

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.

Few-Shot Image Classification Benchmarks are Too Far From Reality: Build Back Better with Semantic Task Sampling

Etienne Bennequin^{1, 2} Myriam Tami¹ Antoine Toubhans² Céline Hudelot¹ ¹CentraleSupélec, Université Paris-Saclay, France

²Sicara, France

{etienneb, antoinet}@sicara.com firstname.name@centralesupelec.fr

Abstract

Every day, a new method is published to tackle Few-Shot Image Classification, showing better and better performances on academic benchmarks. Nevertheless, we observe that these current benchmarks do not accurately represent the real industrial use cases that we encountered. In this work, through both qualitative and quantitative studies, we expose that the widely used benchmark tieredImageNet is strongly biased towards tasks composed of very semantically dissimilar classes e.g. bathtub, cabbage, pizza, schipperke, and cardoon. This makes tieredImageNet (and similar benchmarks) irrelevant to evaluate the ability of a model to solve real-life use cases usually involving more fine-grained classification. We mitigate this bias using semantic information about the classes of tieredImageNet and generate an improved, balanced benchmark. Going further, we also introduce a new benchmark for Few-Shot Image Classification using the Danish Fungi 2020 dataset. This benchmark proposes a wide variety of evaluation tasks with various fine-graininess. Moreover, this benchmark includes many-way tasks (e.g. composed of 100 classes), which is a challenging setting yet very common in industrial applications. Our experiments bring out the correlation between the difficulty of a task and the semantic similarity between its classes, as well as a heavy performance drop of state-ofthe-art methods on many-way few-shot classification, raising questions about the scaling abilities of these methods. We hope that our work will encourage the community to further question the quality of standard evaluation processes and their relevance to real-life applications.

1. Introduction

Few-Shot Learning is the science of learning new concepts with only a few examples. This task is one of the key abilities of humans but one of the major shortcomings of standard deep learning methods [15]. It is also required for the success of many industrial applications of computer vision. Businesses that need automatic image recognition do not necessarily possess hundreds of labeled images for each class. It can be because some (or all) classes are rare, or appeared recently, or because classes change every day. In the last three years, our team was involved in a variety of industrial use cases: retrieving a tool in a merchant's catalog, identifying a floor to facilitate recycling, recognizing dishes in a cafeteria given today's menu, or finding the reference of a printed circuit board. In all of these use cases, some or all classes were represented in our database by only 1 to 5 examples. They also had in common that they consisted in recognizing an object among many classes that were semantically similar to one another.

Sadly, because of this, we could not rely on academic benchmarks to identify the most appropriate methods. The first reason is that, as we show in this paper, standard Few-Shot Image Classification benchmarks generate tasks using uniform random sampling from a wide range of semantically dissimilar or unrelated classes (e.g. tieredImageNet), which leads to evaluating our models mostly on tasks composed of objects that we would never need to distinguish in real-life use cases (see Figure 1a). The second reason is that the Few-Shot Learning community has chosen to formalize the Few-Shot Image classification problem as an accumulation of n-way k-shot classification tasks *i.e.* classifying query images, assuming that they belong to one of nclasses for which we have k labeled examples each. In practice, most works compared their methods on benchmarks for which they fixed n = 5 (sometimes n = 10) and k = 1or $k = 5^1$. To the best of our knowledge, only one method was evaluated with n > 50 [22]. The choice made by the community, while relevant to facilitate experiments in the early stages of Few-Shot Learning research, casts a dark shadow on the robustness of state-of-the-art few-shot learning methods when discriminating between a large number of classes.

How can we improve our current evaluation processes to better fit real-life use cases? In this paper, we bring out limitations of current Few-Shot Classification benchmarks with both quantitative and qualitative studies and propose new

https://paperswithcode.com/task/few-shotimage-classification



(a) Coarsity histogram (top) and an example of task (bottom) of a testbed designed from *tiered*ImageNet with uniform class sampling. This task presents a coarsity of 85.1, which is the median coarsity for this testbed. "We really need a machine that will distinguish bathroom tubs from cabbage, pizzas, cardoon, and some very specific kind of dog!" said no one in the history of humankind.



(b) Coarsity histogram (top) and an example of task (bottom) of our testbed *bettertiered*ImageNet. This task presents a coarsity of 15.8. Tasks with this coarsity never occur in the uniformly sampled testbed, although they are more representative of real few-shot classification use cases.

