

What Matters For Meta-Learning Vision Regression Tasks?

Ning Gao^{1,2} Hanna Ziesche¹ Ngo Anh Vien¹ Michael Volpp² Gerhard Neumann²
¹Bosch Center for Artificial Intelligence ²Autonomous Learning Robots, KIT
 {ning.gao, hanna.ziesche}@de.bosch.com anhvien.ngo@bosch.com
 {michael.volpp, gerhard.neumann}@kit.edu

Abstract

Meta-learning is widely used in few-shot classification and function regression due to its ability to quickly adapt to unseen tasks. However, it has not yet been well explored on regression tasks with high dimensional inputs such as images. This paper makes two main contributions that help understand this barely explored area. First, we design two new types of cross-category level vision regression tasks, namely object discovery and pose estimation of unprecedented complexity in the meta-learning domain for computer vision. To this end, we (i) exhaustively evaluate common meta-learning techniques on these tasks, and (ii) quantitatively analyze the effect of various deep learning techniques commonly used in recent meta-learning algorithms in order to strengthen the generalization capability: data augmentation, domain randomization, task augmentation and meta-regularization. Finally, we (iii) provide some insights and practical recommendations for training meta-learning algorithms on vision regression tasks. Second, we propose the addition of functional contrastive learning (FCL) over the task representations in Conditional Neural Processes (CNPs) and train in an end-to-end fashion. The experimental results show that the results of prior work are misleading as a consequence of a poor choice of the loss function as well as too small meta-training sets. Specifically, we find that CNPs outperform MAML on most tasks without fine-tuning. Furthermore, we observe that naive task augmentation without a tailored design results in underfitting.

1. Introduction

Humans are able to rapidly learn the fundamentals of new tasks within minutes of experience based on prior knowledge. For instance, humans can classify novel objects by capturing the distinguishable properties (e.g., textures, shapes and scales) from only a few examples. Meta-learning is proposed to learn relevant knowledge from various tasks and generalize to unseen tasks with only a few

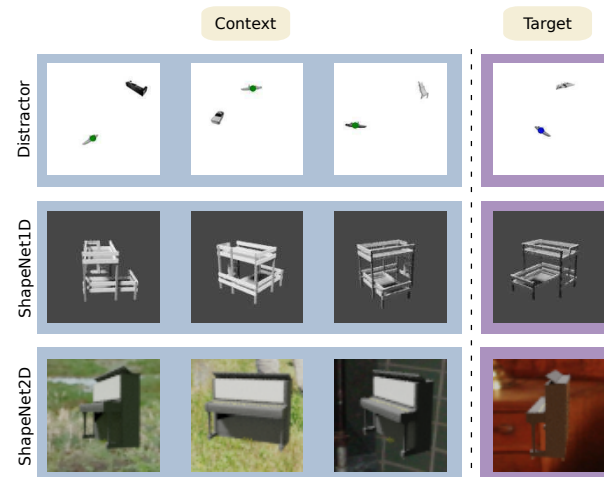


Figure 1. Meta-learning vision regression tasks are designed to i) identify the queried object from context and predict its position for target images (Distractor), ii) identify the object’s canonical pose from context and predict the 1D rotation relative to the canonical pose for target images (ShapeNet1D), iii) predict the 2D rotation w.r.t. the canonical pose with random background (ShapeNet2D). Predictions are performed on unseen objects.

samples. Of the various meta-learning algorithms, MAML-based models [1–4] and Neural Processes (NPs) [5–8] are two variants which are receiving increasing attention in the recent years. Both algorithms try to learn good prior knowledge from related tasks without expanding the learned parameters or sacrificing efficiency at inference. While these methods have shown promising results in many domains, such as few-shot classification [2, 9–12] and hyperparameter optimization [13–15], an extensive study on meta-learning vision regression tasks has not yet been conducted. This is in particular true for NPs which have mostly been investigated on tasks with low-dimensional input such as function regression or pixel-wise completion [16–19].

In this paper, we make two major contributions to the largely unexplored area of meta-learning on high-dimensional input tasks. On the algorithmic level, in-

spired by SimCLR [20], we propose an improvement to NPs by employing contrastive learning at the functional space (FCL) and still train the model in an end-to-end fashion. On the experimental side, we propose two application datasets, object discovery and pose estimation, which are based on high-dimensional inputs and require the meta-learning models to learn and reason at an image level.

For the first application we create a regression task called “Distractor” (see Fig. 1), where each image contains two objects, the queried object and a distractor object, placed at random positions. The goal of this task is to identify the queried object and predict its position in the image plane. Unlike previous tasks such as image completion, where each pixel is considered as an independent input, our task requires the model to learn a high-level representation from the entire image. The second application (i.e., pose estimation) is inspired by prior work [21–24] on the Pascal1D dataset. As this dataset shows limited object variations and features only 1D rotation around the azimuth axis, we generate two new datasets with increasing task diversity, e.g., by introducing random background, cross-categorical object variations and 2D rotation. Since the background is generated from real-world images instead of blank as in prior work, our datasets significantly increase the task difficulty and allow us to perform a thorough investigation of the performance for the considered meta-learning approaches. Examples of our datasets are shown in Fig. 1 where i) ShapeNet1D contains 1D rotations as in Pascal1D, however with larger object variations and ii) ShapeNet2D features 2D rotation and random background.

For both applications, we evaluate the performance on novel objects at both **intra-category** (IC) and **cross-category** (CC) levels. The results on Distractor show that our proposed algorithmic improvements significantly increase the performance, indicating our methods can enhance the task expressivity. The results on pose estimation demonstrate that meta-learning can successfully be applied to predict poses of unknown objects, which has a huge potential in robotic grasping and virtual/augmented reality (VR/AR).

Prior work [21–24] on Pascal1D also demonstrates that meta-learning algorithms suffer from overfitting, especially with limited training data. Our work analyzes the effect of different techniques commonly adopted in recent meta-learning methods (i.e., data augmentation, task augmentation, regularization and domain randomization) on aforementioned datasets. We empirically find that the meta-learning algorithms employed in our work ultimately lead to overfitting regardless of dataset size for both applications. Moreover, our work shows that the results in prior work [21, 22], where MAML typically performs best for such tasks, are misleading. In particular, we find Conditional Neural Processes (CNPs) [5] are more flexible and

efficient than MAML in the investigated pose regression tasks. Additionally, we find that MAML [1] suffers from underfitting especially on large-scale datasets and depends heavily on hyperparameter tuning.

