

Meta Omnium: A Benchmark for General-Purpose Learning-to-Learn

Ondrej Bohdal^{1,*}, Yinbing Tian^{2,*}, Yongshuo Zong¹, Ruchika Chavhan¹,
Da Li³, Henry Gouk¹, Li Guo², Timothy Hospedales^{1,3}

¹The University of Edinburgh ²Beijing University of Posts and Telecommunications

³Samsung AI Center, Cambridge *Joint-First Authors

Abstract

Meta-learning and other approaches to few-shot learning are widely studied for image recognition, and are increasingly applied to other vision tasks such as pose estimation and dense prediction. This naturally raises the question of whether there is any few-shot meta-learning algorithm capable of generalizing across these diverse task types? To support the community in answering this question, we introduce Meta Omnium, a dataset-of-datasets spanning multiple vision tasks including recognition, keypoint localization, semantic segmentation and regression. We experiment with popular few-shot meta-learning baselines and analyze their ability to generalize across tasks and to transfer knowledge between them. Meta Omnium enables meta-learning researchers to evaluate model generalization to a much wider array of tasks than previously possible, and provides a single framework for evaluating meta-learners across a wide suite of vision applications in a consistent manner. Code and dataset are available at <https://github.com/edi-meta-learning/meta-omnium>.

1. Introduction

Meta-learning is a long-standing research area that aims to replicate the human ability to learn from a few examples by learning-to-learn from a large number of learning problems [61]. This area has become increasingly important recently, as a paradigm with the potential to break the data bottleneck of traditional supervised learning [26, 70]. While the largest body of work is applied to image recognition, few-shot learning algorithms have now been studied in most corners of computer vision, from semantic segmentation [37] to pose estimation [49] and beyond. Nevertheless, most of these applications of few-shot learning are advancing independently, with increasingly divergent application-specific methods and benchmarks. This makes it hard to evaluate whether few-shot meta-learners can solve diverse vision tasks. Importantly it also discourages the de-

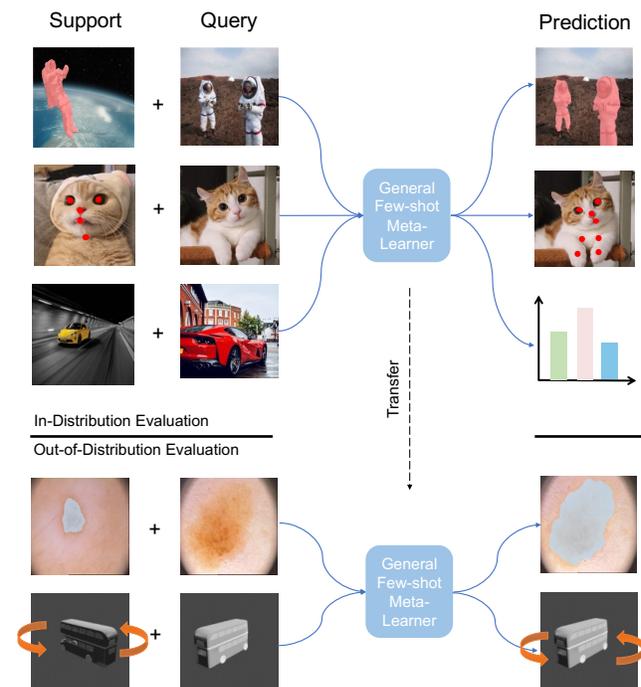


Figure 1. Illustration of the diverse visual domains and task types in Meta Omnium. Meta-learners are required to generalize across multiple task types, multiple datasets, and held-out datasets.

velopment of meta-learners with the ability to learn-to-learn across tasks, transferring knowledge from, e.g., keypoint localization to segmentation – a capability that would be highly valuable for vision systems if achieved.

The overall trend in computer vision [20, 52] and AI [5, 55] more generally is towards more general-purpose models and algorithms that support many tasks and ideally leverage synergies across them. However, it has not yet been possible to explore this trend in meta-learning, due to the lack of few-shot benchmarks spanning multiple tasks. State-of-the-art benchmarks [63, 65] for visual few-shot learning are restricted to image recognition across a handful of visual domains. There is no few-shot benchmark

Dataset	Num Tasks	Num Domains	Num Imgs	Categories	Size	Lightweight	Multi-Task	Multi-Domain
Omniglot [30]	1	1	32K	1623	148MB	✓	✗	✗
miniImageNet [66]	1	1	60K	100	1GB	✓	✗	✗
Meta-Dataset [63]	1	7~10	53M	43~1500	210GB	✗	✗	✓
VTAB [80]	1	3~19	2.2M	2~397	100GB	✗	✗	✓
FSS1000 [37]	1	1	10000	1000	670MB	✓	✗	✗
Meta-Album [65]	1	10~40	1.5M	19~706	15GB	✓	✗	✓
Meta Omnium	4	21	160K	2~706	3.1GB	✓	✓	✓

Table 1. Feature comparison between Meta Omnium and other few-shot meta-learning benchmarks. Meta Omnium uniquely combines a rich set of tasks and visual domains with a lightweight size for accessible use.

that poses the more substantial challenge [57, 77] of generalizing across different *tasks*. We remark that the term *task* is used differently in few-shot meta-learning literature [16, 26, 70] (to mean different image recognition problems, such as cat vs dog or car vs bus) and the multi-task literature [20, 57, 74, 77] (to mean different kinds of image understanding problems, such as classification vs segmentation). In this paper, we will use the term *task* in the multi-task literature sense, and the term *episode* to refer to tasks in the meta-learning literature sense, corresponding to a support and query set.

We introduce Meta Omnium, a dataset-of-datasets spanning multiple vision tasks including recognition, semantic segmentation, keypoint localization/pose estimation, and regression as illustrated in Figure 1. Specifically, Meta Omnium provides the following important contributions: (1) Existing benchmarks only test the ability of meta-learners to learn-to-learn within tasks such as classification [63, 65], or dense prediction [37]. Meta Omnium uniquely tests the ability of meta-learners to learn across multiple task types. (2) Meta Omnium covers multiple visual domains (from natural to medical and industrial images). (3) Meta Omnium provides the ability to thoroughly evaluate both in-distribution and out-of-distribution generalisation. (4) Meta Omnium has a clear hyper-parameter tuning (HPO) and model selection protocol, to facilitate future fair comparison across current and future meta-learning algorithms. (5), Unlike popular predecessors, [63], and despite the diversity of tasks, Meta Omnium has been carefully designed to be of moderate computational cost, making it accessible for research in modestly-resourced universities as well as large institutions. Table 1 compares Meta Omnium to other relevant meta-learning datasets.