Figure 1. Comparison, in terms of our coarsity measure defined in Section 4.1, between a testbed designed with uniform class sampling (left) and a testbed designed with semantic awareness (right, ours). Our testbed gives a better representativity to tasks with low coarsity *i.e.* composed of classes semantically relevant to one another.

benchmarks to get past these limitations. More specifically:

- 1. We use the WordNet taxonomy [19] to evaluate *semantic distances* between classes of the popular Few-Shot Classification benchmark *tiered*ImageNet. Based on these semantic distances we put forward the concept of *coarsity* of an image classification task, which quantifies how semantically close are the classes of the task.
- 2. We conduct both quantitative and qualitative studies of the tasks generated from the test set of *tiered*ImageNet *i.e.* the tasks composing the benchmark on which most papers evaluate different methods. We show that this benchmark is heavily biased towards tasks composed of semantically unrelated classes.
- 3. We harness the semantic distances between classes to generate the improved benchmark *bettertiered*ImageNet reestablishing the balance between fine-grained and coarse tasks. We compare state-ofthe-art Few-Shot Classification methods on this new benchmark and bring out the relation between the *coarsity* of a task and its difficulty.
- 4. We put forward the Danish Fungi 2020 dataset [21] for evaluating Few-Shot Classification models. This dataset offers a wide range of fine-grained classes and therefore allows the sampling of tasks that we deem to be more representative of industrial applications

of Few-Shot Learning. We compare state-of-the-art methods on both 5-way and 100-way tasks generated from this dataset. To the best of our knowledge, these are the first published results of few-shot methods on such wide tasks.

All our implementations, datasets and experiments are publicly available². We hope that our work will drive the research community towards a better awareness of the biases in few-shot evaluation processes and that the new benchmarks that we propose to counterbalance some of these biases will find echoes in the community and be further improved in future works.

2. Related Works

Few-Shot Image Classification Few-Shot Classification methods historically use episodic training *i.e.* training instances are not images but *episodes* that mimic evaluation tasks that will be sampled from the test set [11, 23, 26, 29]. Recent works suggest that simple fine-tuning baselines are competitive with sophisticated episodic methods, and focus on developing a good inference method while keeping a classically trained feature extractor [4, 6, 9, 12, 17]. Note that some of those methods also consider a *transductive* few-shot classification problem [4, 9, 12, 17] *i.e.* they

 $^{^{2}\}mbox{https}$: / / github . com / sicara / semantic - task - sampling

use unlabeled information from the whole query set to help classification.

Benchmarks The most simple few-shot classification benchmarks are built respectively on the basic image recognition dataset CIFAR100 [3] and Omniglot [14], which is a dataset of handwritten characters. Rather quickly, miniImageNet [29] became the reference dataset for Few-Shot Learning. It is a small subset of ImageNet [7] (60k images in 100 classes, split across classes as train/val/test). Later, [24] introduced the larger dataset *tiered*ImageNet, also built from ImageNet but with 608 classes, which are split in a way that preserves the super-category structure of the classes. In this work we focus on *tiered*ImageNet because miniImageNet offers only 20 test classes which clearly do not allow the sampling of any fine-grained 5way task³. Note that the classification benchmark CU-Birds 200 [30], which happens to be a fine-grained benchmark, is now also used for Few-Shot Image Classification, especially to study cross-domain robustness [6]. To our knowledge, all benchmarks built on these datasets follow the same process to generate few-shot classification tasks *i.e.* we sample 5 classes⁴ uniformly at random from the test set, then sample a fixed number of images per class for the support set and for the query set. However, more recently, [28] merged 10 computer vision datasets to build a gigantic benchmark for few-shot classification methods: Meta-Dataset. They introduced some randomness in the shape of the tasks (number of ways, shots and queries) and proposed to study the hierarchy of the methods depending on the number of ways and shots. Despite its ability to benchmark methods on incredibly diverse datasets and tasks, Meta-Dataset remains underused by the community compared to a simpler and more lightweight benchmark like *tiered*ImageNet⁵. In this work we mitigate the biases of current benchmarks while avoiding additional engineering challenges that would make it harder to adopt our novel benchmarks.

Fine-graininess in Few-Shot Learning Recent works proposed specific methods for Fine-Grained Few-Shot Image Classification [25, 27, 31, 32]. These methods are typically compared on CU-Birds [30] or FGVC Aircraft [18]. These datasets propose respectively 50 and 25 test classes. In this work, we propose to use Danish Fungi 2020 [21], a fine-grained image classification dataset equipped with 1604 classes, which allowed us to compare methods on tasks composed of a large number of classes. Also note

that among many other contributions, the original paper for Meta-Dataset [28] proposed a small study of the performance of state-of-the-art few-shot classifiers depending on a measure of *task fine-graininess* on ImageNet. However, they failed to highlight any correlation between the fine-graininess and difficulty of a task, leaving it for future works. We claim to be this future work: compared to [28], we use more precise tools to define the fine-graininess of a task and decorrelate the fine-graininess from the shape of the task (*i.e.* the number of ways). Thereby we successfully show the correlation between fine-graininess and difficulty on *tiered*ImageNet (see Section 6.2).

Characterizing classification tasks Some recent works try and find a way to represent classification tasks so that they can be compared with one another [1,20]. Here we focus on characterizing tasks using class semantics. [8] show that on ImageNet, visual and semantic similarities between classes are linked. They measure the semantic similarity with the Jiang & Conrath pseudo-distance [13], which depends on the WordNet Directed Acyclic Graph and on the number of images in each class. Other methods to evaluate the similarity between categories can be found in [2].

