

Unlocking the Full Potential of Small Data with Diverse Supervision

Ziqi Pang^{*1} Zhiyuan Hu^{*2} Pavel Tokmakov³

Yu-Xiong Wang⁴ Martial Hebert⁵

¹TuSimple ²UCSD ³Toyota Research Institute ⁴UIUC ⁵CMU

ziqi.pang@tusimple.ai, z8hu@ucsd.edu, pavel.tokmakov@tri.global,

yxw@illinois.edu, hebert@cs.cmu.edu

Abstract

Virtually all of deep learning literature relies on the assumption of large amounts of available training data. Indeed, even the majority of few-shot learning methods rely on a large set of “base classes” for pre-training. This assumption, however, does not always hold. For some tasks, annotating a large number of classes can be infeasible, and even collecting the images themselves can be a challenge in some scenarios. In this paper, we study this problem and call it “Small Data” setting, in contrast to “Big Data.” To unlock the full potential of small data, we propose to augment the models with annotations for other related tasks, thus increasing their generalization abilities. In particular, we use the richly annotated scene parsing dataset ADE20K to construct our realistic Long-tail Recognition with Diverse Supervision (LRDS) benchmark, by splitting the object categories into head and tail based on their distribution. Following the standard few-shot learning protocol, we use the head classes for representation learning and the tail classes for evaluation. Moreover, we further subsample the head categories and images to generate two novel settings which we call “Scarce-Class” and “Scarce-Image,” respectively corresponding to the shortage of training classes and images. Finally, we analyze the effect of applying various additional supervision sources under the proposed settings. Our experiments demonstrate that densely labeling a small set of images can indeed largely remedy the small data constraints. Our code and benchmark are available at <https://github.com/BinahHu/ADE-FewShot>.

1. Introduction

The availability of big data is the cornerstone of deep learning research. From early successes on image classification that were enabled by ImageNet [5], to more recent progress in object detection driven by the large-scale COCO dataset [31], thousands of carefully curated and accurately annotated images are essential to outperform the classical, heuristic-based approaches. Indeed, even the latest unsu-

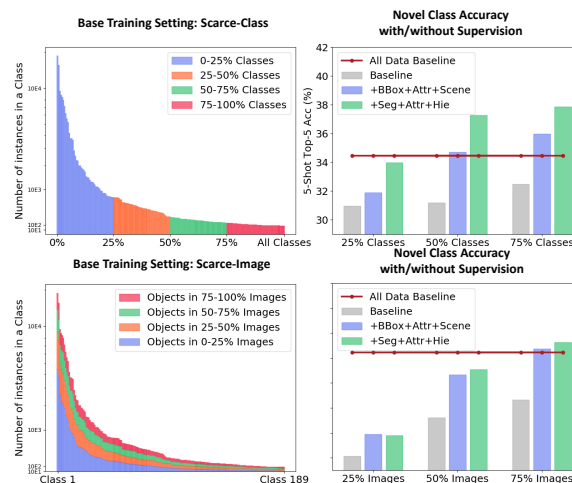


Figure 1. “Small Data” settings and the effect of using diverse supervision. **Left:** We visualize the category-level instance distribution of “Small-Data” regimes. Specifically, Scarce-Class picks the frequent classes and Scarce-Image randomly removes a portion of images. **Right:** Baseline represents the models trained with classification labels only. By incorporating additional supervision, we achieve competitive performance with less data. The types of supervision used: Bounding Box (BBox), Attributes (Attr), Scene Labels (Scene), Segmentation (Seg), and Class Hierarchy (Hie).

pervised learning methods rely on availability of millions of images from the target domain [18]. This requirement, however, cannot always be fulfilled.

First, in many domains the category distribution is highly skewed, with a few categories covering most of the data (so-called head of the distribution), and the rest having only several examples per category (so-called tail) [16, 14]. This setting limits the performance of the methods that rely on a large set of head categories for representation learning. Second, the collection of raw images themselves can be difficult. For instance, in face recognition privacy concerns can limit the data collection [7, 34], and in medical imaging the amount of examples is constrained by the high cost of the image acquisition devices [70, 39].

The problem of learning under data sparsity constraints

has been mostly studied in the few-shot learning setting. Earlier works used toy datasets, like Omniglot [27], or mini-ImageNet [63]. These benchmarks, however, were not realistic both in terms of content, featuring small, object-centric images, and in terms of data distribution, with categories being artificially balanced and randomly split into head and tail. The issue of content realism was partially addressed in the full ImageNet [17], and more recently in MetaDataset [60]. In addition to content realism, the latter benchmark attempted to achieve distribution realism as well by avoiding artificially balancing the categories. Both of these datasets, however, still remain object-centric, and ignore the role of context in recognition. More importantly, by virtue of their size, they have set a new standard of pre-training representations on a large collection of head classes, which, as we have described above, is not always possible in practice.

In this work, we address the limitation of existing few-shot learning benchmarks by introducing a new, realistic dataset – Long-tail Recognition with Diverse Supervision (LRDS). Instead of collecting such a dataset from scratch, we re-purpose the existing ADE20K dataset [71], originally for scene parsing. This dataset was collected by labeling all the objects in a diverse set of images with an open vocabulary, which resulted in a natural long-tail distribution of categories (see Figure 2, left). Moreover, since the objects were labeled in context, we are able to sample them together with a significant portion of the background, removing the object-centric bias of existing benchmarks (see Figure 3).