We expect Meta Omnium to advance the field by encouraging the development of meta-learning algorithms capable of knowledge transfer across different tasks – as well as across learning episodes within individual tasks as is popularly studied today [16, 70]. In this regard, it provides the next step of the level of a currently topical challenge of dealing with heterogeneity in meta-learning [1, 35, 63, 67]. While existing benchmarks have tested *multi-domain* heterogeneity (e.g., recognition of written characters and plants

within a single network) [63, 65] and shown it to be challenging, Meta Omnium tests *multi-task learning* (e.g., character recognition vs plant segmentation). This is substantially more ambitious when considered from the perspective of common representation learning. For example, a representation tuned for recognition might benefit from rotation *invariance*, while one tuned for segmentation might benefit from rotation *equivariance* [11, 15, 71]. Thus, in contrast to conventional within-task meta-learning benchmarks that have been criticized as relying more on common representation learning than learning-to-learn [53, 62], Meta Omnium better tests the ability of learning-to-learn since the constituent tasks require more diverse representations.

2. Related Work

2.1. Meta-learning Benchmarks

The classic datasets in few-shot meta-learning for computer vision are Omniglot [30] and miniImageNet [66]. Later work criticized these for having insufficient task (episode) diversity and tieredImageNet [56] used the class hierarchy of ImageNet to enforce more diversity between meta-train and meta-test episodes. The main contemporary benchmarks are CD-FSL [22], which challenges few-shot learners to generalize to new visual domains; and Meta-Dataset [63] and Meta-Album [65], which go further in requiring few-shot learners to learn from a mixture of visual domains. Such multi-domain heterogeneous meta-learning turns out to be challenging. A related benchmark to Meta-Dataset is VTAB [80], which similarly provides multiple domains for evaluating data-efficient visual recognition, but their focus is on evaluating representation transfer from large-scale pre-training rather than learning-to-learn and meta-learning. VTAB+MD [13] compare representation transfer and meta-learning approaches on the Meta-Dataset tasks. However, none of these benchmarks address multi-task meta-learning as considered here (Figure 1).

Outside of recognition, task-specific few-shot benchmarks have been proposed in vision problems of semantic segmentation [37], regression [19], pose/keypoint estimation [73], etc. These are mostly slightly behind the complexity of the recognition benchmarks with regards to being

single-domain, with the exception of [73]. With regard to multi-task meta-learning as considered here, the only existing benchmark is meta-world [78], which is specific to robotics and reinforcement learning rather than vision.

We also mention taskonomy [79] as a popular dataset that has been used for multi-task learning. However, it is not widely used for few-shot meta-learning. This is because, although taskonomy has many tasks, unlike the main meta-learning benchmarks [37, 63], there are not enough visual concepts within each task to provide a large number of concepts for meta-training and a disjoint set of concepts to evaluate few-shot learning for meta-validation and meta-testing.

2.2. Heterogeneity in Meta-Learning

There are several sophisticated methods in the literature that highlighted the challenge of the task to address heterogeneity in meta-learning. These have gone under various names such as multi-modal meta-learning [1, 38, 67] – in the sense of multi-modal probability distribution (over tasks/episodes). However, with the exception of [38], these have mostly not been shown to scale to the main multi-modal benchmarks such as Meta-Dataset [63]. A more common approach to achieving high performance across multiple heterogeneous domains such as those in Meta-Dataset is to train an ensemble of feature extractors across available training domains, and fuse them during meta-testing [14, 36, 36]. However, this obviously incurs a substantial additional cost of maintaining a model ensemble. In our evaluation, we focus on the simpler meta-learners that have been shown to work in challenging multi-domain learning scenarios [63, 65], while leaving sophisticated algorithmic and ensemble-based approaches for future researchers to evaluate on the benchmark.

3. Meta Omnium Benchmark and Datasets

3.1. Motivation and Guiding Principles

We first explain the motivating goals and guiding principles behind the design of Meta Omnium. The goal is to build a benchmark for multi-task meta-learning that will: (i) Encourage the community to develop meta-learners that are flexible enough to deal with greater task heterogeneity than before, and thus are more likely to be useful in practice with less curated episode distributions. This was identified as a major challenge in the discussion arising in several recent meta-learning and computer vision workshops and challenges¹. (ii) Ultimately progress on this benchmark should provide practical improvements in data-efficient learning

¹ICLR <https://sites.google.com/view/learning-2-learn/>, ECML <https://janvanrijn.github.io/metalearning/workshop2022>, NeurIPS <https://metalearning.chalearn.org/metadlneurips2021>, and ECCV <http://www.ood-cv.org/>

for computer vision through the development of methods that can better transfer across different task types.

In developing this benchmark, we established a few principles that we used to guide design choices. These include: (i) The benchmark should be lightweight in terms of storage and computing, making it accessible to a broad range of researchers and not only large corporations. (ii) The benchmark should cover multiple tasks with heterogeneous output spaces (as opposed to all classification, all regression, or all dense prediction), as well as multiple visual domains. In these regards, Meta Omnium is compared to alternatives in Table 1. (iii) The initial baselines should have only minimal *task-specific decoders*. This is in contrast to the state of the art within various sub-disciplines of FSL such as segmentation [25, 42], keypoint [39, 72], and classification [2, 75] where specially designed decoders are often used. This is to evaluate and encourage future research on *learning-to-learn* across tasks, rather than primarily benchmarking how well we can manually engineer prior knowledge of optimal task-specific decoders. While we are not opposed to future competitors on this benchmark developing task-specific decoders, these should be evaluated separately against the minimal-decoder competitors. (iv) The benchmark should provide distinct datasets for in-distribution (ID) training and out-of-distribution (OOD) evaluation, to evaluate the robustness of the distribution-shift. This is already provided by [63, 65] for classification, and we extend such an ID and OOD dataset ensemble to multiple tasks. Figure 2 illustrates our dataset and task-split. (v) The benchmark should provide a clear hyper-parameter tuning protocol. With a number of recent studies showing that hyper-parameter tuning can dominate other effects of interest in computer vision [21, 32, 45], this is important for a future-proof meta-learning benchmark. This is also related to the first cost point (i) above: Only for a benchmark with a modest cost can most institutions realistically expect to conduct hyper-parameter tuning. We provide the hyper-parameter tuning protocol. (vi) Finally, following the debate in [35, 53, 62] as to the value of meta-learning vs conventional transfer learning, the dataset should support both episodic meta-learning and conventional transfer learning approaches.