3. Problem formalization

Few-Shot Image Classification Consider a training set $\mathcal{D}_{train} = \{(x, y) \mid x \in \chi, y \in \mathcal{C}_{train}\}$ and a test set $\mathcal{D}_{test} = \{(x, y) \mid x \in \chi, y \in \mathcal{C}_{test}\}$ where \mathcal{C}_{train} and \mathcal{C}_{test} are their respective class sets, with $\mathcal{C}_{train} \cap \mathcal{C}_{test} = \emptyset$, and χ is the image space. Few-Shot Classification consists in learning on \mathcal{D}_{train} a classification model that can generalize to the unseen classes in \mathcal{C}_{test} with only a few training examples per class.

Given a subset $\mathbf{C} \subset \mathcal{C}_{test}$, we note $\mathcal{D}_{test|\mathbf{C}} = \{(x, y) \in \mathcal{D}_{test} \mid y \in \mathbf{C}\}$. Then a few-shot classification task \mathbf{T} is defined by a subset of classes $\mathbf{C} \subset \mathcal{C}_{test}$, a small support set $\mathbf{S} \subset \mathcal{D}_{test|\mathbf{C}}$ of labeled examples (with at least one example for each class *i.e.* $\forall c \in \mathbf{C}, \mathbf{S} \cap D_{test|\{c\}} \neq \emptyset$) and a query set $\mathbf{Q} \subset \mathcal{D}_{test|\mathbf{C}}$, such that $\mathbf{S} \cap \mathbf{Q} = \emptyset$ *i.e.* images in the query set cannot appear in the support set. The task's objective is to maximize the model's accuracy on \mathbf{Q} .

Model evaluation We define $\mathsf{E}_{test}(n)$ (and similarly $\mathsf{E}_{train}(n)$) the set of *n*-way classification tasks that can be sampled from \mathcal{D}_{test} *i.e.*

$$\begin{split} \mathsf{E}_{test}(n) &= \{ \mathbf{T} = (\mathbf{S}, \mathbf{Q}) \mid \\ \mathbf{S}, \mathbf{Q} \subset \mathcal{D}_{test|\mathbf{C}}, \\ \mathbf{S} \cap \mathbf{Q} &= \emptyset, \\ \mathbf{C} \subset \mathcal{C}_{test} \text{ and } |\mathbf{C}| = \end{split}$$

In practice, most benchmarks are limited to $E_{test}(5)$ (which we will note E_{test} when there is no ambiguity), and further

n

³The classes are nematode, king crab, golden retriever, malamute, dalmatian, African hunting dog, lion, ant, black-footed ferret, bookshop, crate, cuirass, electric guitar, hourglass, mixing bowl, school bus, scoreboard, theater curtain, vase, and trifle.

⁴Rarely 10, or 20 for Omniglot

⁵21 papers used Meta-Dataset between 2020 and 2021 v.s. 92 for *tiered*ImageNet. Source: *PapersWithCode*



Figure 2. Subgraph of WordNet spanning the 160 classes of *tiered*ImageNet's test set, which are shown in red. The root (in green) corresponds to the concept of "entity". This is a Directed Acyclic Graph. Best viewed in color.

limited to tasks with 1 or 5 support images per class, and 10 query images per class. The number of possible tasks is still untractable (~10¹⁷³ on *tiered*ImageNet). Therefore we evaluate few-shot classification models on a subset $\tilde{E}_{test} =$ { $\mathbf{T} \in E_{test} | \mathbf{T} \sim \mathbf{U}$ } with U the uniform distribution. This design choice gives the same weight to all tasks from E_{test} , regardless of how informative they are on a model's ability to perform on real-life few-shot learning problems. In this work, we study alternative distributions for \tilde{E}_{test} .

Class sampling and instance sampling Here we limit ourselves to the sampling of the classes that constitute each episode. We iterate only on the probability distribution over classes that constitute a task. Once the classes are sampled, we use uniform sampling over all instances for each class to constitute the support and query sets.

4. Building the *better-tiered*Imagenet benchmark

4.1. Measuring task coarsity with WordNet taxonomy

Since *tiered*ImageNet is a subset of ImageNet, its classes are the leaves of a directed acyclic graph which is a subgraph of the WordNet graph [19] (see Figure 2). Using this graph, it is possible to establish a semantic similarity between classes. We use the Jiang & Conrath pseudo-distance between classes [13], which is defined for two classes c_1 and c_2 as:

$$D^{JC}(c_1, c_2) = 2\log|lso(c_1, c_2)| - (\log|c_1| + \log|c_2|)$$
(1)

where |c| is the number of instances of the dataset with class c, and $lso(c_1, c_2)$ is the lowest superordinate, *i.e.* the

most specific common ancestor of c_1 and c_2 in the directed acyclic graph. Our choice of pseudo-distance⁶ was motivated by the results of [8], who showed that this semantic pseudo-distance between classes is correlated with visual similarity between images of these classes on ImageNet. We insist, however, that the definition of the coarsity and the subsequent sampling strategy are agnostic of the choice of the distance, and that other semantic distances may be used in future works.