The primary contributions of this work can be summarized as follows: (1) We investigate meta-learning algorithms on vision regression tasks and demonstrate their ability to tackle structured problems. (2) We propose functional contrastive learning on the task representation of CNPs and thereby improve its expressivity. (3) We quantitatively analyze various deep learning techniques to alleviate meta overfitting. Our results rectify misleading conceptions from prior work, e.g., that MAML performs best for such tasks. We also present insights and practical recommendations on designing and implementing meta-learning algorithms on vision regression tasks.

2. Related Work

Meta-Learning. In meta-learning, also known as *learning to learn*, a learning agent gains meta knowledge from previous learning episodes or different domains and then uses this acquired knowledge to improve the learning on future tasks [25]. MAML is an optimization-based meta-learning method and represents the meta knowledge as the model parameters, where learning good initial parameters can enable quick adaptation to new tasks with only few update steps on a small number of samples [1]. Different from MAML, Neural Processes (NPs) constitute a class of neural latent variable models and interpret meta-learning as conditional few-shot function regression [6]. Similar to Gaussian Processes, NPs model distributions over functions conditioned on contexts [6, 7, 26]. Meta-learning algorithms have been applied successfully in low-dimensional function regression [5–7, 16], image completion [8, 18, 19], few-shot classification [2, 9–12, 27], reinforcement learning [4, 28–31], and neural architecture search (NAS) [32–35]. Recent works [3, 21–23] go one step further and apply meta-learning to pose estimation using gray-scale images. However, in these studies, the prediction is restricted to 1D rotation and the employed loss function is ill-posed as it does not take the periodicity of rotation into consideration. Moreover, [36] proposes to improve meta-learning by adding contrastive representation learning from disjoint context sets. A follow-up work [37] further extends this idea to time series data by combining contrastive learning with ConvNP [8]. However, in contrast to these two methods which need to learn a representation in a self-supervised way and fine-tune on downstream tasks subsequently, we use functional contrastive learning (FCL) between context and target sets and train in an end-to-end fashion.

Meta Overfitting. It is well-known that meta-learning al-

gorithms suffer from two notorious types of overfitting: i) **Memorization overfitting** occurs when the model only conditions on the input to predict the output instead of relying on the context set [21]; ii) **Learner overfitting** happens when the prediction model and meta-learner overfit only to the training tasks and cannot generalize to novel tasks even though the prediction can condition on the context set [22]. Recently, different methods have been proposed to mitigate those overfitting issues, e.g., adding a regularization term on weights to restrict the memorization [21]. However, tuning a regularization term between underfitting and overfitting is challenging [38]. Subsequently, a related work [22] applied task augmentation which helps both memorization and learner overfitting. Meanwhile, [23] proposed MetaMix and Channel Shuffle to linearly combine features of context and target sets and replace channels with samples from different tasks. Furthermore, Ni et al. [24] empirically showed that data augmentation can also alleviate meta overfitting. Moreover, they find that employing data augmentation on target set achieves better performance. However, extensive comparisons on how these methods perform individually or combinedly are missing. In this work, we separate these techniques into data augmentation (DA), task augmentation (TA), meta-regularization (MR) and domain randomization (DR), and quantitatively compare them in different combinations on the two aforementioned applications in order to arrive at a better understanding and consistent comparisons.

3. Study Design

We now briefly describe both MAML and CNP in a unified way. We assume that all tasks are sampled from the same distribution $p(\mathcal{T})$, each task \mathcal{T}_i includes a context set $\mathcal{D}_C^i = \{(x_{C,1}, y_{C,1}), \dots, (x_{C,K}, y_{C,K})\}_i$ and a target set $\mathcal{D}_T^i = \{(x_{T,1}, y_{T,1}), \dots, (x_{T,M}, y_{T,M})\}_i$ where K and M are the number of samples in each set which could be different for each task. The entire training dataset is denoted as $\mathcal{D} = \{\mathcal{D}_C^i, \mathcal{D}_T^i\}_{i=1}^N$ where N is the number of tasks sampled for training. During inference, the model is tested on a new task $\mathcal{T}^* \sim p(\mathcal{T})$ given a small context set, from which it has to infer a new function $f^* : (\mathcal{D}_C^*, x_T^*) \rightarrow \hat{y}_T^*$. In meta-learning, there are two types of learned parameters, the first is the meta-parameters θ , which are learned during a meta-training phase using \mathcal{D} . The second is task-specific parameters ϕ^* which are updated based on samples from a new task \mathcal{D}_C^* conditioned on the learned meta-parameters θ . Predictions can be constructed as $\hat{y}_T^* = f_{\theta, \phi^*}(x_T^*)$, where f is the meta-model parameterized by θ and ϕ^* .

MAML considers both θ and ϕ^* as weights of neural networks while CNP considers only θ as neural weights. Different from MAML, which updates ϕ^* by gradient optimization on the new task samples, CNP takes ϕ^* as task representation and predicts it from the context set as $\phi^* =$

$\bigoplus_{i=1}^K h_{\theta}(x_{C,i}^*, y_{C,i}^*)$. Here \bigoplus is a permutation invariant operator, h is an encoder parameterized by θ . Subsequently, a decoder g_{θ} will take ϕ^* as an additional input and output $\hat{y}_T^* = g_{\theta}(x_T^*, \phi^*)$. Note that meta-parameters θ are fixed after meta-training phase, therefore CNPs don't not require any fine-tuning as MAML.

