

---

# Function Contrastive Learning of Transferable Meta-Representations

---

Muhammad Waleed Gondal<sup>1</sup> Shruti Joshi<sup>1</sup> Nasim Rahaman<sup>1,2</sup> Stefan Bauer<sup>1,3</sup> Manuel Wüthrich<sup>1</sup>  
Bernhard Schölkopf<sup>1</sup>

## Abstract

Meta-learning algorithms adapt quickly to new tasks that are drawn from the same task distribution as the training tasks. The mechanism leading to fast adaptation is the conditioning of a downstream predictive model on the inferred representation of the task’s underlying data generative process, or *function*. This *meta-representation*, which is computed from a few observed examples of the underlying function, is learned jointly with the predictive model. In this work, we study the implications of this joint training on the transferability of the meta-representations. Our goal is to learn meta-representations that are robust to noise in the data and facilitate solving a wide range of downstream tasks that share the same underlying functions. To this end, we propose a decoupled encoder-decoder approach to supervised meta-learning, where the encoder is trained with a contrastive objective to find a good representation of the underlying function. In particular, our training scheme is driven by the self-supervision signal indicating whether two sets of examples stem from the same function. Our experiments on a number of synthetic and real-world datasets show that the representations we obtain outperform strong baselines in terms of downstream performance and noise robustness, even when these baselines are trained in an end-to-end manner.

## 1. Introduction

Many supervised learning problems are concerned with approximating a data-generating function  $f : \mathcal{X} \rightarrow \mathcal{Y}$  given a finite set of  $N$  samples,  $\{x_i, y_i = f(x_i)\}_{i=1}^N$ . Expressive models, such as deep neural networks, are known to excel at this function approximation task, but they often heavily

rely on the number of samples  $N$  being large. This poses further challenges: in many domains of interest, sourcing enough data is a challenging endeavour; further, the process of training such models can be prohibitively slow for many applications. This is exacerbated by the fact that in the typical setting, each new data-generating function encountered requires that the model be retrained. In other words, the model is not shared between data-generating functions, even when training a model to approximate one function can potentially be beneficial for approximating another function.

To overcome these challenges, a variety of meta-learning methods have been proposed (Vinyals et al., 2016; Snell et al., 2017; Garnelo et al., 2018a; Ravi & Larochelle, 2016; Finn et al., 2017). In the present work, we are interested in a class of models that use encoder-decoder architectures such as Conditional Neural Processes (CNPs) (Garnelo et al., 2018a) and Generative Query Networks (GQNs) (Eslami et al., 2018). In the first stage, an encoder is used to infer a fixed-dimensional representation of a given function  $f$  from just a few input-output examples  $O^f = \{(x_i, y_i)\}_i$ , the *context dataset*. We call it the *meta-representation* of the function  $r = h_\phi(O^f)$ , where  $h$  is an encoder parameterized by  $\phi$ . In the second stage, the meta-representation is then used to condition a predictive model in order to solve a downstream prediction task related to that function. For instance, the task may consist of predicting the function value  $y$  at unseen locations  $x$  or classifying images after observing only a few pixels (in that case,  $x$  is the pixel location and  $y$  is the pixel value). This two-stage process has multiple benefits. First, the extraction of prior knowledge about  $f$  directly from the training data, in the form of meta-representation, reduces the need for specifying inductive biases (model architectures, training details, etc.) particular to  $f$ . Thus, it allows learning to be shared between functions such that a single model can be trained on a distribution over functions. Second, the computation of meta-representations is efficient and can be done online. Third, the computation of meta-representations provides flexibility to solve a variety of downstream tasks concerning a specific function.

However, CNPs optimize encoder jointly with the decoder on the downstream prediction task, i.e., prediction of function values  $y$  at unseen locations  $x$ , as illustrated in Figure 1(a). This ties the meta-representation’s quality to the

---

<sup>1</sup>Max Planck Institute for Intelligent Systems, Tübingen, Germany <sup>2</sup>Mila, University of Montreal, Montreal, Canada <sup>3</sup>CIFAR Azrieli Global Scholar. Correspondence to: Muhammad Waleed Gondal <waleed.gondal@tue.mpg.de>.