3.2. Data Splits and Tasks

For each main task (classification, segmentation, keypoint localization), we split the datasets into seen datasets available for meta-training, and unseen datasets that are completely held out for out-of-domain meta-validation and meta-testing. Similarly to [63, 65], for the seen datasets, we construct category-wise splits into meta-train/val/test. While for the unseen datasets, there is no category-wise split as episodes from all categories from the whole dataset will be used for validation and testing respectively. The overall split organization is illustrated in Figure 2. We

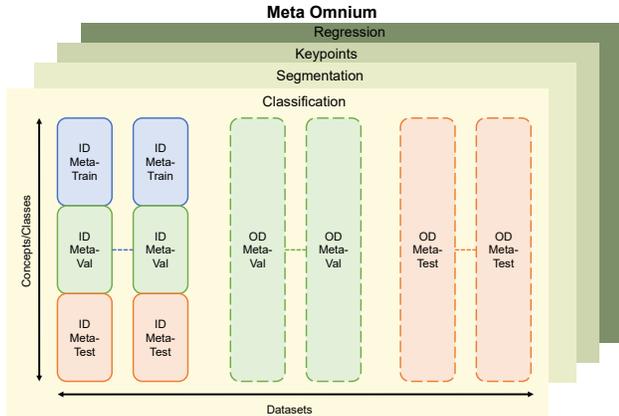


Figure 2. Schematic of benchmark and dataset splits. For each task, there are multiple datasets, which are divided into seen (solid border) and unseen (dash border) datasets. The seen (ID) datasets are divided class-wise into meta-train/meta-val/meta-test splits. The unseen datasets are held out for out-of-distribution (OD) evaluation. Meta-training is conducted on the ID-meta-train split of the seen datasets (blue). Models are validated on ID validation class splits, or OOD validation datasets (green). Results are reported on the ID test class splits and OOD test datasets (orange). We also hold out an entire task family (regression) for evaluating novel task generalisation.

additionally have a completely held-out task: regression. Datasets from this task are not used during meta-training.

Our multi-task setup enables us to define and compare two training protocols: **Single-task meta-learning** which evaluates how well meta-learning performs when trained and tested within a particular task family (within each plate in Figure 2); and **Multi-task meta-learning** which evaluates how well meta-learning performs when trained across all available task families (across all plates in Figure 2).

With this organization we can separately evaluate: **Within-distribution generalization (ID)**: How well do meta-learners generalize to novel test concepts within the seen datasets?; and **Out-distribution generalization (OOD)**: How well do meta-learners generalize to novel concepts in unseen datasets?

We provide two sources of validation data: ID and OOD, and our models are selected based on the combined performance across both. OOD validation is not supposed by the most popular Meta-Dataset benchmark [63] as despite its larger size it does not provide OOD validation datasets.

3.3. Datasets and Metrics

Given the considerations in Section 3.1, our benchmark consists of three main tasks (classification, segmentation, keypoints/pose) and one held-out task (regression).

Classification For classification we take the 10 datasets from the initial public release of Meta-Album [65]. These

images are all 128×128 and contain 19–706 classes per dataset, with 40 images per class. Three of these datasets are reserved for out-of-distribution meta-validation, and four for out-of-distribution meta-test.

Segmentation For segmentation, we take FSS1000 [37] for in-distribution (10,000 images, 1,000 classes), and combine it with VizWiz [64] for OOD meta-validation (862 images, 22 classes), and modified Pascal5i [58] (7,242 images, 6 classes) and the very distinct medical imaging dataset PH2 [41] (200 images, 3 classes) for meta-testing. The segmentation images originally were of diverse sizes. We resize them all to 224×224 for Meta Omnium. Note that VizWiz and Pascal datasets originally contain more classes and images. We exclude the classes that overlap with that in the FSS1000 dataset for few-shot learning, and thus there are no classes overlapping among all the datasets.

Keypoints For keypoints/pose, we take animal-pose [10] for in-distribution, synthetic animal-pose [44] for OOD meta-validation, and MPII human-pose [4] for OOD meta-testing. All images are resized to 128×128 . MPII includes about 40k people in over 25k images with annotated body keypoints. Animal Pose includes 5 animal categories for 6K instances in over 4k images. Each animal is cropped from the original image. We keep cats and dogs for training, horses and sheep for in-domain validation, and cow for in-domain testing. Synthetic animal pose generates synthetic images using animal CAD models rendered from various viewpoints and lightings on a random background. We keep only the horse and tiger categories in our final datasets.

Regression For evaluating regression as a held-out task, we use four datasets corresponding to the test splits of [19]: ShapeNet1D, ShapeNet2D, Distractor and Pascal1D [76]. All images are resized to 128×128 . ShapeNet1D aims to predict azimuth angles. It contains 30 categories in total and we keep the 3 categories from the test set. ShapeNet2D further includes 2D rotation with azimuth angles and elevation. The test set of ShapeNet2D contains 300 categories in total with 30 images per category. Distractor aims to predict the position of a target object in the presence of a distractor. It contains 12 categories in total and the test set has 2 categories. Each category contains 1000 objects with 36 images for each. Pascal1D aims to predict azimuth angle. The whole Pascal1D contains 65 objects from 10 categories. The test set contains 15 objects with 100 images for each object. The supplementary material provides full details of all datasets and splits.