From this pseudo-distance, we define the *coarsity* κ of a task $\mathbf{T}_{\mathbf{C}}$ constituted of instances from a set of classes \mathbf{C} as the mean square distance between the classes of \mathbf{C} *i.e.*

$$\kappa(\mathbf{T}_{\mathbf{C}}) = \underset{c_i, c_j \in \mathbf{C} \\ c_i \neq c_j}{\operatorname{mean}} D^{JC}(c_i, c_j)^2$$
(2)

This coarsity is an indicator of how semantically close are the classes that constitute a task. As shown in [8], on datasets derived from ImageNet, this measure is closely linked to the visual similarity between items of these classes.

4.2. Generate a more informative benchmark using class semantics

As discussed in Section 3, finding an appropriate subset of E_{test} is a key point to ensure that we provide an accurate evaluation of a model. In the literature, testing tasks are sampled uniformly at random from E_{test} [29]. However, we observed that the resulting testbeds are biased towards tasks with high coarsity *i.e.* composed of classes semantically far from each other with respect to the Jiang & Conrath pseudodistance (see Figure 1a). We argue that in practice, fewshot learning models are often used to distinguish between similar objects rather than distinguishing between objects that have nothing to do with one another (*e.g.* circuit boards from circuit boards, carpets from carpets, or people from people). A testbed presenting this type of bias is therefore irrelevant to evaluate a model's ability to solve this family of problems.

In this work, we define a unique, reproducible set of testing tasks to evaluate all models. This testbed is built with a dual objective:

 We want tasks with a smooth repartition in terms of coarsity to ensure that the testbed also evaluates the ability of a model to distinguish between classes close to each other. Providing a good span of coarsities also allows to compare models on different types of tasks: a model might be better for coarse tasks but not for fine-grained tasks.

⁶We call it a pseudo-distance since it is positive, symmetric and separated. However, due to the weak assumption on the directed acyclic graph, it is possible to find c_1, c_2, c_3 such that $D^{JC}(c_1, c_3) > D^{JC}(c_1, c_2) + D^{JC}(c_2, c_3)$.

• This first objective inherently creates a bias towards classes with many neighbouring classes. However we want our testbed to be balanced, *i.e.* all images must be sampled roughly as many times as the others ⁷.

To achieve these goals, we define a semantic task sampler based on the Jiang & Conrath pseudo-distance. We build an initial potential matrix [16] \mathcal{P}_0 such that $\mathcal{P}_0(i, j) = e^{-\alpha D^{J^C}(c_i, c_j)}$ with $\alpha \in \mathbb{R}_+$ an arbitrary scalar. For the first task, the probability for a pair of classes (c_i, c_j) to be sampled together is proportional to $\mathcal{P}_0(i, j)$. To enforce that the testbed is balanced, once the $t - 1^{\text{th}}$ task is sampled we update the number $occ_t(i)$ of occurrences of class c_i in previous tasks. Then we update the potential matrix to penalize classes with higher values of occ_t :

$$\mathcal{P}_t(i,j) = \mathcal{P}_0(i,j) \times \exp(-\beta \frac{occ_t(i) + occ_t(j)}{\max_k (occ_t(k))})$$
(3)

with $\beta \in \mathbb{R}_+$ an arbitrary scalar. Intuitively, a larger α gives more weight to pairs of semantically close classes, while a larger β forces a stricter balance between classes. The class sampling process is detailed in Algorithm 1⁸.

We then sample instances from these classes uniformly at random. As shown in Figure 1b, our 5000-tasks testbed gives far greater representation to fine-grained tasks compared to a uniformly sampled testbed. Our testbed offers a wide range and balance of task coarsities, allowing to test models on both coarse and fine-grained tasks, while the uniformly sampled testbed only allows the evaluation on coarse tasks. Figure 3 shows that the occurrences-based penalty successfully enforces the balance between classes in our testbed.

5. Fungi: a large fine-grained dataset for Few-Shot Image Classification

5.1. Danish Fungi 2020

Danish Fungi 2020 (DF20) [21] is an image recognition dataset of 295 938 images of fungi distributed in 1604 finegrained classes, with no overlap with ImageNet. The dataset offers insightful metadata, such as the object's geographical location, habitat, and substrate. DF20's classes are equipped with a seven-level hierarchical structure. Note that Meta-Dataset [28] also include a Fungi dataset, from the FGVCx 2018 Fungi classification challenge⁹. This Fungi dataset came from the same source as DF20 but offered fewer images, fewer classes, no metadata, and no taxonomy, which convinced us to push forward DF20. The semantic tree of DF20 is shown in Figure 4.

⁷In the case of *tiered*ImageNet, which presents as many images for each class, this is equivalent to ensuring the balance between classes.

⁸We use the convention $\mathcal{P}(i) = [\mathcal{P}(i,j)]_{j=1}^{|\mathcal{C}_{test}|}$. Also, **occ** represents the whole vector, while **occ**(*i*) represents the element at index *i*.