To better emulate the real-world challenges of learning from small data, we propose two new evaluation settings in LRDS: **Scarce-Class** and **Scarce-Image**, which are illustrated in the left column of Figure 1. The former emulates the heavy skew towards the tail categories by removing the least frequent classes from the head, whereas the latter targets the difficulty of image collection by randomly subsampling images in the training set of the head categories. As can be seen in the right column of Figure 1, both these settings lead to a significant performance degradation on the tail classes with respect to the baseline trained on the full head set, emphasizing the difficulty of the problem.

To unlock the full potential of small data, we explore additional sources of supervision that can be helpful to enrich the representations without the need to collect new images or labels for the main task (such additional supervision can be a lot easier to obtain – one can always annotate bounding boxes in existing images, whereas collecting new images may be infeasible). These experiments are enabled by the wide variety of annotations available in ADE20K. In particular, we study the effect of localization supervision in the form of object masks and bounding boxes, background segmentation, scene-level labels, and object part annotations, as well as semantic supervision in the form of attributes and

class hierarchy (see Figure 2, right). In addition, we show that these forms of supervision can be combined to further improve the performance. As illustrated in the right column of Figure 1, diverse supervision can cover up the degradation from lacking training data to some extent.

To sum up, our contributions are three-fold: (1) We propose a novel benchmark for evaluating long-tail and few-shot learning in a realistic setting. It is based on the ADE20K dataset for scene parsing and features a natural long-tail distribution of object categories; (2) We introduce two new evaluation settings – Scarce-Class and Scarce-Image, which emulate the realistic issues of class and image scarcity; (3) We demonstrate that incorporating diverse supervision can help unlock the full potential of small data by enriching the learned representations, thus increasing their generalization abilities.

2. Related Work

Related visual benchmarks. Our proposed LRDS benchmark aims to facilitate further investigation on few-shot learning (FSL) and long-tail learning (LTL). Most of the existing benchmarks for FSL and LTL are constructed from balanced datasets like ImageNet [5]. Then the data distribution is *manually* modified to simulate the data imbalance and/or scarcity, by splitting the classes into head and tail and subsampling the tail classes. Representative benchmarks include mini-ImageNet [63], tiered-ImageNet [44], and CUB [64] for FSL, Places-LT [32] and ImageNet-LT [32] for LTL, and BSCD-FSL [15], MetaDataset [60], and [61] that involve multiple datasets for cross-domain FSL. Despite the progress made possible by these benchmarks, they are still not practical – their random splitting and subsampling strategies fail to properly reflect how the visual concepts are naturally organized. Also, either common or rare categories have inherent properties (e.g., scale, context, or super-category), but these important properties are lost during such process. Our LRDS instead captures the real-world frequency from ADE20K [71], and is thus more natural and realistic for both FSL and LTL. In addition, no existing benchmark investigates the practical settings with Scarce-Class and Scarce-Image.

Most relevant to our benchmark, iNaturalist [66] is also a realistic long-tail dataset [21]. However, it only contains animals, while our LRDS benchmark has 482 categories covering diverse concepts of living and non-living, objects indoor and outdoor. Moreover, LRDS also exhibits a more severe head/tail imbalance, with the frequencies ranging from 15 to more than 20,000, which poses serious challenge to existing long-tail techniques. LVIS [16] is a recent dataset also capturing a natural long-tail distribution. However, LVIS does not have additional annotation information, and is thus not suitable for our exploration. By contrast, our LRDS contains rich annotations reflecting the complexity

of the visual world.

Multi-task learning (MTL) is also related. To the best of our knowledge, current MTL benchmarks [51, 49, 8] focus on pixel-level RGBD tasks with relatively abundant data. Our benchmark is the first to incorporate the few-shot classification problem into the context of MTL.

Few-shot learning. FSL is a classical problem of recognizing novel objects from few training examples [57, 36, 38, 28, 67, 40, 46]. Some representative approaches include metric learning [24, 63, 50, 55], meta-learning [9, 41, 65], and their combination [60]. However, some most recent work [11, 3, 58] shows that the performance of these complex models can be matched by simple representation learning on base classes and classifier fine-tuning on novel classes. Consistent with and further supportive to this observation, in the paper we mainly focus on the incorporation of supervision with the linear classifier baseline [3, 58], though we also provide results for representative FSL methods on our LRDS benchmark.

Some FSL work improves feature learning with external cues, which is the most relevant to ours. [10, 6, 53] use self-supervised tasks. [66, 30] utilize localization information. Many others explore semantic information: [29] uses the hierarchical structure of concepts, [59, 22] use attributes, and [47, 68, 37] exploit more extensive semantic information. However, ours is the first that systematically investigates extensive and diverse sources of supervision and further addresses their combination, thus pushing the state of the art in this research direction.

Multi-task learning. MTL aims at learning a single model to handle multiple tasks [1]. Most of the work seeks to jointly maximize the performance of each individual task. Existing approaches can be mainly grouped into two categories: balancing the weights of loss functions [48, 23, 4] and introducing better feature sharing and transferring mechanisms [62, 52, 2, 35, 42, 43, 69, 25]. In contrast to their objective of boosting all the involved tasks, we combine multiple sources of supervision to improve the representation learning for classification. In our experiments, we show that a simple sequential and multi-task training approach achieves competitive performance under this objective.

More broadly, our work is related to the general investigation of learning with varying amount of data and annotation, including fully-supervised learning [5, 26], semi-supervised learning, and self-supervised learning. While such work [54, 13] is still in the regime of big data with weak or no annotation, here we in particular focus on the “opposite” direction: “Small Data” with rich supervision.