3.4. Training API

For episodic learning, we proceed by (i) sampling a task, (ii) sampling a dataset, (iii) sampling an episode. Under our main protocol we consider variable 1 to 5-shot evaluation, but also evaluate separate 1 and 5 shot settings (training is always done with a variable number of shots – sup-

port examples). For classification tasks, we follow [65] in generating 5-way episodes. For segmentation tasks, we follow [37] in considering each episode to be a binary foreground/background classification problem for a novel class and generate 2-way episodes. For keypoint, we form episodes by randomly selecting a class (e.g., animal category) and then randomly selecting a subset of 5 keypoints to localize for each episode. For regression tasks, we generate variable 5 to 25-shot episodes because it is a common practice to use more shots for regression tasks [19].

For non-episodic/transfer learning, we provide access to the meta-train portions of the seen datasets in conventional mini-batches for conventional single task and multi-task supervised learning.

Evaluation Metrics For classification tasks, we use standard top-1 accuracy; for segmentation tasks, we use standard mean intersection-over-union (mIOU) that averages over IoU values of all object classes; for keypoint prediction, we report the Percentage of Correct Keypoints (PCK). In detail, a detected joint is considered correct if the distance between the predicted and the true joint is within a certain threshold. In our experiments, the threshold is 0.01 for normalized value, which stands for about 12.8 pixel of input image resolution. For regression tasks, we follow [19] and use the same metrics.

3.5. Architecture and Baseline Competitors

As discussed in Section 3.1, we aim to establish baselines that can be adapted to tasks with heterogeneous outputs, with minimal reliance on task-specific decoders. We follow [63, 65] in using a ResNet-18 CNN [23] as a feature extractor architecture. For recognition tasks, we perform multi-class classification immediately after ResNet’s GAP. For regression tasks, we perform linear regression directly after the ResNet’s GAP. For keypoint tasks, we consider them to be a regression problem from the feature map to the keypoint location. For segmentation tasks, we use a simplified PSPNet-like [81] strategy. We concatenate the extracted feature maps from ResNet’s feature pyramid, with upsampling where appropriate, to generate a feature map of size $w \times h$, and then do pixel-wise classification with 1×1 convolutional layer to obtain the final segmentation map. All tasks thus have only *one* learnable weight as a minimal classifier/decoder after the common ResNet feature encoder. Based on this common encoder and minimal decoder architecture, we describe our meta-learning baselines. **Prototypical Network** [59] is a classic meta-learner that exploits nearest-centroid metric learning for few-shot classification. ProtoNets were adapted to segmentation tasks in PANet [68], by performing pixel-level feature matching between support prototypes and query pixels. We use the same principle together with the PSPNet-like features described earlier. To generalize ProtoNets to regression tasks such as

keypoint prediction, we must relax the prototype assumption, and use them as simple Gaussian kernel-regression models [7]. Specifically, we generate a feature embedding for each support example, and then for query examples, we calculate the negative exponential distance to the support examples, and use this inverse distance-weighted sum of support set labels as the prediction. Thus for regression tasks with a support set $S = \{(x_i, y_i)\}$ and query example x_q , ProtoRegression predicts

$$f(y_q|x_q, S) \propto \sum_i y_i \exp(-\tau(f_\theta(x_i) - f_\theta(x_q))^2) \quad (1)$$

enabling us to learn deep feature f_θ in the usual episodic meta-learning way. We use cross-entropy loss for classification and segmentation, and MAE loss for regression tasks.

DDRR Deep differentiable ridge-regression has been considered for few-shot recognition [6], tracking [82], and other tasks. It is related to ProtoNet in that the feature is not adapted after the meta-train stage, but different in that the decoder/classifier layer is learned by differentiable ridge-regression rather than nearest centroid or kernel regression. An elegant property of DDRR methods is that they naturally address regression tasks, although they have been repurposed for classification [6] by conducting MSE-loss regression to a target 1-hot vector. Thus they are a natural choice for our benchmark. For application to segmentation, we apply DDRR in a 1×1 convolution-like way, to perform pixel classification for the output mask with a DDRR classifier at each pixel. Further, we calibrate the prediction for binary cross entropy loss with a learnable scale and bias following [6]. DDRR uses MAE loss for only regression tasks and use MSE loss for other tasks.

MAML The seminal few-shot meta-learner MAML [16] aims to learn an initial condition for per-episode gradient-descent. MAML is straightforward to adapt to different types of tasks. Based on each episode’s support set, a new output layer is learned, and the feature extractor is updated, both by a few steps of gradient descent. Similarly to Meta-Dataset [63], we do not learn an initialization for the output layer, since it can change size between episodes drawn from multiple tasks. To alleviate this challenge, we also follow Meta-Dataset in evaluating Proto-MAML – a variant that initializes the MAML output layer based on the linear classifier/regressor suggested by nearest-centroid prior to gradient descent. Going beyond this, to adapt Proto-MAML to regression tasks, we also initialize the output layer based on the ridge-regression solution to the support set.

Meta-Curvature Meta-Curvature [48] is an enhancement of MAML that learns a pre-conditioning matrix to improve inner-loop adaptation, as well as an initial condition as in standard MAML. Meta-Curvature outperforms MAML in simpler single-task few-shot benchmarks.

Transfer Learning We also consider standard supervised

learning on the meta-train tasks for transfer to the target tasks, a strategy reported to be competitive with meta-learning [62]. For adaptation, we explore both linear read-out [62, 69] and fine-tuning [65]. Besides learning a new output layer from scratch, we also consider a fine-tuning version that initializes the classifier weights using class prototypes (recognition/segmentation) or ridge regression weights (keypoints/regression), inspired by Proto-MAML. **Train-from-Scratch (TFS)** We lastly consider training each episode from scratch using only the support set [65].

3.6. Hyperparameter Optimization

As part of our benchmark, we perform hyperparameter optimization (HPO) to ensure we select appropriate hyperparameters for the diverse tasks that we consider. Multi-objective HPO under restricted resources is challenging, so we devise the following HPO protocol: estimate the performance of each candidate configuration on a lower fidelity (lower number of iterations) and then identify the configuration that works the best across all validation datasets considered (combination of in-domain and out-domain datasets, across various task families). Note that fast multi-fidelity methods such as Hyperband [34], ASHA [33] or PASHA [8] are not applicable out of the box in our multi-objective setup, so we decided to train each candidate configuration using fixed 5,000 training iterations. Since different tasks/datasets are of different difficulties (and use different metrics), we normalize the score of each configuration for each validation dataset by the best score for that dataset across all candidate configurations. We then select the configuration with the best average normalized score.