3.1. Problem Setting

In this paper, we consider two types of image-based regression tasks, namely object discovery and pose estimation. First, we propose a non-trivial object discovery task called Distractor, which is only used for evaluating CNP variants. In contrast to existing object detection tasks [39–42] that are designed to specify all object instances from an input image, our task aims to i) distinguish the queried object from other distractors and additionally ii) predict its 2D location in the image plane. Therefore, it is essential to learn a distinctive embedding ϕ^* that can represent various queried objects given their associated context images $\{x_{C,i}^*\}_{i=1}^K$ and corresponding positions $\{y_{C,i}^*\}_{i=1}^K$. Note that the distractors are sampled randomly from all categories and in many cases their appearances closely resemble the queried object. Hence, it is expected that aggregating multiple context pairs helps extracting expressive information to disambiguate the tasks and thus improve the performance.

The second task, pose estimation, is evaluated on three datasets, namely Pascal1D, ShapeNet1D and ShapeNet2D with incremental difficulty, caused e.g., by extending inference to unseen cross-category objects, adding random backgrounds and extending 1D rotations to 2D rotations. Note that in this task, each object has a random canonical pose, which has to be learned from a context set \mathcal{D}_C^* where $\{y_{C,i}^*\}_{i=1}^K$ are the ground-truth rotations of context images $\{x_{C,i}^*\}_{i=1}^K$.

We use these tasks for an exhaustive evaluation of meta-learning algorithms: i) We evaluate the performance of CNPs using different aggregation operators, i.e., mean [5], max, bayesian aggregation (BA) [43] and cross-attention (CA) [7]. ii) We evaluate MAML on Pascal1D and ShapeNet1D following [21, 22] and compare it with different CNP variants. iii) Furthermore, we investigate meta overfitting with respect to different choices, e.g., augmentations, regularization, aggregation operators and task properties. iv) Moreover, we combine functional contrastive learning (FCL) with CNPs and compare it with original CNPs.

3.2. Datasets

We generate **Distractor** that contains 12 object categories from ShapeNetCoreV2 [44], where each category includes 1000 randomly sampled objects. For each object we create 36 128×128 gray-scale images, containing two objects with random azimuth rotation and 2D position (see

Fig. 1). The data generation is based on an extended version of a prior open-source pipeline [26]. We choose 10 categories for training, where we reserve 20% of the data for intra-category (IC) evaluation. The remaining 2 categories are only used for cross-category (CC) evaluation. The second dataset, **Pascal1D** [21], contains 65 objects from 10 categories. We randomly select 50 objects for training and the other 15 objects for testing. 100 128×128 gray-scale images are rendered for each object with a random rotation in azimuth angle normalized between $[0, 10]$. Since the performance is limited due to the size of the dataset, we generate a larger dataset, **ShapeNet1D** which includes 30 categories. 27 of these are used during training and IC evaluation, the other 3 categories are used for CC evaluation. For each training category, we randomly sample 50 objects for training and 10 for IC evaluation while CC evaluation is performed on 20 objects for each unseen category. To further increase the task difficulty, we create **ShapeNet2D** which includes 2D rotations. We restrict the azimuth angles to the range $[0^\circ, 180^\circ]$ in order to reduce the effect of symmetric ambiguity while elevations are restricted to $[0^\circ, 30^\circ]$. Furthermore, we use RGB images and employ randomly sampled real-world images from SUN2012 [45] as background instead of static background.

3.3. Data Augmentation, Domain Randomization, Task Augmentation and Meta Regularization

Data Augmentation (DA). We use standard image augmentation techniques in our work, i.e., *Dropout* and *Affine* for all tasks, and an additional *CropAndPad* for all pose regression tasks. Furthermore, we employ *Contrast*, *Brightness* and *Blur* for ShapeNet2D. Details are presented in ??.

Domain Randomization (DR). For ShapeNet2D, we additionally employ DR [46] by regenerating background images for all training data after every 2k training iterations while the data used for evaluation remain the same.

Task Augmentation (TA). Task augmentation adds randomness to each task in order to encourage the meta-learner to learn non-trivial solutions instead of simply memorizing the training tasks. Following [22], we sample random noise $\epsilon^{(t)}$ from a discrete set for each task and create new tasks by adding the noise to the regression targets: $D_C^{(t)} = \{x_{C,i}^{(t)}, y_{C,i}^{(t)} + \epsilon^{(t)}\}_{i=1}^K$ and $D_T^{(t)} = \{x_{T,i}^{(t)}, y_{T,i}^{(t)} + \epsilon^{(t)}\}_{i=1}^M$. Specifically, we sample 2D position noise from a discrete set $\epsilon \in \{0, 1, 2, \dots, 16\}^2$ for Distractor. For Pascal1D, we use the same noise set $\{0., 0.25, 0.5, 0.75\}$ as proposed in [21, 22] while $\{0., 0.125, 0.25, \dots, 2\}$ for ShapeNet1D. In ShapeNet2D, we first only add random noise in the azimuth angle from the discrete set $\{-10^\circ, -9^\circ, \dots, 20^\circ\}$ and in a second step add additional elevation noise from the set $\{-5^\circ, -4^\circ, \dots, 10^\circ\}$ for further comparison.

Meta Regularization (MR). Following Yin et al. [21], we employ MR on the weights θ of the neural networks.

Furthermore, we find that it is crucial to fine-tune the coefficient β which modulates the regularizer and task information stored in the meta-parameters θ . In our experiments, we use $\beta = 1e^{-4}$ for Pascal1D, $1e^{-7}$ for ShapeNet1D and ShapeNet2D. More details about MR are presented in ??.