3. The LRDS Benchmark

In this section, we construct a new benchmark for recognizing rare categories in the wild. The benchmark features a realistic and diverse class distribution, with objects cap-

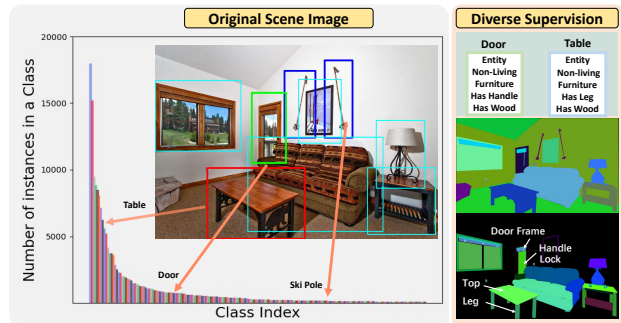


Figure 2. Our LRDS benchmark. **Left:** LRDS is constructed by turning the scene parsing dataset ADE20K into a long-tailed image classification benchmark. **Right:** We present three types of supervision for example: Attributes, Segmentation, and Parts, which are used for learning rich features in the “Small Data” regime.

tured in context. Additionally, we use this benchmark to study Scarce-Class and Scarce-Image regimes, and investigate to what extent can diverse supervision help alleviate the issues caused by data insufficiency. Below, we describe the steps we take to satisfy these objectives.

3.1. Data Selection and Splitting

We choose the ADE20K [71] dataset as the basis for constructing our few-shot learning benchmark, due to its diversity and richness of the vocabulary. ADE20K has more than 3,000 classes, covering objects, object parts, and background. It also exhibits a highly imbalanced category distribution shown in the left column of Figure 2. Below we summarize the filtering steps to select the categories for our benchmark.

First, we manually split the classes into objects, parts, and stuff. Since we want to focus on object classification, only the corresponding categories are included in our label set, resulting in 1,971 classes. The part and stuff labels are instead used as additional sources of supervision.

Second, we further filter the object categories based on their frequency. In particular, we only keep the classes with at least 15 instances in the dataset, resulting in 482 categories. While this excludes the most challenging classes from the benchmark, we argue that including them would introduce significant noise in the evaluation. Indeed, measuring performance on 2-3 images is dominated by noise and is not informative of the model’s recognition ability.

Finally, to mimic the typical setup in the few-shot classification benchmarks, we split the categories into base and novel subsets, where the former is used for representation learning, and the latter for evaluation. Instead of splitting the categories at random, we follow the natural distribution of objects in the world, and select the classes that have more than 100 instances in the dataset as base and the remaining ones as novel. This results in a realistic few-shot learning benchmark to date, where the natural regularities between frequent and infrequent categories are captured.

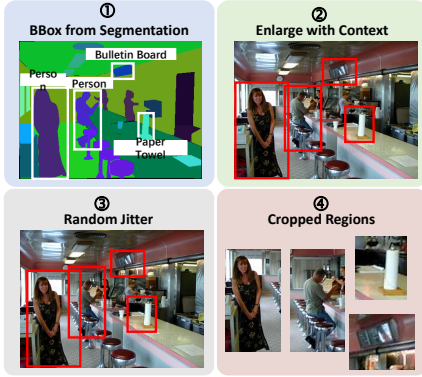


Figure 3. Converting ADE20K into an image classification benchmark. With the original segments annotated in ADE20K, we first crop out every object using tight bounding boxes. Then we imitate realistic data distribution by enlarging the boxes with context and applying random jitter to avoid center bias.

Overall, our dataset contains 189 base and 293 novel categories. For each of the categories in the base set, 1/6 of the data is held out for validation. In the novel set, we randomly select 5 instances in each category for training and the rest are used for evaluation. We further divide the novel set into 100 novel-val and 193 novel-test categories at random, where the former is used for hyper-parameter selection, and the latter for reporting the final performance.

3.2. Adapting ADE20K for Classification

We convert the original ADE20K designed for scene parsing into a classification dataset to serve our needs. Since ADE20K provides instance-level masks for the objects, it can be easily transformed by cropping the regions around the masks and treating them as independent images. However, using tight bounding boxes is an unrealistic data distribution, since objects typically appear in context. Therefore, we simulate a more realistic distribution by box enlargement and random jitter. In particular, we compute the average context ratio (area of context divided by area of tight bounding box) in the ImageNet [5] dataset and expand the original bounding boxes accordingly. We then apply a random shift to the box to avoid center bias (see Figure 3).

3.3. Defining the Small Data Regimes

To experiment under the “Small Data” settings, we propose two new evaluation regimes, Scarce-Class and Scarce-Image, to simulate the corresponding real-world constraints. For the former, we sort the base categories according to the number of instances, and remove the lower 25%, 50%, and 75% for varying degrees of class sparsity. This indeed corresponds to highly-skewed long-tailed distributions in the real world when increasingly fewer categories occupy a larger fraction of the data. For the latter, we simply remove the images from the training set of the base categories at random using the same ratios as for the

Scarce-Class regime. This approach simulates scenarios in which the images are hard to obtain, while preserving the natural, long-tailed category distribution.

Importantly, these two settings have a different effect on the overall size of the training set. Recall that in Scarce-Class the least frequent categories are removed first, which constitute only a small fraction of the dataset (see Figure 1, top left). Removing images, on the other hand, has a more monotonous effect on the dataset size, as shown in the bottom left part of Figure 1. As a result, the two settings have a different effect on the model’s performance, with Scarce-Image being harder in the most constrained scenario, when only 25% of the data is preserved (see Figure 1, right).