Note that due to resource constraints we are only able to sample a relatively smaller number of candidates (30), so we utilize a sample efficient state-of-the-art Multi-Objective TPE method [47], available from the Optuna library [3]. We perform HPO for multi-task and single-tasks setups separately, so single task classification, segmentation and keypoint estimation have their own set of hyperparameters; and the multi-task case has its own set. The hyperparameters include the meta-learning rate and optimizer, momentum, and various method-specific hyperparameters (full details are in the supplementary). Once the hyperparameters are chosen, we perform standard training of the model for the full number of iterations.

4. Experiments

In this section, we aim to use our benchmark to answer the following questions: (1) *Which meta-learner performs best on average across a heterogenous range of tasks?* Existing benchmarks have evaluated meta-learners for one task at a time, we now use our common evaluation platform to find out if any meta-learner can provide general-purpose learning to learn across different task types, or whether

each task type prefers a different learner. Similarly, we can ask *which meta-learner is most robust to out-of-distribution tasks?* (2) Having defined the first multi-task meta-learning benchmark, and generalizations of seminal meta-learners to different kinds of output spaces, we ask *which meta-learner performs best for multi-task meta-learning?* More generally, *is there a trend in gradient-based vs metric-based meta-learner success?* (3) *Does multi-task meta-learning improve or worsen performance compared to single-task?* The former obviously provides more meta-training data, which should be advantageous, but the increased heterogeneity across meta-training episodes in the multi-task case also makes it harder to learn [63, 67]. (4) *How does meta-learning perform compared to simple transfer learning, or learning from scratch?*

4.1. Experimental Settings

We train each meta-learner for 30,000 meta-training iterations, with meta-validation after every 2,500 iterations (used for checkpoint selection). For evaluation during meta-testing we use 600 tasks for each corresponding dataset, and for meta-validation we use 1200 tasks together. We use random seeds to ensure that the same tasks are used across all methods that we compare. For transfer learning approaches (fine-tuning, training from scratch, etc.), we use 20 update steps during evaluation. We only retain the meta-learned shared feature extractor across tasks, and for each new evaluation task, we randomly initialize the output layer so that we can support any number of classes as well as novel task families during meta-testing (in line with [65]).

4.2. Results

The main experimental results are shown in Table 2, where rows correspond to different few-shot learners, and columns report the average performance of test episodes, aggregated across multiple datasets in each task family, and broken down by "seen" datasets (ID) and "unseen" datasets (OOD). The table also reports the average rank of each meta-learner across each dataset, both overall and broken down by ID and OOD datasets. More specifically, for each setting (e.g. cls. ID) we calculate the rank of each method (separately for STL and MTL), and then we average those ranks across cls., seg. and keypoints. From the results, we can draw the following main conclusions:

(1) *ProtoNet is the most versatile meta-learner*, as shown by its highest average rank in the single-task scenario. This validates our novel Kernel Regression extension of ProtoNet for tackling regression-type keypoint localization tasks. Somewhat surprisingly, *ProtoNet is also the most robust to out-of-distribution episodes (OOD)* which is different from the conclusion of [63] and others who suggested that gradient-based adaptation is crucial to adapt to OOD data. However, it is also in line with the results of [65] and

		Classification		Segmentation		Keypoints		Average Rank		
		ID	OOD	ID	OOD	ID	OOD	ID	OOD	AVG
Single-Task	MAML	58.7	61.6	54.7	42.1	25.4	33.0	4.3	3.3	3.8
	Proto-MAML	50.5	49.7	46.4	44.1	23.6	22.5	6.0	6.3	6.2
	Meta-Curvature	64.8	61.4	65.6	49.8	43.5	16.0	2.0	4.3	3.2
	ProtoNet	70.4	59.4	75.8	57.2	27.8	33.3	1.3	1.7	1.5
	DDRR	63.1	58.7	66.7	48.0	20.5	31.9	4.7	3.7	4.2
	Proto-FineTuning	50.8	50.7	60.0	43.4	21.3	33.1	5.3	4.3	4.8
	FineTuning	42.3	48.2	50.5	40.0	25.7	30.0	5.7	6.7	6.2
	Linear-Readout	48.6	53.4	34.0	22.7	22.1	26.9	7.3	6.7	7.0
	TFS	31.5	42.0	42.8	37.6	21.0	26.0	8.3	8.0	8.2
Multi-Task	MAML	59.1	58.5	43.3	37.4	24.3	23.9	2.7	4.7	3.7
	Proto-MAML	58.5	63.7	53.0	43.2	21.6	33.3	3.0	1.7	2.3
	Meta-Curvature	70.4	66.9	42.6	34.5	18.2	25.3	4.3	4.7	4.5
	ProtoNet	65.9	58.8	63.3	49.7	20.1	33.0	2.7	2.0	2.3
	DDRR	52.8	51.9	40.4	37.3	22.8	30.1	5.0	4.7	4.8
	Proto-FineTuning	52.4	53.2	44.8	37.8	21.2	30.0	4.3	4.0	4.2
	FineTuning	44.1	51.2	41.3	36.1	18.1	20.5	7.7	7.0	7.3
	Linear-Readout	46.0	50.9	41.5	32.6	19.9	23.5	6.3	8.0	7.2
	TFS	21.9	23.8	38.7	35.8	14.1	11.0	9.0	8.3	8.7

Table 2. Main Results. Results are presented as averages across the datasets within each task type and separately for in-distribution (ID) and out-of-distribution (OOD) datasets. Classification, segmentation, and keypoint results are reported in accuracy (%), mIOU (%), and PCK (%) respectively. The upper and lower groups correspond to multi-task and single-task meta-training prior to evaluation on the same set of meta-testing episodes. Upper and lower sub row groups correspond to meta-learners and non-meta learners respectively. See the appendix for a full breakdown over individual datasets.

the strong performance of prototypes more broadly [9].

(2) Coming to multi-task meta-learning the situation is similar in that ProtoNet dominates the other competitors, but now sharing the first place with Proto-MAML.

(3) To compare single-task and multi-task meta-learning (top and bottom blocks of Table 2) more easily, Figure 3 shows the difference in meta-testing episode performance after STL and MTL meta-training for each method. Overall STL outperforms the MTL condition, showing that the difficulty of learning from heterogeneous tasks [57, 77] outweighs the benefit of the extra available multi-task data.

