu, S. Bengio, and G. Brain, “Rapid Learning or Feature Reuse? Towards  
345 Understanding the Effectiveness of MAML,” tech. rep., 2020.
- 346 [19] Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, and P. Isola, “Rethinking Few-Shot Image  
347 Classification: a Good Embedding Is All You Need?,” 2020.
- 348 [20] A. Achille UCLA, M. Lam AWS, R. Tewari AWS, A. Ravichandran AWS, S. Maji UMass,  
349 S. Soatto UCLA, and P. Perona Caltech, “TASK2VEC: Task Embedding for Meta-Learning  
350 Charless Fowlkes UCI and AWS,” tech. rep., 2019.
- 351 [21] M. Raghu, J. Gilmer, J. Yosinski, and J. Sohl-Dickstein, “SVCCA: Singular Vector Canonical  
352 Correlation Analysis for Deep Learning Dynamics and Interpretability,” tech. rep., 2017.

- 353 [22] A. S. Morcos, Deepmind, M. Raghu, S. Bengio, and G. Brain, “Insights on representational  
354 similarity in neural networks with canonical correlation,” tech. rep., 2018.
- 355 [23] S. Kornblith, M. Norouzi, H. Lee, and G. Hinton, “Similarity of Neural Network Representations  
356 Revisited,” tech. rep., may 2019.
- 357 [24] F. Ding, J.-S. Denain, and J. Steinhardt, “Grounding Representation Similarity with Statistical  
358 Testing,” 2021.
- 359 [25] S. Ravi and H. Larochelle, “Optimization as a model for few-shot learning,” tech. rep., 2017.
- 360 [26] B. Miranda, “An empirical study of the properties of meta-learning - presentation,” *Illinois  
361 Digital Environment for Access to Learning and Scholarship (IDEALS)*, dec 2020.
- 362 [27] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein,  
363 J. Bohg, A. Bosselut, E. Brunskill, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. Chatterji,  
364 A. Chen, K. Creel, J. Q. Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon,  
365 J. Etchemendy, K. Ethayarajh, L. Fei-Fei, C. Finn, T. Gale, L. Gillespie, K. Goel, N. Goodman,  
366 S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D. E. Ho, J. Hong, K. Hsu,  
367 J. Huang, T. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani,  
368 O. Khattab, P. W. Kohd, M. Krass, R. Krishna, R. Kuditipudi, A. Kumar, F. Ladhak, M. Lee,  
369 T. Lee, J. Leskovec, I. Levent, X. L. Li, X. Li, T. Ma, A. Malik, C. D. Manning, S. Mirchandani,  
370 E. Mitchell, Z. Munyikwa, S. Nair, A. Narayan, D. Narayanan, B. Newman, A. Nie, J. C.  
371 Niebles, H. Nilforoshan, J. Nyarko, G. Ogut, L. Orr, I. Papadimitriou, J. S. Park, C. Piech,  
372 E. Portelance, C. Potts, A. Raghunathan, R. Reich, H. Ren, F. Rong, Y. Roohani, C. Ruiz,  
373 J. Ryan, C. Ré, D. Sadigh, S. Sagawa, K. Santhanam, A. Shih, K. Srinivasan, A. Tamkin,  
374 R. Taori, A. W. Thomas, F. Tramèr, R. E. Wang, and W. Wang, “On the Opportunities and Risks  
375 of Foundation Models,” aug 2021.
- 376 [28] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell,  
377 P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning Transferable Visual Models From  
378 Natural Language Supervision,” feb 2021.
- 379 [29] “The Dynamic Distance Between Learning Tasks: \* From Kolmogorov Complexity to Transfer  
380 Learning via Quantum Physics and the Information Bottleneck of the Weights of Deep Networks,”  
381 2018.
- 382 [30] A. Achille, G. B. Mbeng, and S. Soatto, “Dynamics and Reachability of Learning Tasks,” 2019.
- 383 [31] A. Achille, G. Paolini, G. Mbeng, and S. Soatto, “The Information Complexity of Learning  
384 Tasks, their Structure and their Distance,” 2020.
- 385 [32] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, “Building Machines That  
386 Learn and Think Like People,” *Behavioral and Brain Sciences*, vol. 40, 2016.
- 387 [33] B. Miranda, “An empirical study of the properties of meta-learning - presentation,” *Illinois  
388 Digital Environment for Access to Learning and Scholarship (IDEALS)*, 2020.
- 389 [34] K. Liang, J. Y. Zhang, B. Wang, Z. Yang, O. Koyejo, and B. Li, “Uncovering the Connections  
390 Between Adversarial Transferability and Knowledge Transferability,” 2021.

## 391 Checklist

392 The checklist follows the references. Please read the checklist guidelines carefully for information on  
393 how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or  
394 **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing  
395 the appropriate section of your paper or providing a brief inline description. For example:

- 396 • Did you include the license to the code and datasets? **[Yes]** See Section ??.
- 397 • Did you include the license to the code and datasets? **[No]** Code and data will be released if  
398 accepted.

399

- Did you include the license to the code and datasets? [N/A]

400

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

401

402

403

1. For all authors...

404

(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]

405

406

(b) Did you describe the limitations of your work? [Yes]

407

(c) Did you discuss any potential negative societal impacts of your work? [No]

408

(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

409

410

2. If you are including theoretical results...

411

(a) Did you state the full set of assumptions of all theoretical results? [Yes]

412

(b) Did you include complete proofs of all theoretical results? [Yes]

413

3. If you ran experiments...

414

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No]

415

416

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]

417

418

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]

419

420

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]

421

422

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

423

(a) If your work uses existing assets, did you cite the creators? [Yes]

424

(b) Did you mention the license of the assets? [No]

425

(c) Did you include any new assets either in the supplemental material or as a URL? [Yes]

426

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

427

428

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes]

429

430

5. If you used crowdsourcing or conducted research with human subjects...

431

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [TODO]

432

433

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [TODO]

434

435

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [TODO]

436

437