3.4. Functional Contrastive Learning (FCL)

The representations learned by CNP are invariant under permutation of the elements within a given context set. This property is achieved by a permutation invariant aggregation mechanism, e.g., max aggregation. However, another desirable property of the representation is invariance across context sets of the same task. In particular, the representations of different context sets belonging to the same task should be close to each other in the embedding space, while representations of different tasks should be farther apart. To achieve this, we add an additional contrastive loss at the functional space and train the model in an end-to-end fashion. The contrastive cross-entropy loss is defined as follows [20]:

$$\mathcal{L}_{\text{FCL}} = -\frac{2}{N} \sum_{t=1}^N \log \frac{\exp(\text{sim}(\phi_C^{(t)} \cdot \phi_T^{(t)})/\tau)}{D(\phi_C^{(t)})D(\phi_T^{(t)})}, \quad (1)$$

where N denotes the number of tasks per batch. $(\phi_C^{(t)}, \phi_T^{(t)})$ denotes a positive pair of latent representations of a given task obtained from context and target set respectively. More specifically, the pairs are obtained via max aggregation $\phi_C^{(t)} = \max(r_{C,1}^{(t)}, \dots, r_{C,K}^{(t)})$ and $\phi_T^{(t)} = \max(r_{T,1}^{(t)}, \dots, r_{T,M}^{(t)})$, where K denotes the number of context pairs per task and M the number of target pairs per task. \max returns the element-wise maximum over the latent variables $r_i = h_\theta(x_i, y_i)$ which are output by the encoder network h_θ for each context pair (x_i, y_i) . τ is a temperature parameter, which is crucial for learning good representations (details are presented in ??). $\text{sim}(\cdot)$ is the cosine similarity and $D(\phi_i^t)$ sums the similarity of all positive and negative pairs for ϕ_i^t :

$$D(\phi_i^t) = \sum_{k=1}^N \sum_{j \in \{C, T\}} \mathbb{1}_{\{k \neq t\} \vee \{j \neq i\}} \exp\left(\frac{\text{sim}(\phi_i^t \cdot \phi_j^k)}{\tau}\right), \quad (2)$$

where $\mathbb{1}_{\{k \neq t\} \vee \{j \neq i\}} \in \{0, 1\}$ is an indicator evaluating to 1 only if the representations are sampled from different tasks or different sets. The log-value in Eq. (1) can be interpreted as the weighted importance of the positive pair. Therefore, this loss function encourages the model to obtain large similarity for positive pairs and small for negative pairs.

3.5. Objective Functions and Evaluation Metrics

Pascal1D. Following prior work [21–23] we conduct experiments using the MSE score between predicted

and ground-truth azimuth rotation for both training and evaluation. However, this loss function does not take the ambiguity of coterminal angles into account. Hence, it can hamper the training process, e.g., predicting 359° for a ground-truth angle of 0° incurs a higher loss than predicting 180° . Nevertheless, we follow the same setup to obtain a fair comparison to prior works.

ShapeNet1D. Instead of using MSE score, we use the “cosine-sine-loss” for training and a prediction error defined in terms of the angular degree for evaluation. The loss of a single sample is defined as:

$$\mathcal{L} = |\cos(y) - \cos(y^*)|^2 + |\sin(y) - \sin(y^*)|^2, \quad (3)$$

where y^* is the ground-truth rotation and y the predicted rotation. The prediction error used for evaluation is defined as follows:

$$\mathcal{E} = \min\{\mathcal{E}_{y^+,y^*}, \mathcal{E}_{y^-,y^*}, \mathcal{E}_{y,y^*}\}, \quad (4)$$

where

$$\mathcal{E}_{y^\pm,y^*} = |y \pm 360 - y^*|, \mathcal{E}_{y,y^*} = |y - y^*|.$$

ShapeNet2D. We represent the 2D rotation as quaternion in both training and evaluation. The loss of a single sample is accordingly defined as follows:

$$\mathcal{L} = \min \left\{ \left| q^* - \frac{q}{\|q\|} \right|, \left| q^* - \frac{-q}{\|q\|} \right| \right\}, \quad (5)$$

where q^* denotes the ground-truth unit quaternion and q denotes the predicted quaternion. We empirically find that using this objective function achieves a better performance than constraining the scalar part of q to be positive. We hypothesize that enforcing the scalar constraint breaks the continuity of the rotation representation and therefore hampers training.

4. Experiments

In this section we present experimental results¹, perform a thorough analysis and provide insights and recommendations. Instead of presenting the results following the task sequence, we structure this section by different algorithmic choices and perform a systematic comparison over all tasks by raising different questions. ?? provides visualization examples of different tasks.

MAML or CNPs? We compare MAML and CNPs on two pose estimation datasets, Pascal1D and ShapeNet1D.

¹Codes and data are available at <https://github.com/boschresearch/what-matters-for-meta-learning>

Methods	Mean	Max	BA	CA	Max _{FCL}
No Aug	6.02	5.11	4.63	5.13	3.70
	6.89	6.17	5.91	6.39	4.61
DA	2.67	2.45	2.44	2.65	2.00
	4.10	3.75	3.97	4.08	3.05
TA	6.29	6.18	6.33	6.32	5.45
	7.19	7.04	7.02	7.02	6.66
TA+DA	3.20	3.09	2.65	3.05	2.60
	6.07	5.14	4.67	4.98	3.90

Table 1. Prediction error (pixel) on euclidean distance in the 2D image plane for Distractor. Different aggregation methods and augmentations are employed. The first row shows results for intra-category (IC) evaluations, the second row for cross-category (CC).

Methods	MAML	CNP (Mean)	CNP (CA)
No Aug	1.69 (0.22)	5.28 (0.51)	4.66 (0.74)
MR	1.90 (0.27)	2.96 (0.21)	3.33 (0.27)
TA	1.02 (0.06)	1.98 (0.22)	1.36 (0.25)
DA	2.10 (0.09)	3.69 (0.13)	2.90 (0.03)
TA+DA	1.31 (0.14)	2.29 (0.19)	1.77 (0.33)

Table 2. Pascal1D pose estimation error. MSE and standard deviations are calculated with 5 random seeds.

Methods	MAML	CNP (Max)	CNP (CA)
No Aug	25.27	14.97 (0.37)	8.19 (0.30)
	21.63	18.09 (0.21)	9.13 (0.18)
MR	13.23	12.71 (0.26)	8.87 (0.36)
	16.55	14.77 (0.35)	8.43 (0.39)
TA	23.01	10.89 (0.27)	7.92 (0.25)
	20.59	14.43 (0.55)	9.18 (0.50)
DA	14.69	8.64 (0.21)	6.24 (0.15)
	16.02	9.87 (0.35)	6.54 (0.19)
TA+DA	17.96	7.66 (0.18)	5.81 (0.23)
	18.79	8.66 (0.19)	6.23 (0.12)
TA+DA+FCL	—	7.82 (0.08)	6.44 (0.36)
	—	8.84 (0.04)	6.74 (0.20)
TA+DA+MR	13.45	10.54 (0.37)	8.28 (0.17)
	14.44	10.76 (0.30)	8.04 (0.10)

Table 3. ShapeNet1D pose estimation error($^\circ$). Results are calculated with 5 random seeds except for MAML. The first row presents results for IC and the second row for CC.