⁹https://github.com/visipedia/fgvcx_fungi_comp

Algorithm 1 How classes of a task are sampled using the potential matrix, following [16]

Input: potential matrix \mathcal{P}_0 , number of tasks T, number of classes per task n

Output: set of sampled tasks \mathcal{T}

- 1: $\mathcal{T} \longleftarrow \{\}$
- 2: $\mathcal{P} \longleftarrow \mathcal{P}_0$
- 3: occ $\leftarrow [1]_i$
- 4: **for** t < T **do**
- 5: $\mathbf{p} \leftarrow \exp(-\beta \frac{\mathbf{occ}}{\max(\mathbf{occ})})$
- 6: $\mathbf{C} \leftarrow \{c_0\}$ with c_0 sampled according to a distribution of probability proportional to \mathbf{p}
- 7: $\mathbf{p} \leftarrow \mathbf{p} \odot \mathcal{P}_0(c_0)$
- 8: while $|\mathbf{C}| < n$ do
- 9: $\mathbf{C} \leftarrow \mathbf{C} \cup \{c\}$ with *c* sampled according to a distribution of probability proportional to \mathbf{p}

10:
$$\mathbf{p} \longleftarrow \mathbf{p} \odot \mathcal{P}_0(c)$$

11: $\mathcal{T} \longleftarrow \mathcal{T} \cup \{\mathbf{C}\}$
12: $\forall i \in \mathbf{C}, \mathbf{occ}(i) \longleftarrow \mathbf{occ}(i) + 1$



Figure 3. Comparison of class imbalance in testbeds. We show the proportion of tasks containing each one of the 160 classes of *tiered*ImageNet's test set, with uniform task sampling (blue) and our semantic task sampling (red). A perfectly balanced testbed will show a flat line *i.e.* each class is equally represented in the testbed. We see that our sampling does not raise class imbalance compared to uniform sampling.

5.2. DF20 for Few-Shot Image Classification Benchmark

Why do we use it? Following our observations on the high coarsity of tasks sampled from *tiered*ImageNet, we propose to use DF20 as a test set for few-shot image classification models. This dataset allows sampling a wide variety of fine-grained few-shot classification tasks, which are to our experience more representative of real-world applications. It also offers a taxonomy allowing to further study



Figure 4. Semantic tree spanning the 1604 classes of DF20 (in red). The root is shown in green. DF20 contains images of species that belong to three kingdoms: Chromista (purple), Protozoa (cyan), and Fungi (orange). As we can see from the tree, the vast majority of classes are spanned from the Fungi kingdom. Best viewed in color.

the performance of few-shot models depending on the coarsity of the task. Since we consider the whole dataset as a test set, we allow the comparison of methods with parameters optimized on various training sets, provided that they do not overlap with DF20. This brings the few-shot learning methodology closer to the neighboring field of transfer learning, in which training data is not part of the benchmark [10]. This can also be considered as a generalization of the cross-domain few-shot learning setting, in which we train a model on the training set of one benchmark (e.g. miniImageNet) and test it on the testing set of another benchmark (e.g. CUB [6]). Note that in this work, compared methods all share the same backbone with parameters classically trained on ImageNet's training set. This was a convenient baseline, but we insist that this does not make it mandatory for future work to train on ImageNet. In fact, we believe that it is of prior importance to study the effect of the choice of the training data on the few-shot classification model's performance.

How do we use it? We sampled four benchmarks: 5-way 1-shot, 5-way 5-shot, 100-way 1-shot, and 100-way 5-shot. The 5-way settings are very common in the few-shot learning literature [28]. However, to the best of our knowledge, very few works evaluate their method on *wide* (*i.e.* more than 10-way) few-shot classification tasks¹⁰. We believe that such tasks are at least equally interesting as 5-way tasks. We assume that this setting was avoided by

early works because GPU memory limitations made it very hard to use back-propagation on batches mimicking 100way tasks during episodic training. We claim that these constraints do not justify overlooking such an interesting problem. Specifically for DF20, the task of recognizing an image among a wide variety of fungi makes way more sense than recognizing an image among 5 random species¹¹.

6. Experiments on new benchmarks

6.1. Implementation details

We conducted the necessary experiments to bring out the need for novel few-shot classification benchmarks and showcase the limitations of state-of-the-art methods on more challenging settings. We restricted the comparison to methods allowing classical training (*i.e.* non episodic) and to a unique set of hyper-parameters. This comparison will be enriched in future works. All parameters of our experiments can be found on our publicly available code ¹².

tieredImageNet We followed the original split of [24]. All methods tested in our benchmark use a common ResNet12 with parameters trained for 500 epochs with classical cross-entropy among the 351 classes of the train set using stochastic gradient descent with a batch size of 512 and learning rate of 0.1 with a decreasing factor of 0.1 after 350, 450 and 480 epochs. The trained weights are directly downloadable from our code. We built two testbeds with uniform class sampling (1-shot and 5-shot), and two testbeds (1-shot and 5-shot) with semantic task sampling (see Section 4.2) with $\alpha = 0.383$ and $\beta = 100.0$. These hyperparameters were selected to enforce the sampling of tasks with small coarsity while ensuring that all classes were equally represented in the testbed (with a small margin). This selection was monitored with visual observations shown in Figures 1b and 3. We upsampled 10000 tasks, then we removed all duplicate tasks and downsampled them to 5000 tasks. All tasks present 10 queries per class.