3.4. Collecting Diverse Supervision

Evaluating the effect of diverse supervision on the model’s few-shot classification ability is one of the main goals of this work. To this end, we accumulate all the labels provided in ADE20K, which include localization supervision for the object categories (both in the form of masks and bounding boxes), object part annotations, stuff category segmentation, as well as scene labels. In addition, inspired by [59], we collect category-level attribute annotations by inheriting the same attribute set defined on ImageNet [5], and manually assigning attributes for each base category. Finally, we collect class hierarchy labels by extracting the WordNet Tree provided by ADE20K. Notice that *no additional annotations are used for the novel categories*.

Combining all these diverse forms of supervision in a single image classification framework is non-trivial. In the next section we discuss our experimental setup as well as evaluation protocol.

4. Experimental Protocol

The pipeline and setup of our experiments are demonstrated in Figure 4, where the standard few-shot learning protocol in [3] is adopted. We first learn the feature extractor F , such as a ResNet-18 [20], on the training set of base classes. Then F is frozen and evaluated on novel classes for few-shot learning. The accuracy on the test set of novel classes indicates the generalization of the learned feature. The evaluation focuses on both 1-shot and 5-shot settings. Note that, different from [3], we perform classification on the full set of classes (100-way for the validation set and 193-way for the test set).

During both training and evaluation, a model takes as input a region B in a scene image I , and outputs the category C of the object in that area. Therefore, our backbone is identical to the classification branch of Faster R-CNN [45], which performs object classification on the input image regions. The scene image I is propagated through F , producing a feature map. Then the feature map of the region B is fed into an RoI-Align layer [19], whose output is the feature

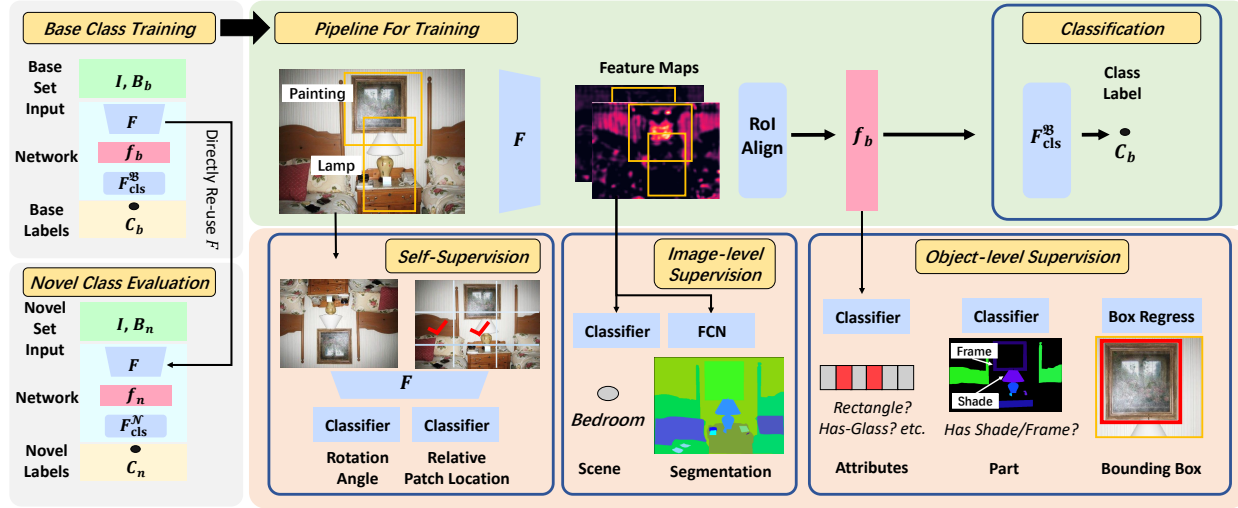


Figure 4. Pipeline for training and evaluation with diverse supervision. **Left:** We learn the feature extractor F on base classes and evaluate it on novel classes for few-shot learning. **Upper-Right:** The model is based on the classification branch of Faster R-CNN [45]. **Lower-Right:** We show some representative examples of introducing supervision at different levels. Self-supervision modifies the input image, image-level supervision operates on the entire feature map, and object-level supervision focuses on the feature vectors of the objects.

vector f of the region. Finally, the model predicts the category C with a classifier, denoted as F_{cls}^B and F_{cls}^N for base and novel classes, respectively. More formally, assume that the image I consists of object set \mathcal{O}_I ; then the classification loss for the base objects is as Eq. (1), where \mathcal{L}_{CE} and RA denote the cross-entropy loss and the RoI-Align operation:

$$\mathcal{L}_{cls}(I) = \sum_{o \in \mathcal{O}_I} \mathcal{L}_{CE}(F_{cls}^B \circ \text{RA}(F(I), B_o), C_o). \quad (1)$$

When combining different types of supervision, each type of supervision s has its own task head F_s , supervision label C_s , and task-specific metric \mathcal{M}_s . Taking the *Scene Label* supervision as an example, we have a classifier, scene label, and cross-entropy loss.