We obtain similar results as [21, 22] on Pascal1D, where MAML performs better than CNPs and the latter shows more severe overfitting (see Tab. 2). However, Tab. 3 illustrates that both CNP variants outperform MAML with a large margin on ShapeNet1D. It is good to note that, the prediction errors of all methods in Tab. 2 after denormal-

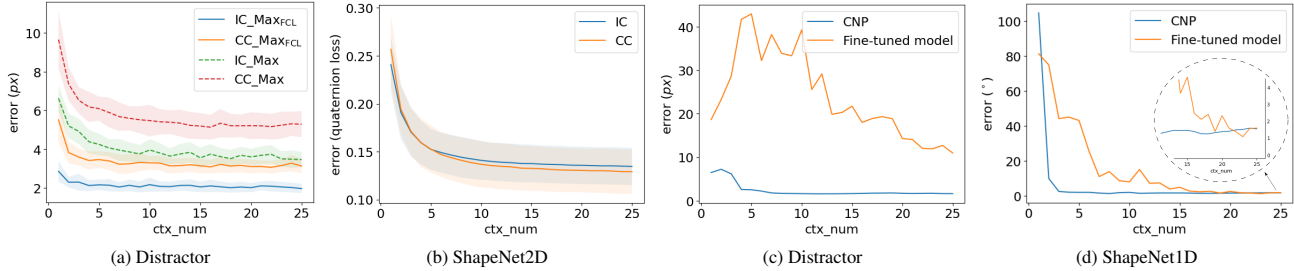


Figure 2. (a) CNP Prediction error (pixel) vs context number for the Distractor task using Max aggregation and Max + FCL (Max_{FCL}). Results are evaluated on novel objects from both intra-category (IC) and unseen cross-category (CC) levels. (b) CNP (CA) Prediction error vs context number for ShapeNet2D using DA + TA. (c) We compare a classical object detection method and CNP (Max) using same dataset for training on Distractor. The classical model is further fine-tuned on each new task. The results are shown in dependence of the number of images used for fine-tuning or as context set. (d) Prediction error between the fine-tuned model and CNP (CA) on ShapeNet1D.

izing are larger than 30° , indicating the experiments of prior work on Pascal1D simply used too little meta-data to make informative conclusions about the quality of different algorithms. Our interpretation is that MAML tries to learn a good initial prior (global optimum) which needs to be optimized on each specific task (fine-tuned optimum) within few samples and updates. On small datasets, MAML can easily find a global optimum that satisfies all the training tasks. At the same time MAML also overfits less, since the fine-tuning from global to fine-tuned optimum happens during inference time. However, finding a global optimum is getting difficult for large-scale datasets due to the increasing task diversity. Consequently, more samples and updates are necessary to fine-tune the task-specific parameters ϕ (see Sec. 3), which also explains why MAML is sensitive to hyperparameter tuning [47]. Furthermore, MAML shows much longer training times than CNPs, which limits us to conduct exhaustive comparisons on more complicated tasks such as Distractor or ShapeNet2D. In contrast, CNPs use the local parameterization ϕ as a fixed dimensional output of the encoder, which forces the model to learn an informative low-rank representation from the contexts. Meanwhile, increasing data and task diversity will encourage the model to extract more expressive and mutual-exclusive task representations.

DA, DR, TA or MR? From the results of different experiments presented in Tab. 1, Tab. 2, Tab. 3 and Tab. 4, it is obvious that DA improves the performance across all tasks and methods. Tab. 4 also shows the importance of DR on ShapeNet2D which cannot be simply compensated by DA. TA hinders the performance on Distractor but benefits all pose regression tasks. The reason for this is that, for Distractor, TA increases task complexity by shifting the origin of the image plain by the sampled noise, thus creating N^2 copies of the original task, where $N = 16$ is the number of non-zero elements in the noise set. However, since these task copies live in independent coordinate frames, the

Methods	IC ($1e^{-2}$)	CC ($1e^{-2}$)
None	38.33 (0.33)	39.81 (0.31)
DR	18.67 (0.13)	20.05 (0.12)
DR+MR	27.89 (0.61)	28.99 (0.46)
DR+TA _{azi}	16.94 (0.13)	18.42 (0.26)
DR+TA _{azi+ele}	16.62 (0.12)	17.76 (0.35)
DA	19.32 (0.09)	17.98 (0.09)
DR+DA	14.26 (0.09)	13.91 (0.14)
DR+DA+TA _{azi+ele}	14.12 (0.14)	13.59 (0.10)
DR+DA+TA _{azi+ele} + FCL	14.01 (0.09)	13.32 (0.18)

Table 4. Comparison of different augmentation techniques on ShapeNet2D. Results are calculated with 3 random seeds using CNP (CA) as baseline.

increased task diversity is irrelevant to the original task. For pose regression tasks, by contrast, TA augments the canonical poses of the existing data, which coherently benefits the original task as the augmented canonical poses remain in the coordinate frame of the original task. Therefore, even though TA increases the cross-entropy $\mathcal{H}(Y|X)$ for both cases as demanded in [22], only the pose regression tasks gain additional benefits. MR results in underfitting as combining MR with augmentations leads to worse performance than using the same augmentations alone for both ShapeNet1D and ShapeNet2D (see Tab. 3 and Tab. 4). Furthermore, MR requires extensive fine-tuning on the regularization parameter β (see Sec. 3.3) to modulate between underfitting and overfitting.