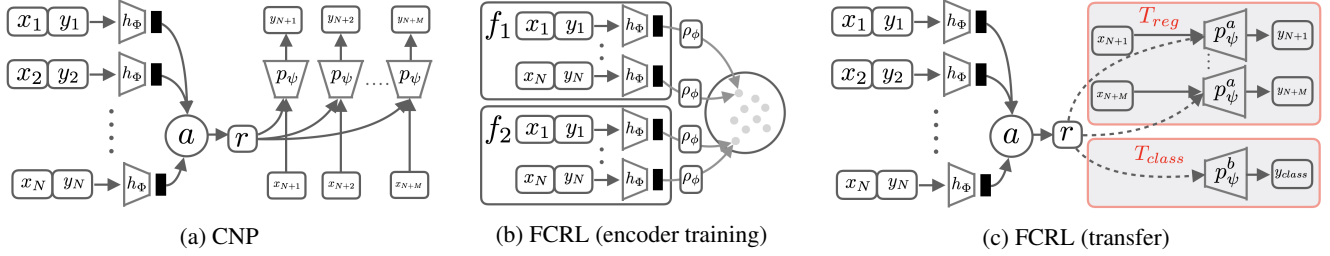


Figure 1. The difference in the training of CNP (Garnelo et al., 2018a) and FCRL for learning meta-representations  $r$  of functions. (shown left) CNP learns the aggregated representation  $r$  of the context set by maximizing the conditional likelihood of the target data. (shown center) Training of FCRL encoder  $h_\phi$  by contrasting context sets of different functions. Note that the target inputs  $x_{N+1}$  etc., are not used at this stage. (shown right) Using the pretrained FCRL encoder  $h_\phi$ , we train separate decoders  $p_\psi^*$  for each downstream task, shown in grey boxes. The dotted arrows indicate the transfer of inferred meta-representations to the tasks.

combined encoder and decoder performances on this particular task and thereby makes it susceptible to supervision collapse, i.e. the representations lose any information which is irrelevant for solving the training task, but may be necessary for the transfer to new tasks (Doersch et al., 2020). Moreover, many real-world tasks are noisy, and the prediction task might entail reconstructing high dimensional data, such as images in GQNs (Eslami et al., 2018). The corresponding objective function can cause the model to waste its capacity on reconstructing unimportant features such as static backgrounds and noise, while ignoring visually small but important details in its learned representation (Anand et al., 2019; Kipf et al., 2019). This is crucial for many real-world applications; for instance, in order to manipulate a small object in a complex scene, the model’s ability to infer the object’s shape carries more importance than inferring its color or reconstructing the static background.

In this work, we study the generalization of a function’s meta-representations in terms of their transferability to downstream tasks and their robustness to noise. We empirically show that the joint optimization of meta-representations and a prediction task is detrimental to the transferability of meta-representations and makes them vulnerable to noise. To address this issue, we propose a decoupled encoder-decoder training scheme, wherein the encoder is exclusively trained by a novel contrastive learning framework which we call FCRL (Function Contrastive Representation Learning). Instead of relying on reconstructions, it learns by contrasting sets of input-output pairs sampled from different functions. The key idea is that two sets of samples from the same function should have similar latent representations, while representations of sets of samples from different functions should remain easily distinguishable. FCRL retains the useful properties of meta-representations such as shared learning and sample efficiency while improving its transferability to downstream tasks and robustness to noise. Unlike contemporary meta-learning algorithms, meta-representations in FCRL are explicitly optimized over a distribution of functions rather than tasks.

To evaluate the effectiveness of the proposed method, we conduct comprehensive experiments on diverse downstream problems, including classification, regression, parameter identification, scene understanding, scene reconstruction and reinforcement learning. We consider different datasets, ranging from simple 1D and 2D regression to challenging simulated and real-world scenes. In particular, we find that a downstream predictor trained with our (pre-trained) encoder compares favorably to related methods on these tasks, including ones where the predictor is trained jointly with the encoder.

## 2. Preliminaries

### 2.1. Problem Setting

Consider a distribution over data-generating functions  $p(f)$ . Let  $f$  be a sample from this distribution  $f \sim p(f)$ , where  $f : \mathcal{X} \rightarrow \mathcal{Y}$  with  $\mathcal{X} = \mathbb{R}^d$  and  $\mathcal{Y} \subseteq \mathbb{R}^{d'}$ :

$$y = f(x, \xi); \quad \xi \sim \mathcal{Z} \quad (1)$$

where  $\xi$  is sampled from some noise distribution  $\mathcal{Z}$ . Let  $O^f = \{(x_i, y_i)\}_{i=1}^N$  be a set of few observed examples of a function  $f$ , referred to as the context set, and consider a set of downstream tasks  $\mathcal{T}$ . Here, each task  $T \in \mathcal{T}$  can be defined as a mapping defined over  $f$ . In the case of few shot regression (see Section 4.1),  $T$  maps from  $f$  to a predictive model  $p_\psi(y|x)$ . In the case of parameter identification,  $T$  maps from  $f$  to some scalar or vector valued parameter of  $f$ . Our goal is therefore to learn an encoder which maps a context set  $O^f$  to a representation of  $f$  that can interchangeably be used for multiple downstream tasks  $\mathcal{T}$  defined on the same function (without requiring retraining).