Danish Fungi 2020 All methods tested in our benchmark use the built-in ResNet18 from PyTorch with weights trained on ImageNet. Since DF20 is already fine-grained, we built four 5000-tasks testbeds (5-way 1-shot, 5-way 5shot, 100-way 1-shot, and 100-way 5-shot) with uniform class sampling. All tasks present 10 queries per class.

Methods For Finetune [6] we use 10 fine-tuning steps with a learning rate of 10^{-3} . For Transductive Information Maximization (TIM) [4] we use 100 fine-tuning steps

 $^{^{10}\}ensuremath{\mathsf{https:}}\xspace$ / paperswithcode . com / task / few - shot - image-classification

¹¹As of 2020, experts have identified \sim 148 000 species of fungi [5].

¹²https://github.com/sicara/semantic-tasksampling

	1-shot							
	better-tieredImageNet					Uniform testhed		
	Whole testbed	1st. Qtl	2nd. Qtl	3rd Qtl	4th Qtl			
ProtoNet [26]	53.10 ± 0.40	42.01	51.11	56.35	62.88	65.24 ± 0.35		
Finetune [6]	56.60 ± 0.41	44.56	54.29	60.12	67.43	69.96 ± 0.35		
BD-CSPN [17]	$\textbf{59.99} \pm \textbf{0.45}$	46.69	57.50	63.80	71.95	$\textbf{74.55} \pm \textbf{0.37}$		
TIM [4]	$\textbf{59.19} \pm \textbf{0.43}$	46.74	56.76	63.03	70.21	73.09 ± 0.35		
Transductive Finetuning [9]	53.01 ± 0.40	41.98	51.74	56.31	62.82	65.27 ± 0.35		
PT MAP [12]	56.74 ± 0.38	45.81	54.38	59.76	66.96	69.58 ± 0.37		

Table 1. Top-1 accuracy of various few-shot learning methods on 1-shot tasks sampled from the tieredImageNet test set with uniform and semantic sampling strategies. For better-tieredImageNet, we show the average accuracy on the whole testbed, along with the average accuracy on the four quartiles of the testbed, when sorted by coarsity of the task (1st quartile contains the tasks with the smallest coarsity). For the testbed of uniformly sampled tasks, we only show the average accuracy on the whole testbed, as the results on each quartile are very similar to one another. We separate inductive (top) from transductive (bottom) methods. Best method(s) in each column is shown in bold.

	5-shot						
	better-tieredImageNet					Uniform testhed	
	Whole testbed	1st. Qtl	2nd. Qtl	3rd Qtl	4th Qtl		
ProtoNet [26]	70.77 ± 0.37	59.72	68.95	74.05	80.41	82.79 ± 0.25	
Finetune [6]	71.60 ± 0.37	60.63	69.68	74.86	81.26	83.66 ± 0.25	
BD-CSPN [17]	72.50 ± 0.37	61.23	70.47	75.94	82.38	84.70 ± 0.25	
TIM [4]	$\textbf{73.32} \pm \textbf{0.37}$	62.40	71.48	76.53	82.92	$\textbf{85.49} \pm \textbf{0.24}$	
Transductive Finetuning [9]	70.79 ± 0.37	59.73	68.98	74.08	80.43	82.79 ± 0.25	
PT MAP [12]	69.45 ± 0.36	58.97	67.40	72.36	79.09	81.54 ± 0.26	

Table 2. Top-1 accuracy of various few-shot learning methods on 1-shot tasks sampled from the tieredImageNet test set with uniform and semantic sampling strategies. For better-tieredImageNet, we show the average accuracy on the whole testbed, along with the average accuracy on the four quartiles of the testbed, when sorted by coarsity of the task (1st quartile contains the tasks with the smallest coarsity). For the testbed of uniformly sampled tasks, we only show the average accuracy on the whole testbed, as the results on each quartile are very similar to one another. We separate inductive (top) from transductive (bottom) methods. Best method(s) in each column is shown in bold.

with a learning rate of 10^{-3} and put a 0.1 weight on the conditional entropy term of the loss. For Transductive Finetuning [9] we use 25 inference steps with a learning rate of 5×10^{-5} . These hyperparameters were selected to fit the original implementations of these methods. We insist that we didn't put any additional effort into further optimizing any few-shot method on our benchmarks.

6.2. Results

tieredImageNet Results for *tieredImageNet* are shown in Tables 1 and 2. The immediate observation that we can make is that our benchmark *better-tieredImageNet* is much more challenging than uniform task sampling, with a performance drop of 12 to 15% in top-1 accuracy for all settings and methods. For further details, we sorted all 5000 tasks with respect to their coarsity and grouped them into four quartiles. The 1st quartile contains the most fine-grained tasks, and the 4th quartile contains the coarsest tasks. From these results, we can confirm that coarsity is indeed correlated to the difficulty of the task since the performance con-

sistently improves when moving towards coarser tasks. We finally observed that even the 4th quartile seems to be more challenging than the uniform benchmark. This is consistent with the demography of tasks shown in Figure 1, since the tasks constituting the 4th quartile of our testbed show a smaller average coarsity than the uniform testbed.