Effect of the context set size in CNPs. We compare the prediction error w.r.t. the size of the context set for Distractor (see Fig. 2a) and ShapeNet2D (see Fig. 2b). Both figures show that increasing the context set size benefits the performance, indicating that both Max and CA aggregations can merge useful information from different

Methods	CA _S	CA _M	CA _L	Max _S	Max _M	Max _L
No Aug	18.60 (0.78)	12.08 (0.44)	8.19 (0.30)	30.44 (0.82)	18.86 (0.34)	14.97 (0.37)
	19.95 (1.08)	12.62 (0.87)	9.13 (0.18)	30.59 (1.14)	21.78 (0.47)	18.09 (0.21)
TA	18.69 (0.87)	10.70 (0.98)	7.92 (0.25)	21.67 (0.66)	13.69 (0.27)	10.89 (0.27)
	19.24 (0.79)	12.05 (0.73)	9.18 (0.50)	23.60 (0.88)	16.76 (0.62)	14.43 (0.55)
TA+DA	7.86 (0.21)	6.32 (0.11)	5.81 (0.23)	11.00 (0.16)	8.23 (0.34)	7.66 (0.18)
	7.49 (0.35)	6.48 (0.41)	6.23 (0.12)	12.98 (0.48)	9.65 (0.40)	8.66 (0.19)

Table 5. Performance on ShapeNet1D using small (S), medium (M) and large (L) training dataset sizes for CNP with cross-attention (CA) and Max aggregation. The first row presents results for intra-category (IC) and the second row for cross-category (CC) evaluation. MSE and standard deviations are calculated with 5 random seeds.

context pairs and thereby reduce the task ambiguity. In addition, we find that the model can further improve the performance given the size of context set surpasses the maximum number used for training (15 for both tasks). In particular, there is a small performance gap between intra- and cross-category evaluation for Distractor which is however absent for ShapeNet2D. We believe this indicates that Distractor has more task ambiguity than pose estimation and thus explains why Distractor gains more benefits from FCL than ShapeNet2D (see Tab. 1 and Tab. 4).

CNPs vs pretrained models. It is a common practice in vision task to pretrain a model on a large-scale dataset (e.g., ImageNet [48]) in order to obtain good prior features and reduce training time. To conduct a fair comparison of this approach to our model regarding data efficiency, we first pretrain a classical object detection model jointly over all tasks using the same training data as for CNPs. After training has finished, we fine-tune the pretrained model further on each specific new task using different numbers of images. Results are shown in Fig. 2c for Distractor and Fig. 2d for ShapeNet1D, where the horizontal axis denotes the number of images used for fine-tuning or as contexts for CNPs, respectively. Both figures show that CNPs outperform the pretrained model especially for small numbers of contexts. In the Distractor task, CNP (Max) outperforms the fine-tuned model with a large margin after 25 context images are given. Note that CNPs are capable of transferring to various tasks simultaneously. In contrast, the pretrained model requires separate tuning on each given task, which results in a decreased performance on prior learning tasks.

Which aggregation methods should I use? Cross-attention (CA) performs better than mean aggregation on Pascal1D (see Tab. 2) and Max on ShapeNet1D (see Tab. 3), while it achieves a similar performance to Max aggregation and BA on ShapeNet2D (see Tab. 6). In contrast, mean aggregation used in the original CNP performs the worst on both Pascal1D and ShapeNet2D. Our interpretation is that Mean assigns the same importance to each context while the

Methods	IC ($1e^{-2}$)	CC ($1e^{-2}$)
CNP+Mean	15.04 (0.08)	15.45 (0.13)
CNP+Max	14.20 (0.06)	13.56 (0.28)
CNP+BA	14.16 (0.08)	13.56 (0.18)
CNP+CA	14.12 (0.14)	13.59 (0.10)

Table 6. Comparison of aggregation methods on ShapeNet2D using DR+DA+TA. Results are calculated with 3 random seeds.

other aggregation operators can allocate different weights. Max assigns a weight of one to a context and zero to all others for each dimension of the representation while BA assigns the weights predicted by another neural network. Meanwhile, CA assigns importance by comparing the similarity between context inputs $\{x_C^i\}_{i=1}^K$ and target input x_T at the feature space.

Furthermore, we find that CA achieves competitive results on all pose estimation tasks but performs slightly worse than BA and Max on Distractor (see Tab. 1) though still better than mean aggregation. This indicates that CA helps in learning representations for object-centric images. Distractor, however, contains objects with random locations, requiring the model to disregard positional information. Methods like CA, which compare similarity between contexts and target over feature space, face inherent difficulties on Distractor. This is due to the fact that CNNs, owing to their translational equivariant nature, are prone to encode some positional information into the extracted image features. Consequently, CA, which compares the similarity directly on this feature space, inevitably forces the model to focus on positional similarity, which leads to a suboptimal allocation of importance.

How much meta-data is essential? We split the training data of ShapeNet1D into subsets of three different sizes, with 10 objects per category for the small dataset (S), 30 objects per category for the medium dataset (M) and 50 for the large dataset (L). Afterwards, we test the performance of CNP with Max aggregation and CA on each of them.

The results in Tab. 5 show that Max overfits on the small dataset by simply memorizing all training tasks while CA works much better. Moreover, CA trained on small dataset achieves a comparable performance with Max on large dataset after using TA and DA, and even outperforms Max on the cross-category level. Thus, we conclude that using CA in combination with augmentation techniques can drastically alleviate the overfitting problem and therefore requires less meta-data on object-centric vision tasks than Max. In contrast, MAML performs much worse on ShapeNet1D (L) (see Tab. 3) than CNPs and thus hardly profits from an increased dataset.

Data augmentation. Tab. 7 shows the effect of each individual data augmentation technique (see Sec. 3.3) on ShapeNet2D. The first row contains results obtained with all techniques applied jointly. In the other rows, one of the techniques is removed respectively. We find that removing *Affine* leads to the worst performance which indicates that object-centric pose regression tasks are more sensitive to scale. On the other hand, omitting *CropAndPad* even leads to an performance increase.

Methods	Val	Test
All	0.1417	0.1410
w/o CropAndPad	0.1412	0.1368
w/o Affine	0.1623	0.1743
w/o Dropout	0.1452	0.1445
w/o Contrast	0.1482	0.1406
w/o Brightness	0.1454	0.1380
w/o Blur	0.1426	0.1422

Table 7. Comparison of different data augmentation techniques on ShapeNet2D using CNP (CA) + DR as baseline.