(4) Finally, comparing meta-learning methods with simple transfer learning methods as discussed in [13, 62], the *best meta-learners are clearly better than transfer learning for both single and multi-task scenarios*.

We also note that Proto-MAML is better than MAML in the multi-task case, likely due to the importance of a good output-layer initialization in the case of heterogeneous episodes, as per [63]. Meta-Curvature outperforms MAML in single-task in-domain scenarios, in line with previous results [48], but it did not achieve stronger performance out-of-domain or in the multi-task case. Finally, while DDRR is perhaps the most elegant baseline in terms of most naturally spanning all task types, its overall performance is middling.

4.3. Additional Analysis

How well can multi-task meta-learners generalize to completely new held-out tasks? We take the multi-task meta-learners (trained on classification, segmentation, keypoints) and evaluate them on four regression benchmarks inspired by [19]: ShapeNet1D, ShapeNet2D, Distractor and Pascal1D. Because the metrics differ across datasets, we analyse the rankings and summarise the results in Table 3. We see the basic TFS performs the worst, with MAML, ProtoNets, DDRR being the best. However, in several cases the results were not better than predicting the mean (full results in the appendix), showing that learning-to-learn of completely new task families is an open challenge.

MAML	PMAML	MC	PN	DDRR	PFT	FT	LR	TFS
3.3	6.5	4.8	3.5	3.5	3.8	5.8	4.3	8.5

Table 3. Average ranking of the different methods across four out-of-task regression datasets.

How much does external pre-training help? Our focus is on assessing the efficacy of meta-learning rather than representation transfer, but we also aim to support researchers investigating the impact of representation learning on external data prior to meta-learning [13, 27, 80]. We therefore

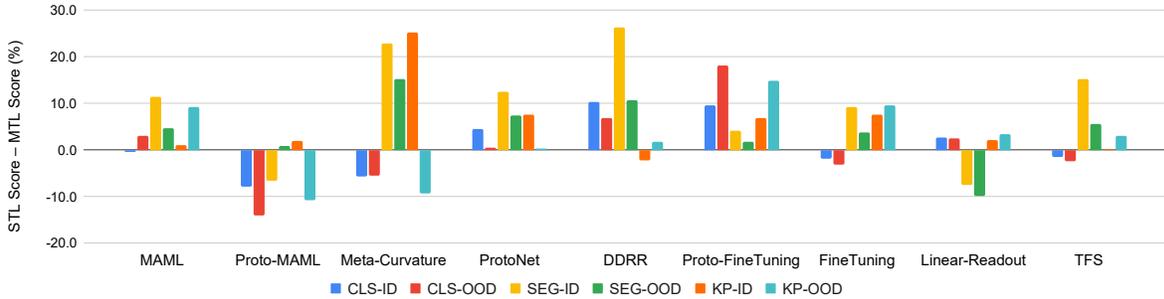


Figure 3. Analysis of the differences in scores between single-task (STL) and multi-task (MTL) learning for different methods.

Method	Pretrain	Cls.		Segm.		Keyp.	
		ID	OOD	ID	OOD	ID	OOD
Proto-MAML	✗	58.5	63.7	53.0	43.2	21.6	33.3
ProtoNet	✗	66.0	58.8	63.3	49.7	20.1	33.0
Proto-MAML	✓	63.9	62.7	56.2	45.3	21.8	33.3
ProtoNet	✓	63.5	58.6	62.0	49.0	20.1	33.1

Table 4. Analysis of the impact of external data pretraining for selected meta-learners in multi-task learning condition. The results show that ImageNet pretraining does not necessarily help improve performance. Cls., Segm., Keyp. represent classification, segmentation and keypoint respectively.

specify evaluation conditions where external data outside our defined in-distribution meta-training set is allowed.

We take two high-performing approaches in the multi-task scenario, Proto-MAML and ProtoNet, and we investigate to what extent external pre-training helps. We use the standard ImageNet1k pre-trained ResNet18, prior to conducting our meta-learning pipeline as usual. We use the same hyperparameters as selected earlier for these models to ensure consistent evaluation, and ensure that the differences in performance are not due to a better selection of hyperparameters. The results in Table 4 show that pretraining is not necessarily helpful in the considered multi-task setting, in contrast to purely recognition-focused evaluations [13, 27, 80], which were unambiguously positive about representation transfer from external data.

Analysis of Runtimes We analyze the times that the different meta-learning approaches spend on meta-training, meta-validation and meta-testing in the multitask learning case of our benchmark. The results in Table 5 show that all experiments are relatively lightweight, despite the ambitious goal of our benchmark to learn a meta-learner that can generalize across various task families. Most notably we observe that ProtoNets are the fastest approach, alongside being the best-performing one. Note that fine-tuning and training from scratch are expensive during the test time as they use backpropagation with a larger number of steps.

Discussion and Future Work In future Meta Omnium can be used in a variety of ways beyond benchmarking

Method	Train Time	Val Time	Test Time	Total Time
MAML	1.8h	1.9h	0.9h	5.0h
Proto-MAML	1.9h	1.9h	0.9h	5.1h
Meta-Curvature	3.4h	2.6h	1.3h	7.6h
ProtoNet	0.8h	0.4h	0.2h	1.8h
DDRR	1.4h	0.6h	0.3h	2.7h
Proto-FineTuning	1.7h	4.5h	2.3h	8.9h
FineTuning	1.5h	8.1h	4.9h	14.9h
Linear-Readout	1.2h	5.1h	2.8h	9.6h
TFS	0.0h	0.8h	6.2h	7.0h

Table 5. Analysis of times needed by different algorithms in the multitask setting (using one NVIDIA 1080 Ti GPU and 4 CPUs).

multi-task meta-learning per-se. These including: studying the multi-task optimisation [77] in meta-learning, studying HPO for meta-learning, developing validation strategies in meta-learning (using ID vs OD val sets [32]), and studying the benefit of task-specific decoders and external data.

5. Conclusion

We have introduced Meta Omnium, the first multi-task few-shot meta-learning benchmark for computer vision. The benchmark is challenging in multiple highly topical ways including requiring learning on heterogeneous task distributions, evaluating generalization to out-of-distribution datasets, and uniquely challenging meta-learners to learn-to-learn and transfer knowledge across tasks with heterogeneous output spaces. Meta Omnium is nevertheless lightweight enough to be of broad interest and use for driving future research, and even to support future research in hyper-parameter optimization for meta-learning.