We observed that transductive methods [4, 9, 12, 17], which use the unlabeled information from the query set, unsurprisingly show the best results on both set-ups but especially in 1-shot classification. The leaderboard seems to be consistent on all quartiles, suggesting that none of these methods are "specialized" towards a particular demographic of tasks.

Danish Fungi 2020 Results for DF20 are shown in Table 3. They show that while being more challenging than *tiered*ImageNet and *better-tiered*ImageNet, our DF20 benchmark still constitutes an achievable task. We also report results showing that all methods struggle in the more

	5-v	vay	100-way					
	1-shot 5-shot		1-:	shot	5-shot			
	Top-1		Top-1	Top-5	Top-1	Top-5		
ProtoNet [26]	37.55 ± 0.25	60.53 ± 0.27	7.81 ± 0.08	20.12 ± 0.13	17.69 ± 0.12	40.20 ± 0.16		
Finetune [6]	47.00 ± 0.30	65.06 ± 0.29	$\textbf{9.70} \pm \textbf{0.09}$	25.94 ± 0.15	$\textbf{19.60} \pm \textbf{0.12}$	$\textbf{44.25} \pm \textbf{0.17}$		
BD-CSPN [17]	47.81 ± 0.33	$\textbf{66.32} \pm \textbf{0.30}$	$\textbf{9.75} \pm \textbf{0.09}$	24.11 ± 0.15	$\textbf{19.52} \pm \textbf{0.13}$	41.79 ± 0.17		
TIM [4]	40.73 ± 0.28	62.89 ± 0.28	8.36 ± 0.09	21.30 ± 0.14	18.53 ± 0.12	41.47 ± 0.17		
Trans. Finetuning [9]	37.54 ± 0.25	60.54 ± 0.27	7.71 ± 0.08	20.13 ± 0.13	17.69 ± 0.12	40.21 ± 0.16		
PT MAP [12]	$\textbf{52.08} \pm \textbf{0.35}$	$\textbf{66.78} \pm \textbf{0.29}$	9.54 ± 0.09	$\textbf{26.37} \pm \textbf{0.15}$	18.50 ± 0.12	43.36 ± 0.16		

Table 3. Accuracy of various few-shot learning methods on DF20. For 100-way tasks, we report both top-1 and top-5 accuracy. For 5-way tasks, we do not report top-5 accuracy as we found that it was always 100%. Best method(s) in each column is shown in bold.



Figure 5. Correlation between the coarsity of the 5000 tasks sampled from DF20 for our 5-way 5-shot testbed and the accuracy of the PT-MAP [12] method. The plot is smoothed using the rolling average over a window of 200 tasks.

challenging problem of 100-way classification, especially in the 1-shot setting (less than 10% top-1 accuracy for the best model, less than 20% in the 5-shot setting). We believe that this should stand as a red flag regarding the ability of state-of-the-art few-shot classification methods to scale to real-life use cases.

To complete the study, we show in Figure 5 the correlation between the coarsity of a task (based on the taxonomy of DF20) and the accuracy of the PT-MAP [12] method. We observe that, as was the case for *better-tiered*ImageNet, the closest the classes sampled from Fungi are from one another, the harder the task composed of these classes.

7. Conclusion

We showed that the widely used *tiered*ImageNet benchmark with uniform sampling of classes led to evaluating few-shot learning models on disproportionately coarse tasks. We used semantic task sampling to generate a more informative testbed from *tiered*ImageNet's test set. We also pushed forward as a new benchmark for Few-shot Image Classification models the Danish Fungi 2020 dataset, which we believe to be an incredibly promising playing field for future research in Few-Shot Learning. Finally, we showed that state-of-the-art methods dramatically fail when confronted with many-way classification (here many is 100) of fine-grained objects. We insist that this setting is **not** farfetched and fits tangible industrial use cases. We believe that these results should push us to take a step back and reassess the way we currently think about Few-Shot Image Classification.

We used to define a few-shot classification task by its number of ways and its number of shots, addressing n-way k-shot classification as an indivisible problem. What we did here can be seen as a novel framework, in which the number of classes is not sufficient to define a task: we need to know what these classes are. In the same fashion, future works may go beyond tasks defined by their number of shots, and consider which images are chosen for the support set.

In this work, we addressed what we believe to be a very limiting bias of current Few-Shot Learning benchmarks *i.e.* a bias towards coarse tasks. We chose to tackle this particular shortcoming because we observed that it was the main difference between academic benchmarks and the industrial applications of Few-Shot Learning that we encountered. However, many more limitations of few-shot learning benchmarks are yet to address: the fixed shape of the tasks, the strict balance in both support and query sets, the empty overlap between large-scale classes (currently only used for base training) and few-shot classes, no prior in the choice of support instances, and many other of which we did not think yet. We believe that addressing these shortcomings must be considered a priority in our field, and we encourage any and all who agree to join us in this effort.