Methods	IC	CC
Same Ctx	2.30 (0.04)	3.46 (0.06)
Diff Ctx	2.16 (0.05)	3.25 (0.05)
Ctx & Target	2.00 (0.02)	3.05 (0.08)

Table 8. Analysis of FCL + CNP on different choices of positive pairs using: i) the same context set with different augmentations (Same Ctx), ii) different context sets from the same task (Diff Ctx), iii) context and target sets (Ctx & Target). Prediction error (pixel) is calculated with 3 random seeds.

Does FCL improve CNPs? Tab. 1 shows the evaluation on Distractor using different aggregation methods where Max_{FCL} denotes Max aggregation with FCL. Modulating task representation by functional contrastive learning (FCL) alleviates meta overfitting across all augmentation

levels and thus achieves a significant improvement in performance. Fig. 2a further compares the performance of Max and Max_{FCL} for different context set sizes, showing that our methods can differentiate the queried object and distractors well, even for very small context sets. Furthermore, we investigate the influence of FCL on the predicted task representations over all 12 categories using different clustering metrics, where the results show that FCL leads to a more dispersed latent distribution compared to the original CNPs, which can improve generalization capability to unseen tasks. T-SNE visualizations of the task representations along with the results of cluster metrics are provided in ??.

FCL on different sets. We compare FCL on three choices of positive pairs: i) We use the same context set but with different data augmentations. ii) We use different context sets sampled from the same task. iii) We use context and target sets from the same task. We test the performance on Distractor using Max aggregation and DA. For each choice, we run three experiments with different seeds and present the average performance in Tab. 8. Compared to Tab. 1, all three choices consistently outperform CNP (Max) while using FCL on context and target sets achieves the best performance.

Limitations. Since meta-learning algorithms are data-driven methods, generalization depends on the diversity of training tasks. However, the augmentation methods used in our work are limited in their capability of creating new and diverse training tasks. Therefore, using a generative model to enrich training data could be one possibility to achieve higher diversity. Furthermore, concerning the class of NPs, we restricted ourselves to the deterministic CNPs in the experiments since we empirically found that using stochastic variants such as [6] does not benefit our tasks. We leave an in-depth exploration of stochastic NPs on vision tasks to future work.

5. Conclusion

In this paper, we investigate MAML and CNPs on several image-level regression tasks and analyze the importance of different choices in mitigating meta overfitting. Furthermore, we provide insights and practical recommendations of different algorithmic choices for CNPs with respect to various task settings. In addition, we combine CNPs with functional contrastive learning in task space and train in an end-to-end manner, which significantly improves the task expressivity of CNPs. We believe that our work can lay the basis for future work on designing and implementing meta-learning algorithms in image-based regression tasks.

References

- [1] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 06–11 Aug 2017, pp. 1126–1135. [Online]. Available: <https://proceedings.mlr.press/v70/finn17a.html> 1, 2
- [2] A. Nichol, J. Achiam, and J. Schulman, “On first-order meta-learning algorithms,” *ArXiv*, vol. abs/1803.02999, 2018. 1, 2
- [3] C. Finn, A. Rajeswaran, S. Kakade, and S. Levine, “Online meta-learning,” in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 1920–1930. [Online]. Available: <https://proceedings.mlr.press/v97/finn19a.html> 2
- [4] J. Yoon, T. Kim, O. Dia, S. Kim, Y. Bengio, and S. Ahn, “Bayesian model-agnostic meta-learning,” in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018. [Online]. Available: <https://proceedings.neurips.cc/paper/2018/file/e1021d43911ca2c1845910d84f40aeae-Paper.pdf> 1, 2
- [5] M. Garnelo, D. Rosenbaum, C. Maddison, T. Ramalho, D. Saxton, M. Shanahan, Y. W. Teh, D. Rezende, and S. M. A. Eslami, “Conditional neural processes,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 1704–1713. [Online]. Available: <https://proceedings.mlr.press/v80/garnelo18a.html> 1, 2, 3
- [6] M. Garnelo, J. Schwarz, D. Rosenbaum, F. Viola, D. J. Rezende, S. M. A. Eslami, and Y. W. Teh, “Neural processes,” in *ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*, 2018. 2, 8
- [7] H. Kim, A. Mnih, J. Schwarz, M. Garnelo, A. Eslami, D. Rosenbaum, O. Vinyals, and Y. W. Teh, “Attentive neural processes,” in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=SkE6PjC9KX> 2, 3
- [8] J. Gordon, W. P. Bruinsma, A. Y. K. Foong, J. Requeima, Y. Dubois, and R. E. Turner, “Convolutional conditional neural processes,” in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=SkE6PjC9KX> 1, 2
- [9] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, “Learning to compare: Relation network for few-shot learning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 1199–1208. 1, 2
- [10] J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/cb8da6767461f2812ae4290eac7cbc42-Paper.pdf>
- [11] M. Ren, S. Ravi, E. Triantafillou, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel, “Meta-learning for semi-supervised few-shot classification,” in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=HJcSzz-CZ>
- [12] O. Vinyals, C. Blundell, T. Lillicrap, k. kavukcuoglu, and D. Wierstra, “Matching networks for one shot learning,” in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016. [Online]. Available: <https://proceedings.neurips.cc/paper/2016/file/90e1357833654983612fb05e3ec9148c-Paper.pdf> 1, 2
- [13] Y. Wei, P. Zhao, and J. Huang, “Meta-learning hyperparameter performance prediction with neural processes,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 11058–11067. [Online]. Available: <https://proceedings.mlr.press/v139/wei21c.html> 1
- [14] M. Volpp, L. P. Fröhlich, K. Fischer, A. Doerr, S. Falkner, F. Hutter, and C. Daniel, “Meta-learning acquisition functions for transfer learning in bayesian optimization,” in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=ryeYpJSKwr>
- [15] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil, “Bilevel programming for hyperparameter optimization and meta-learning,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 1568–1577. [Online]. Available: <https://proceedings.mlr.press/v80/franceschi18a.html> 1
- [16] Q. Wang and H. Van Hoof, “Doubly stochastic variational inference for neural processes with hierarchical latent variables,” in *Proceedings of the 37th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 10018–10028. [Online]. Available: <https://proceedings.mlr.press/v119/wang20s.html> 1, 2
- [17] A. Norcliffe, C. Bodnar, B. Day, J. Moss, and P. Liò, “Neural {ode} processes,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=27acGyyI1BY>