Acknowledgements This work was supported by the UKRI and EPSRC Grant numbers EP/S000631/1, EP/L016427/1 and EP/S02431X/1; the MOD University Defence Research Collaboration (UDRC) in Signal Processing; the Royal Academy of Engineering under the Research Fellowship programme; Samsung AI Center, Cambridge; China Grant No.2020KJ010802 and No.BYSYZHKC2021115.

References

- [1] Milad Abdollahzadeh, Toubia Malekzadeh, and Ngai-Man Man Cheung. Revisit multimodal meta-learning through the lens of multi-task learning. In *NeurIPS*, 2021. 2, 3
- [2] Arman Afrasiyabi, Hugo Larochelle, Jean-François Lalonde, and Christian Gagné. Matching feature sets for few-shot image classification. In *CVPR*, 2022. 3
- [3] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *KDD*, 2019. 6
- [4] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014. 4, 12
- [5] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatuo Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. 2022. 1
- [6] Luca Bertinetto, João F. Henriques, Philip H.S. Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *ICLR*, 2019. 5
- [7] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006. 5
- [8] Ondrej Bohdal, Lukas Balles, Martin Wistuba, Beyza Ermis, Cedric Archambeau, and Giovanni Zappella. PASHA: Efficient HPO and NAS with progressive resource allocation. In *ICLR*, 2023. 6
- [9] Ondrej Bohdal, Da Li, and Timothy Hospedales. Feed-forward source-free domain adaptation via class prototypes. In *ECCV OOD-CV workshop*, 2022. 7
- [10] Jinkun Cao, Hongyang Tang, Hao-Shu Fang, Xiaoyong Shen, Cewu Lu, and Yu-Wing Tai. Cross-domain adaptation for animal pose estimation. In *ICCV*, 2019. 4, 12
- [11] Ruchika Chavhan, Jan Stuehmer, Calum Heggan, Mehrdad Yaghoobi, and Timothy Hospedales. Amortised invariance learning for contrastive self-supervision. In *ICLR*, 2023. 2
- [12] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 2017. 12
- [13] Vincent Dumoulin, Neil Houlsby, Utku Evci, Xiaohua Zhai, Ross Goroshin, Sylvain Gelly, and Hugo Larochelle. A unified few-shot classification benchmark to compare transfer and meta learning approaches. In *NeurIPS Datasets and Benchmarks*, 2021. 2, 7, 8
- [14] Nikita Dvornik, Cordelia Schmid, and Julien Mairal. Selecting relevant features from a universal representation for few-shot classification. In *ECCV*, 2020. 3
- [15] Linus Ericsson, Henry Gouk, and Timothy M. Hospedales. Why do self-supervised models transfer? Investigating the impact of invariance on downstream tasks. *BMVC*, 2022. 2
- [16] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 2, 5
- [17] Mario Fritz, E. Hayman, B. Caputo, and J. Eklundh. THE KTH-TIPS database. 2004. 12
- [18] Geetharamani G. and Arun Pandian J. Identification of plant leaf diseases using a nine-layer deep convolutional neural network. *Computers and Electrical Engineering*, 76:323–338, 2019. 12
- [19] Ning Gao, Hanna Ziesche, Ngo Anh Vien, Michael Volpp, and Gerhard Neumann. What matters for meta-learning vision regression tasks? In *CVPR*, 2022. 2, 4, 5, 7, 12
- [20] Golnaz Ghiasi, Barret Zoph, Ekin D Cubuk, Quoc V Le, and Tsung-Yi Lin. Multi-task self-training for learning general representations. In *ICCV*, 2021. 1, 2
- [21] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *ICLR*, 2021. 3
- [22] 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
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [24] Emily F. Brownlee Heidi M. Sosik, Emily E. Peacock. Annotated plankton images - data set for developing and evaluating classification methods., 2015. 12
- [25] Sunghwan Hong, Seokju Cho, Jisu Nam, Stephen Lin, and Seungryong Kim. Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In *ECCV*, 2022. 3
- [26] Timothy M Hospedales, Antreas Antoniou, Paul Micaelli, and Amos J. Storkey. Meta-learning in neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2021. 1, 2
- [27] Shell Hu, Da Li, Jan Stuehmer, Minyoung Kim, and Timothy Hospedales. Pushing the limits of simple pipelines for few-shot learning. In *CVPR*, 2022. 7, 8
- [28] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, 2013. 12
- [29] Gustaf Kylberg. The Kylberg texture dataset v. 1.0. Technical report, 2011. 12
- [30] Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 2015. 2
- [31] Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. A sparse texture representation using local affine regions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1265–1278, Aug. 2005. 12
- [32] Da Li, Henry Gouk, and Timothy Hospedales. Finding lost DG: Explaining domain generalization via model complexity. In *arXiv*, 2022. 3, 8
- [33] Liam Li, Kevin Jamieson, Afshin Rostamizadeh, Ekaterina Gonina, Moritz Hardt Hardt, Benjamin Recht, and Ameet Talwalkar. A system for massively parallel hyperparameter tuning. In *MLSys*, 2020. 6
- [34] Lisha Li, Kevin G. Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *JMLR*, 2018. 6