The object-level supervision applies F_s on top of the feature vectors of objects. This set of supervision includes Attributes, Class Hierarchy, Parts, and Bounding Box. For an arbitrary supervision s , the loss on image I is computed as:

$$\mathcal{L}_s(I) = \sum_{o \in \mathcal{O}_I} \mathcal{M}_s(F_s \circ \text{RA}(F(I), B_o), C_{o,s}). \quad (2)$$

For the type of image-level supervision which is relevant to the entire image, such as Scene Label and Segmentation, we apply F_s on top of the feature map $F(I)$. Their loss functions are in the form of

$$\mathcal{L}_s(I) = \mathcal{M}_s(F_s \circ F(I), C_s). \quad (3)$$

In addition, we further consider using pretext tasks as self-supervision [10, 6]. This type of supervision defines an editing operation E on the images, and predicts labels corresponding to E . The losses for self-supervision s are

$$\mathcal{L}_s(I) = \mathcal{M}_s(F_s \circ F \circ E_s(I), C_s). \quad (4)$$

When multiple types of supervision are combined, the final loss function \mathcal{L} becomes the weighted sum of each individual loss term. The combination weights are hyper-parameters that balance different supervision and are selected on the validation set. We elaborate on this and other implementation details in the appendix.

5. Experimental Evaluation

We begin by evaluating state-of-the-art approaches for few-shot learning on our benchmark in Section 5.1. Then we use the full training set of the base classes to study the effect of using additional sources of supervision in Section 5.2 and lowering the cost of supervision in Section 5.3. In advance, we explore the optimal ways of combining them in Section 5.4. Finally, we evaluate the best combinations in the very challenging Scarce-Class and Scarce-Image regimes in Section 5.5 (more analysis is reported in the appendix).

5.1. Benchmarking Few-shot Learning Methods

We begin by providing an evaluation of several recent few-shot learning methods on LRDS. In particular, we focus on the best performing approaches from the recent study of Chen *et al.* [3], and report the results of prototypical networks [50], relational networks [55], and the linear and cosine classifier baselines that have shown promising results in [3]. In addition, we evaluate the very recent ProtoMAML approach [60] which achieves top performance on the realistic MetaDataset. We train all the models under the full data setting and report 5-shot accuracy on the novel set in Table 1.

Similar to [3], we observe that prototypical networks

Model	Top-1	Top-5
Linear Classifier [3]	17.10	34.46
Prototypical Networks [50]	17.14	33.81
Relational Networks [56]	8.34	25.53
Cosine Classifier [3]	13.65	32.33
Proto-MAML [60]	14.54	30.37

Table 1. Benchmarking few-shot learning methods. 5-shot accuracy for models trained with full data is reported.

show strong results; however, the gap between prototypical and relational networks is more significant on our benchmark. We hypothesize that this is due to the fact that relational networks do not generalize well to more realistic data distributions. A similar observation can be made for Proto-MAML. While this complex approach outperforms the linear classifier baseline on MetaDataset, it struggles LRDS, further illustrating the unique challenges presented by our benchmark. Based on these observations, we use linear classifier for the rest of the experiments in the paper.

5.2. Exploration of Individual Supervision Sources

We now explore a diverse set of supervision sources on the proposed LRDS benchmark, while using the full training set of the base classes. In Table 2 we report results for 5 types of supervision, including attributes, class hierarchy, bounding boxes, segmentation, and scene labels. In addition, we report results for a few more sources of supervision, including self-supervised objectives like rotation and relative patch location in the appendix.

Attributes. Following [59], we use a multi-label classification loss for the attributes as shown in Figure 4. The results in the second row of Table 2 demonstrate that this approach indeed improves the performance on the novel classes. In particular, top-5 classification performance in the 5-shot scenario is increased by more than 3%, and the improvements on the base set are a lot less significant, confirming the observations of [59].

Class Hierarchy. We first try to incorporate the hierarchical structure of categories into the feature space. Following [29], we utilize a hierarchical embedding space, transforming the feature representation to different levels of semantic hierarchy to classify corresponding concepts.

We further experiment with a simplified version of class hierarchy supervision, that is, classifying an object on multiple concept levels. For instance, we can classify a cat as `cat`, `mammal`, and `animal`. In implementation, we split the WordNet of ADE20K [71] into four levels and learn four independent classifiers on each of the levels. This naive approach outperforms the hierarchical embedding on LRDS, and we use it for the rest of the experiments.

Bounding Boxes. Unlike [66], we study the effect of providing bounding box labels for the training on base classes. The intuition is that this will allow the object representation to focus on the objects and not on the background, thus

generalizing better to unseen object distributions. To this end, we add a bounding box regression layer identical to R-CNN [12] after the RoI-Align stage. As can be seen from the corresponding row of Table 2, this results in a significant performance improvement on the novel categories, confirming our intuition.

Segmentation. We experiment with two variants of providing segmentation supervision. The first version follows the protocol in Mask R-CNN [19], where we predict a binary mask inside an RoI-Aligned region. The second version operates in the semantic segmentation regime [33], appending an additional convolutional layer after the feature map. Notice that we only apply segmentation supervision to the objects of base classes in the training set and ignore the other pixels. The numbers in Table 2 demonstrate that both approaches result in an improvement over the baseline and the semantic segmentation model achieves top performance. We hypothesize that the problem of semantic labeling of the pixels in a large feature map is more difficult than the alternative formulation of binary pixel classification in a cropped region, providing a stronger regularization signal for representation learning.

Scene Labels. As shown in Figure 4, we supervise scene labels by applying a global pooling layer and linear scene classifier after the feature map of the image. Although scene labels are not directly related to the object categories, we observe a significant improvement in the novel classification performance. We attribute it to the fact that scenes are correlated with certain groups of object categories in the same way as class attributes or elements of the class hierarchy, regularizing learning of object representation.

Overall, the effect of additional supervision sources on the training process varies significantly with their types. We hypothesize that the difficulty of a supervisory task is directly related to how useful it is for learning generalizable representations. For instance, segmentation labels seem to be more useful than bounding boxes, as learning to predict the precise boundaries of the objects is harder than approximately localizing them with a box. The main implication of this observation is that investing in more expensive annotations does pay off when used for representation learning.