- [18] C. Louizos, X. Shi, K. Schutte, and M. Welling, “The functional neural process,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/db182d2552835bec774847e06406bfa2-Paper.pdf> 2
- [19] B.-J. Lee, S. Hong, and K.-E. Kim, “Residual neural processes,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, pp. 4545–4552, Apr. 2020. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/5883> 1, 2
- [20] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *Proceedings of the 37th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 1597–1607. [Online]. Available: <https://proceedings.mlr.press/v119/chen20j.html> 2, 4
- [21] M. Yin, G. Tucker, M. Zhou, S. Levine, and C. Finn, “Meta-learning without memorization,” in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=BkIEFpEYwS> 2, 3, 4, 5
- [22] J. Rajendran, A. Irpan, and E. Jang, “Meta-learning requires meta-augmentation,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 5705–5715. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/3e5190eeb51ebef6c5bbc54ee8950c548-Paper.pdf> 2, 3, 4, 5, 6
- [23] H. Yao, L.-K. Huang, L. Zhang, Y. Wei, L. Tian, J. Zou, J. Huang, and Z. Li, “Improving generalization in meta-learning via task augmentation,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 11887–11897. [Online]. Available: <https://proceedings.mlr.press/v139/yao21b.html> 2, 3, 4
- [24] R. Ni, M. Goldblum, A. Sharaf, K. Kong, and T. Goldstein, “Data augmentation for meta-learning,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 8152–8161. [Online]. Available: <https://proceedings.mlr.press/v139/ni21a.html> 2, 3
- [25] T. M. Hospedales, A. Antoniou, P. Micaelli, and A. J. Storkey, “Meta-learning in neural networks: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021. 2
- [26] J. Gordon, J. Bronskill, M. Bauer, S. Nowozin, and R. Turner, “Meta-learning probabilistic inference for prediction,” in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=HkxStoC5F7> 2, 4
- [27] R. Vuorio, S.-H. Sun, H. Hu, and J. J. Lim, “Multi-modal model-agnostic meta-learning via task-aware modulation,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/e4da3b7fbbce2345d7772b0674a318d5-Paper.pdf> 2
- [28] K. Rakelly, A. Zhou, C. Finn, S. Levine, and D. Quillen, “Efficient off-policy meta-reinforcement learning via probabilistic context variables,” in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 5331–5340. [Online]. Available: <https://proceedings.mlr.press/v97/rakelly19a.html> 2
- [29] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, “Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning,” *CoRR*, vol. abs/1910.10897, 2019. [Online]. Available: <http://arxiv.org/abs/1910.10897>
- [30] A. Gupta, R. Mendonca, Y. Liu, P. Abbeel, and S. Levine, “Meta-reinforcement learning of structured exploration strategies,” in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018. [Online]. Available: <https://proceedings.neurips.cc/paper/2018/file/4de754248c196c85ee4fbdcee89179bd-Paper.pdf>
- [31] M. W. Gondal, S. Joshi, N. Rahaman, S. Bauer, M. Wuthrich, and B. Schölkopf, “Function contrastive learning of transferable meta-representations,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 3755–3765. [Online]. Available: <https://proceedings.mlr.press/v139/gondal21a.html> 2
- [32] H. Liu, K. Simonyan, and Y. Yang, “DARTS: Differentiable architecture search,” in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=SlEYHoC5FX> 2
- [33] B. Zoph and Q. V. Le, “Neural architecture search with reinforcement learning,” in *International Conference on Learning Representations*, 2017. [Online]. Available: <https://openreview.net/forum?id=r1Ue8Hcxg>
- [34] T. Elsken, B. Staffler, J. H. Metzen, and F. Hutter, “Meta-learning of neural architectures for few-shot learning,” in

IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020.

- [35] H. Lee, E. Hyung, and S. J. Hwang, “Rapid neural architecture search by learning to generate graphs from datasets,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=rkQuFUuUOg3> 2
- [36] M. W. Gondal, S. Joshi, N. Rahaman, S. Bauer, M. Wutrich, and B. Schölkopf, “Function contrastive learning of transferable meta-representations,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 3755–3765. [Online]. Available: <https://proceedings.mlr.press/v139/gondal21a.html> 2
- [37] K. Kallidromitis, D. Gudovskiy, K. Kazuki, O. Iku, and L. Rigazio, “Contrastive neural processes for self-supervised learning,” in *Asian Conference on Machine Learning*, 2021. [Online]. Available: <http://www.acml-conf.org/2021/conference/accepted-papers/266/> 2
- [38] H. Li, P. Chaudhari, H. Yang, M. Lam, A. Ravichandran, R. Bhotika, and S. Soatto, “Rethinking the hyperparameters for fine-tuning,” in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=B1g8VkhHFPH> 3
- [39] T.-Y. Lin, M. Maire, S. J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft coco: Common objects in context,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014, pp. 740–755. 3
- [40] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3354–3361, 2012.
- [41] M. Everingham, L. V. Gool, C. K. I. Williams, J. M. Winn, and A. Zisserman, “The pascal visual object classes (voc) challenge,” *International Journal of Computer Vision*, vol. 88, pp. 303–338, 2010.
- [42] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and L. Fei-Fei, “Imagenet large scale visual recognition challenge,” *International Journal of Computer Vision*, vol. 115, pp. 211–252, 2015. 3
- [43] M. Volpp, F. Flürenbrock, L. Grossberger, C. Daniel, and G. Neumann, “Bayesian context aggregation for neural processes,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=ufZN2-ahFa> 3
- [44] A. X. Chang, T. A. Funkhouser, L. J. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, “Shapenet: An information-rich 3d model repository,” *ArXiv*, vol. abs/1512.03012, 2015. 3
- [45] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, “Sun database: Large-scale scene recognition from abbey to zoo,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 3485–3492. 4
- [46] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 23–30. 4
- [47] A. Antoniou, H. Edwards, and A. Storkey, “How to train your MAML,” in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=HJGven05Y7> 6
- [48] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 248–255. 7