References

 Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charless C Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6430–6439, 2019. **3**

- [2] Guilherme Alves, Miguel Couceiro, and Amedeo Napoli. Sélection de mesures de similarité pour les données catégorielles. In EGC 2020-20ème édition de la conférence Extraction et Gestion des Connaissances, 2020. 3
- [3] Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. arXiv preprint arXiv:1805.08136, 2018. 3
- [4] Malik Boudiaf, Ziko Imtiaz Masud, Jérôme Rony, José Dolz, Pablo Piantanida, and Ismail Ben Ayed. Transductive information maximization for few-shot learning. *arXiv preprint* arXiv:2008.11297, 2020. 2, 6, 7, 8
- [5] M Cheek, E Nic Lughadha, P Kirk, H Lindon, J Carretero, B Looney, B Douglas, D Haelewaters, E Gaya, T Llewellyn, et al. New discoveries: plants and fungi. plants. *People Planet*, 2020. 6
- [6] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019. 2, 3, 6, 7, 8
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009. 3
- [8] Thomas Deselaers and Vittorio Ferrari. Visual and semantic similarity in imagenet. In *CVPR 2011*, pages 1777–1784. IEEE, 2011. 3, 4
- [9] Guneet Singh Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. In *ICLR*, 2020. 2, 7, 8
- [10] Vincent Dumoulin, Neil Houlsby, Utku Evci, Xiaohua Zhai, Ross Goroshin, Sylvain Gelly, and Hugo Larochelle. Comparing transfer and meta learning approaches on a unified few-shot classification benchmark. arXiv preprint arXiv:2104.02638, 2021. 6
- [11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Modelagnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017. 2
- [12] Yuqing Hu, Vincent Gripon, and Stéphane Pateux. Leveraging the feature distribution in transfer-based few-shot learning. In *International Conference on Artificial Neural Networks*, pages 487–499. Springer, 2021. 2, 7, 8
- [13] Jay J Jiang and David W Conrath. Semantic similarity based on corpus statistics and lexical taxonomy. arXiv preprint cmp-lg/9709008, 1997. 3, 4
- [14] Brenden Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua Tenenbaum. One shot learning of simple visual concepts. In *Proceedings of the annual meeting of the cognitive science society*, volume 33, 2011. 3
- [15] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015. 1

- [16] Chenghao Liu, Zhihao Wang, Doyen Sahoo, Yuan Fang, Kun Zhang, and Steven CH Hoi. Adaptive task sampling for meta-learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 752–769. Springer, 2020. 5
- [17] Jinlu Liu, Liang Song, and Yongqiang Qin. Prototype rectification for few-shot learning. In *Computer Vision–ECCV* 2020: 16th European Conference, Glasgow, UK, August 23– 28, 2020, Proceedings, Part I 16, pages 741–756. Springer, 2020. 2, 7, 8
- [18] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
 3
- [19] George A Miller. Wordnet: a lexical database for english. Communications of the ACM, 38(11):39–41, 1995. 2, 4
- [20] Cuong Nguyen, Thanh-Toan Do, and Gustavo Carneiro. Similarity of classification tasks. arXiv preprint arXiv:2101.11201, 2021. 3
- [21] Lukáš Picek, Milan Šulc, Jiří Matas, Thomas S Jeppesen, Jacob Heilmann-Clausen, Thomas Læssøe, and Tobias Frøslev. Danish fungi 2020-not just another image recognition dataset. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1525–1535, 2022. 2, 3, 5
- [22] Tiago Ramalho and Marta Garnelo. Adaptive posterior learning: few-shot learning with a surprise-based memory module. arXiv preprint arXiv:1902.02527, 2019. 1
- [23] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In 5th International Conference on Learning Representations, 2017. 2
- [24] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised fewshot classification. *ICLR*, 2019. 3, 6
- [25] Xiaoqian Ruan, Guosheng Lin, Cheng Long, and Shengli Lu. Few-shot fine-grained classification with spatial attentive comparison. *Knowledge-Based Systems*, 218:106840, 2021. 3
- [26] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. arXiv preprint arXiv:1703.05175, 2017. 2, 7, 8
- [27] Luming Tang, Davis Wertheimer, and Bharath Hariharan. Revisiting pose-normalization for fine-grained few-shot recognition. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 14352– 14361, 2020. 3
- [28] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, et al. Meta-dataset: A dataset of datasets for learning to learn from few examples. *ICLR*, 2020. 3, 5, 6
- [29] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. Advances in neural information processing systems, 29:3630– 3638, 2016. 2, 3, 4
- [30] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical

Report CNS-TR-2010-001, California Institute of Technology, 2010. 3

- [31] Jingyi Xu, Hieu Le, Mingzhen Huang, ShahRukh Athar, and Dimitris Samaras. Variational feature disentangling for fine-grained few-shot classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8812–8821, 2021. 3
- [32] Yaohui Zhu, Chenlong Liu, and Shuqiang Jiang. Multiattention meta learning for few-shot fine-grained image recognition. In *IJCAI*, pages 1090–1096, 2020. 3