5.3. Varying the Amount of Supervision

After observing that providing diverse sources of supervision can bring significant improvements in few-shot learning performance, we want to explore whether these improvements can be obtained at a lower cost. In this section, we explore the effect of the fraction of data for which additional labels are provided on the quality of learned representation. In particular, we experiment with three kinds of supervision that have shown strong improvement: attributes, class hierarchy, and segmentation, varying the proportion of labeled data, and report the results in Figure 5.

Type of supervision	Model	Base-Val	Novel-test set			
			1-shot		5-shot	
			Top-1	Top-5	Top-1	Top-5
	Baseline	44.13	7.00	16.36	17.10	34.46
Attributes	+Attribute	45.38	7.56	17.00	19.66	37.62
Class	+Hierarchy Embedding	44.57	7.48	17.34	19.12	36.93
Hierarchy	+Hierarchy Classifier	46.43	7.83	17.97	19.57	37.19
Bounding Box	+Bounding Box	45.97	7.14	17.16	19.64	37.40
Segmentation	+Segmentation Region	45.68	7.68	17.35	18.95	37.46
	+Segmentation FCN	45.82	7.84	17.53	20.02	38.26
Scene Label	+ Scene	45.03	7.37	18.08	18.26	36.45

Table 2. Comparison of different supervision sources on the base-validation set and novel-test set of LRDS. For clarity, the models are trained with full data.

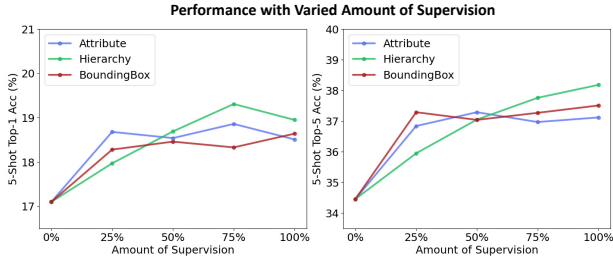


Figure 5. Varying the amount of supervision. The performance increases when additional labels are available even for as little as 25% of the instances.

Model	Method	Top-1	Top-5
Baseline	\	17.10	34.46
+Seg	MTL	18.43	35.49
+Seg	CL	20.02	38.26
+Seg+Attr	MTL	19.59	37.99
+Seg+Attr	CL	20.96	39.41
+Seg+Hier	CL	21.02	39.31
+Seg+BBBox	CL	20.98	39.14
+Seg+Hier+Attr	CL	20.40	39.40
+Seg+Attr+Hier	CL	21.18	39.99

Table 3. Combining multiple types of supervision. We compare both methods for training on multiple tasks (MTL and CL), as well as different task combinations on the novel set of LRDS.

We can observe that for all three types of supervision, labeling as little as 25% of the data already results in significant improvements over the baseline. Indeed, labeling 25% of instances already covers a lot of categories as well as different scenarios, providing a sufficiently strong regularization signal for representation learning. These results demonstrate that the benefits of diverse supervision can be obtained at a relatively low cost. For the rest of the paper, however, we use the full set of of annotations for each source of supervision to demonstrate the full potential of this approach.

5.4. Combining Multiple Sources of Supervision

It is natural to ask whether combining multiple sources of supervision can lead to further improvements. We answer this question in Table 3, and begin by exploring two settings: multi-task learning (MTL), where all types of supervision are applied at once, and curriculum learning (CL), where the additional types of supervision are added one by one (the details of each setting are reported in the appendix). We observe that both for segmentation alone (combined with the main task of object classification), and for a combination of segmentation with attributes, CL achieves a better performance compared to MTL, demonstrating the complexity of multi-task learning. We thus use CL for combining multiple sources of supervision in the rest of the paper.

Next, we notice that combining diverse sources of supervision can indeed have an additive effect on the generalization abilities of the learned representation. For instance, adding both semantic and localization tasks on top of semantic segmentation noticeably improves the performance on the novel set. Moreover, combining three types of supervision results in further improvements, though the returns tend to diminish as more sources of supervision are added.

Finally, the orders in which the sources of supervision are added in our CL framework influence the final performance. In Table 3, two different ways of combining segmentation, attributes, and class hierarchy result in a noticeably different results on the novel set. Overall, we observe that adding stronger forms of supervision first leads to better results in our framework. We hypothesize that this is due to the fact that earlier stages of representation learning are more important, and later fine-tuning stages have only a minor influence. However, optimal rules for selecting this order remain to be thoroughly explored.

5.5. Effect of Diverse Supervision on Small Data

Finally, we answer the question raised at the beginning of this paper: “Can diverse supervision remedy the small data constraints?” To this end, we first establish the baselines

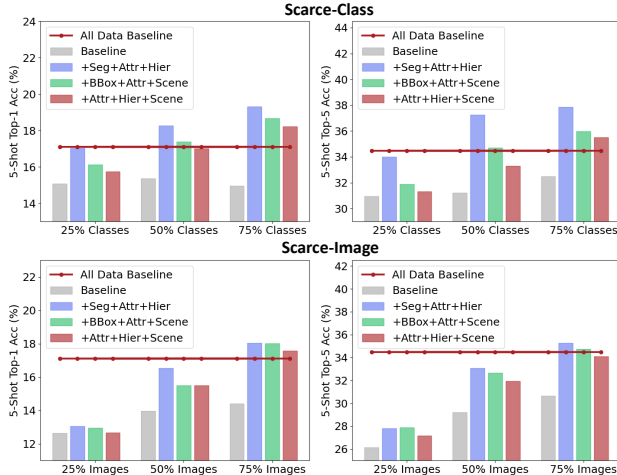


Figure 6. Effect of diverse supervision on small data. We follow the notation introduced in Figure 1. Additional sources of supervision consistently improve the performance under all settings.

for each setting (Scarce-Class and Scarce-Image) and each data ratio (75%, 50%, and 25%) by training on classification labels alone, and then augment those with different combinations of supervisory sources. The particular settings are chosen based on an extensive study of the optimal combinations in the appendix, and include both the most effective but expensive types of supervision (such as bounding box and segmentation) and more cost-effective labels (such as attributes and scene classes). The results are reported in Figure 6.