- [35] Rui Li, Ondrej Bohdal, Rajesh K. Mishra, Hyeji Kim, Da Li, Nicholas D. Lane, and Timothy M. Hospedales. A channel coding benchmark for meta-learning. 2021. [2](#), [3](#)
- [36] Wei-Hong Li, Xialei Liu, and Hakan Bilen. Universal representation learning from multiple domains for few-shot classification. In *ICCV*, 2021. [3](#)
- [37] Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *CVPR*, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [12](#)
- [38] Yanbin Liu, Juho Lee, Linchao Zhu, Ling Chen, Humphrey Shi, and Yi Yang. A multi-mode modulator for multi-domain few-shot classification. In *ICCV*, 2021. [3](#)
- [39] Changsheng Lu and Piotr Koniusz. Few-shot keypoint detection with uncertainty learning for unseen species. In *CVPR*, 2022. [3](#)
- [40] P. Mallikarjuna, Alireza Tavakoli Targhi, Mario Fritz, E. Hayman, B. Caputo, and J. Eklundh. The KTH-TIPS 2 database. 2006. [12](#)
- [41] Teresa Mendonça, Pedro M. Ferreira, Jorge S. Marques, André R. S. Marçal, and Jorge Rozeira. PH2 - a dermoscopic image database for research and benchmarking. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2013. [4](#), [12](#)
- [42] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. In *ICCV*, 2021. [3](#)
- [43] Ari Morcos, Maithra Raghu, and Samy Bengio. Insights on representational similarity in neural networks with canonical correlation. *Advances in Neural Information Processing Systems*, 31, 2018. [12](#)
- [44] Jiteng Mu, Weichao Qiu, Gregory D. Hager, and Alan L. Yuille. Learning from synthetic animals. In *CVPR*, 2020. [4](#), [12](#)
- [45] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. Unsupervised domain adaptation: A reality check. In *arXiv*, 2021. [3](#)
- [46] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, 2008. [12](#)
- [47] Yoshihiko Ozaki, Yuki Tanigaki, Shuhei Watanabe, and Masaki Onishi. Multiobjective tree-structured parzen estimator for computationally expensive optimization problems. In *Genetic and Evolutionary Computation Conference*, 2020. [6](#)
- [48] Eunbyung Park and Junier B. Oliva. Meta-Curvature. In *NeurIPS*, 2019. [5](#), [7](#)
- [49] Massimiliano Patacchiola, Jack Turner, Elliot J Crowley, Michael O’Boyle, and Amos J Storkey. Bayesian meta-learning for the few-shot setting via deep kernels. In *NeurIPS*, 2020. [1](#)
- [50] Gerald Piosenka. 100 sports image classification. [12](#)
- [51] Gerald Piosenka. Birds 400 - species image classification. [12](#)
- [52] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. [1](#)
- [53] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? Towards understanding the effectiveness of maml. In *ICLR*, 2020. [2](#), [3](#), [12](#)
- [54] Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. *Advances in neural information processing systems*, 30, 2017. [12](#)
- [55] Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake Bruce, Ali Razavi, Ashley Edwards, Nicolas Heess, Yutian Chen, Raia Hadsell, Oriol Vinyals, Mahyar Bordbar, and Nando de Freitas. A generalist agent, 2022. [1](#)
- [56] Mengye Ren, Eleni Triantafillou, Jake Snell Sachin Ravi, Kevin Swersky, Joshua B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. Meta-learning for semi-supervised few-shot classification. In *ICLR*, 2018. [2](#)
- [57] Sebastian Ruder. An overview of multi-task learning in deep neural networks. In *arXiv*, 2017. [2](#), [7](#)
- [58] Amirreza Shaban, Shray Bansal, Z. Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *BMVC*, 2017. [4](#), [12](#)
- [59] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017. [5](#)
- [60] Haozhe Sun, Wei-Wei Tu, and Isabelle M Guyon. Omniprint: A configurable printed character synthesizer. In *NeurIPS Datasets and Benchmarks Track*, 2021. [12](#)
- [61] Sebastian Thrun and Lorien Pratt, editors. *Learning to Learn*. Kluwer Academic Publishers, 1998. [1](#)
- [62] 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. [2](#), [3](#), [6](#), [7](#)
- [63] 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. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [64] Yu-Yun Tseng, Alexander Bell, and Danna Gurari. Vizwiz-fewshot: Locating objects in images taken by people with visual impairments. In *ECCV*, 2022. [4](#), [12](#)
- [65] Ihsan Ullah, Dustin Carrion, Sergio Escalera, Isabelle Guyon, Mike Huisman, Felix Mohr, Jan N. van Rijn, Haozhe Sun, Joaquin Vanschoren, and Phan Anh Vu. Meta-album: Multi-domain meta-dataset for few-shot image classification. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [12](#), [14](#)
- [66] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *NIPS*, 2016. [2](#)

- [67] Risto Vuorio, Shao-Hua Sun, Hexiang Hu, and Joseph J. Lim. Multimodal model-agnostic meta-learning via task-aware modulation. In *NeurIPS*, 2019. 2, 3, 6
- [68] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *ICCV*, 2019. 5
- [69] Yan Wang, Wei-Lun Chao, Kilian Q. Weinberger, and Laurens van der Maaten. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning, 2019. 6
- [70] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM Computing Surveys (CSUR)*, 53(3):1–34, 2020. 1, 2
- [71] Tete Xiao, Xiaolong Wang, Alexei A Efros, and Trevor Darrell. What should not be contrastive in contrastive learning. In *ICLR*, 2021. 2
- [72] Yang Xiao and Renaud Marlet. Few-shot object detection and viewpoint estimation for objects in the wild. In *ECCV*, 2020. 3
- [73] Lumin Xu, Sheng Jin, Wang Zeng, Wentao Liu, Chen Qian, Wanli Ouyang, Ping Luo, and Xiaogang Wang. Pose for everything: Towards category-agnostic pose estimation. In *ECCV*, 2022. 2, 3
- [74] Yongxin Yang and Timothy Hospedales. A unified perspective on multi-domain and multi-task learning. In *ICLR*, 2015. 2
- [75] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *CVPR*, 2020. 3
- [76] Mingzhang Yin, George Tucker, Mingyuan Zhou, Sergey Levine, and Chelsea Finn. Meta-learning without memorization. 2020. 4, 12
- [77] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *NeurIPS*, 2020. 2, 7, 8
- [78] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *CORL*, 2019. 3
- [79] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, 2018. 3
- [80] Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruysen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, et al. A large-scale study of representation learning with the visual task adaptation benchmark. In *arXiv*, 2019. 2, 7, 8
- [81] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. 5
- [82] Linyu Zheng, Ming Tang, Yingying Chen, Jinqiao Wang, and Hanqing Lu. Learning feature embeddings for discriminant model based tracking. In *ECCV*, 2020. 5
- [83] Bartosz Zieliński, Anna Plichta, Krzysztof Misztal, Przemysław Spurek, Monika Brzywczy-Włoch, and Dorota Ochońska. Deep learning approach to bacterial colony classification. *PLOS ONE*, 12(9):1–14, 09 2017. 12