First, we observe that decreasing both the number of base classes and the number of images results in a significant drop in the performance of the baseline, confirming the complexity of the small data setting studied in this paper. It is worth noting that the drop in performance is not uniform across the settings. Instead, the lack of images has a larger effect on the model’s performance than the lack of classes. As discussed in Section 3.3, this is due to the fact that removing images reduces the actual number of training instances by a large margin, compared to removing the least frequent classes. We quantify this observation in Table 4.

Next, we can see that various combinations of diverse supervision can indeed significantly improve the baseline performance in these challenging scenarios. For instance, in the Scarce-Class setting, our best combination of supervisory sources (Seg + Attr + Hier) outperforms the model trained on the full base training set by using only 50% of the classes. As discussed above, the Scarce-Image setting is more challenging; however, diverse supervision still provides noticeable improvements, nearly reaching the full data performance by only using half of the images.

Overall, using both expensive and cheap sources of supervision results in improvements across the board. That said, we also observe that in the most challenging scenario,

Settings	Remaining Instance Portion
25% Images	25.14%
50% Images	49.79%
75% Images	75.00%
25% Classes	80.87%
50% Classes	92.73%
75% Classes	97.39%

Table 4. The portion of remaining samples under Scarce-Class and Scarce-Image settings. The former retains significantly fewer instances than the latter.

where only 25% of the images are available, all the models struggle to learn a generalizable representation, emphasizing the complexity of the problem. In future work this limitation can be addressed by combining a small number of densely labeled images, with a large collection of unlabeled ones, using recent advances in self-supervised learning [18].

6. Conclusion and Discussion

In this paper we have introduced LRDS – a realistic few-shot learning benchmark with rich annotations. In addition, we have proposed two new evaluation regimes – Scarce-Class and Scarce-Image, which emulate the real-world issues of class and image scarcity, respectively. We have then explored how these challenging scenarios can be remedied by using diverse supervision available in LRDS. Our experiments have shown that a variety of supervision sources, as well as their combinations, are in fact helpful for the task. When combining multiple sources of supervision, those who are most different in nature, such as semantic and localization cues, are most complementary. However, we have also encountered some open problems which we discuss below.

First, the performance quickly saturates as more supervision sources are added. This is a common problem in multi-task learning, and a subject of an active research. We expect that as better approaches to training multi-task models are developed, they will automatically increase the value of diverse supervision for learning generalizable representations in the small data regime.

Second, as is discussed in Section 5.4, the order in which multiple sources of supervision are applied has a noticeable effect on performance. As the number of supervision sources grows, finding the optimal order in a naive way will become prohibitively expensive. Devising principled or heuristic rules to guide the search is thus very important. While we provide some preliminary intuition into this problem (*i.e.*, “harder” tasks need to be learned first), it still remains to be thoroughly explored.

Acknowledgement: This work was supported in part by ONR MURI N000014-16-1-2007 and by AFRL Grant FA23861714660. We also thank NVIDIA for donating GPUs and AWS Cloud Credits for Research program.

References

- [1] Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Multi-task feature learning. In *NeurIPS*, 2007. 3
- [2] Felix JS Bragman, Ryutaro Tanno, Sebastien Ourselin, Daniel C Alexander, and Jorge Cardoso. Stochastic filter groups for multi-task cnns: Learning specialist and generalist convolution kernels. In *ICCV*, 2019. 3
- [3] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *ICLR*, 2019. 3, 4, 5, 6
- [4] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *ICML*, 2018. 3
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 1, 2, 3, 4
- [6] Nikita Dvornik, Cordelia Schmid, and Julien Mairal. Diversity with cooperation: Ensemble methods for few-shot classification. In *ICCV*, 2019. 3, 5
- [7] Zekeriya Erkin, Martin Franz, Jorge Guajardo, Stefan Katzenbeisser, Inald Lagendijk, and Tomas Toft. Privacy-preserving face recognition. In *International symposium on privacy enhancing technologies symposium*, pages 235–253, 2009. 1
- [8] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 3
- [9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *NeurIPS*, 2017. 3
- [10] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Boosting few-shot visual learning with self-supervision. In *ICCV*, 2019. 3, 5
- [11] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018. 3
- [12] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 6
- [13] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *ICCV*, 2019. 3
- [14] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, 2018. 1
- [15] Yunhui Guo, Noel C Codella, Leonid Karlinsky, James V Codella, John R Smith, Kate Saenko, Tajana Rosing, and Rogerio Feris. A broader study of cross-domain few-shot learning. In *ECCV*, 2020. 2
- [16] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 1, 2
- [17] Bharath Hariharan and Ross Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *CVPR*, 2017. 2
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 8
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *ICCV*, 2017. 4, 6
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [21] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alexander Shepard, Hartwig Adam, Pietro Perona, and Serge J. Belongie. The inaturalist species classification and detection dataset. In *CVPR*, 2018. 2
- [22] Siteng Huang, Min Zhang, Yachen Kang, and Donglin Wang. Attributes-guided and pure-visual attention alignment for few-shot recognition. *arXiv preprint arXiv:2009.04724*, 2020. 3
- [23] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, 2018. 3
- [24] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, 2015. 3
- [25] Iasonas Kokkinos. Ubertnet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *ICCV*, 2017. 3
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. 3
- [27] Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015. 2
- [28] Kwonjoon Lee, Subhansu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *CVPR*, 2019. 3
- [29] Aoxue Li, Tiange Luo, Zhiwu Lu, Tao Xiang, and Liwei Wang. Large-scale few-shot learning: Knowledge transfer with class hierarchy. In *CVPR*, 2019. 3, 6
- [30] Jinfu Lin and Xiaojian He. Few-shot learning with weakly-supervised object localization. *arXiv preprint arXiv:2003.00874*, 2020. 3
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 1
- [32] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. In *CVPR*, 2019. 2
- [33] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 6
- [34] Zhuo Ma, Yang Liu, Ximeng Liu, Jianfeng Ma, and Kui Ren. Lightweight privacy-preserving ensemble classifica-

- tion for face recognition. *IEEE Internet of Things Journal*, 6(3):5778–5790, 2019. 1
- [35] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive single-tasking of multiple tasks. In *CVPR*, 2019. 3
- [36] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. In *ICLR*, 2018. 3
- [37] Jesse Mu, Percy Liang, and Noah D. Goodman. Shaping visual representations with language for few-shot classification. In *ACL*, 2020. 3
- [38] Boris N. Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. TADAM: task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018. 3
- [39] Cheng Ouyang, Carlo Biffi, Chen Chen, Turkay Kart, Huaqi Qiu, and Daniel Rueckert. Self-supervision with superpixels: Training few-shot medical image segmentation without annotation. In *ECCV*, 2020. 1
- [40] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L. Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, 2018. 3
- [41] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. 3
- [42] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. In *NeurIPS*, 2017. 3
- [43] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Efficient parametrization of multi-domain deep neural networks. In *CVPR*, 2018. 3
- [44] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676*, 2018. 2
- [45] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1137–1149, 2017. 4, 5
- [46] Andrei A. Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. In *ICLR*, 2019. 3
- [47] Eli Schwartz, Leonid Karlinsky, Rogério Schmidt Feris, Raja Giryes, and Alexander M. Bronstein. Baby steps towards few-shot learning with multiple semantics. *arXiv preprint arXiv:1906.01905*, 2019. 3
- [48] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *NeurIPS*, 2018. 3
- [49] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from RGBD images. In *ECCV*, 2012. 3
- [50] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017. 3, 5, 6
- [51] Trevor Standley, Amir Roshan Zamir, Dawn Chen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *ICML*, 2020. 3
- [52] Gjorgji Strezoski, Nanne van Noord, and Marcel Worring. Many task learning with task routing. In *ICCV*, 2019. 3
- [53] Jong-Chyi Su, Subhransu Maji, and Bharath Hariharan. When does self-supervision improve few-shot learning? In *ECCV*, 2020. 3
- [54] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *ICCV*, 2017. 3
- [55] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H.S. Torr, and Timothy M. Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018. 3, 5
- [56] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018. 6
- [57] Sebastian Thrun. Is learning the n-th thing any easier than learning the first? In *NeurIPS*, 1996. 3
- [58] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B. Tenenbaum, and Phillip Isola. Rethinking few-shot image classification: A good embedding is all you need? In *ECCV*, 2020. 3
- [59] Pavel Tokmakov, Yu-Xiong Wang, and Martial Hebert. Learning compositional representations for few-shot recognition. In *ICCV*, 2019. 3, 4, 6
- [60] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. Meta-dataset: A dataset of datasets for learning to learn from few examples. In *ICLR*, 2020. 2, 3, 5, 6
- [61] Hung-Yu Tseng, Hsin-Ying Lee, Jia-Bin Huang, and Ming-Hsuan Yang. Cross-domain few-shot classification via learned feature-wise transformation. In *ICLR*, 2020. 2
- [62] Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Mti-net: Multi-scale task interaction networks for multi-task learning. In *ECCV*, 2020. 3
- [63] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *NeurIPS*, 2016. 2, 3
- [64] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 2
- [65] Yu-Xiong Wang and Martial Hebert. Learning from small sample sets by combining unsupervised meta-training with CNNs. In *NeurIPS*, 2016. 3
- [66] Davis Wertheimer and Bharath Hariharan. Few-shot learning with localization in realistic settings. In *CVPR*, 2019. 2, 3, 6
- [67] Zhirong Wu, Alexei A. Efros, and Stella X. Yu. Improving generalization via scalable neighborhood component analysis. In *ECCV*, 2018. 3
- [68] Chen Xing, Negar Rostamzadeh, Boris N. Oreshkin, and Pedro O. Pinheiro. Adaptive cross-modal few-shot learning. In *NeurIPS*, 2019. 3
- [69] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *CVPR*, 2018. 3

- [70] Amy Zhao, Guha Balakrishnan, Fredo Durand, John V Guttag, and Adrian V Dalca. Data augmentation using learned transformations for one-shot medical image segmentation. In *CVPR*, 2019. [1](#)
- [71] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ADE20K dataset. *International Journal of Computer Vision*, 127(3):302–321, 2019. [2](#), [3](#), [6](#)