### 2.2. Background

In this section, we briefly discuss a class of meta-learning methods that are particularly relevant to our encoder-decoder setting, namely conditional neural processes (CNPs) and generative query networks (GQNs) (Garnelo et al., 2018a;b; Eslami et al., 2018).

**Conditional Neural Processes (CNPs).** The key proposal in CNPs (applied to few-shot learning) is to express a distribution over predictor functions given a context set. They learn the meta-representations  $r$  by jointly training the encoder and decoder, as illustrated in Figure 1(a). To this end, they first encode the context  $O^f$  into individual representations  $r_i = h_\Phi(x_i, y_i) \forall i \in [N]$ , where  $h_\Phi$  is a neural network. The representations are then aggregated via a mean-pooling operation into a fixed size vector  $r = 1/N(r_1 + r_2 + \dots + r_N)$ . The idea is that  $r$  captures all the relevant information about  $f$  from the context set  $O^f$ ; accordingly, the predictive distribution is approximated by maximizing the conditional likelihood of the target distribution  $p(y|x, O^f)$ , where  $y = f(x)$ .

**Generative Query Networks (GQN).** GQN (Eslami et al., 2018) can be seen as an extension of NPs (Garnelo et al., 2018b) for learning 3D scenes representations. The context dataset  $O^f$  in GQN consists of tuples of camera viewpoints in 3D space ( $\mathcal{X}$ ) and the images taken from those viewpoints ( $\mathcal{Y}$ ). Like NPs, GQNs learn to infer the latent representation of the scene (a function) by conditioning on the aggregated context and maximizing the likelihood of generating the correct image corresponding to a queried viewpoint.

### 3. Function-Contrastive Representation Learning (FCRL)

We take the perspective here that the sets of context points  $O^f$  provide a partial observation of an underlying function  $f$ . Our goal is to find an encoder  $g_{(\phi, \Phi)}$  which maps such partial observations to low-dimensional representations of the underlying function. The key idea is that a good encoder  $g_{(\phi, \Phi)}$  should map different context sets (i.e. partial observations) of the same function to be close in the latent space, such that they can easily be identified among context sets of different functions. This motivates the contrastive-learning objective which we will detail in the following.

**Encoder Structure.** Since the inputs to the encoder  $g_{(\phi, \Phi)}$  are sets, it needs to be permutation invariant with respect to input order and able to process inputs of varying sizes. We enforce this permutation invariance in  $g_{(\phi, \Phi)}$  via sum-decomposition, proposed by (Zaheer et al., 2017). We first average-pool the point-wise transformations of  $O^f$  to get the encoded representations

$$r^f = \frac{1}{|O^f|} \sum_{(x,y) \in O^f} h_\Phi(x, y) \quad (2)$$

where  $h_\Phi(\cdot)$  is the encoder network. We then obtain a nonlinear projection of this encoded representation  $g_{(\phi, \Phi)}(O^f) = \rho_\phi(r^f)$ . Note that the function  $\rho_\phi$  can be any nonlinear function. We use an MLP with one hidden layer which also acts as the projection head for learning the




Figure 2. Inner-workings of the FCRL objective function. We split the context set of each function into  $J$  disjoint views, and align the aggregated representations of those views. Shown here is the example of two functions, with two views each i.e.,  $J = 2$ .

representation. Similar to (Chen et al., 2020), we found that it is beneficial to define the contrastive objective on these projected representations  $\rho_\phi(r^f)$ , rather than directly on the encoded representations  $r^f$ . More details can be found in our ablation study in Appendix A.1.

**Encoder Training.** At training time, we are provided with partial observations  $O^{1:K}$  from  $K$  functions. Each observation is a set of  $N$  i.i.d. (independent and identically distributed) samples  $O^k = \{(x_i^k, y_i^k)\}_{i=1}^N$ . To encourage that different observations of the same functions are mapped to similar representations, we will now formulate a contrastive learning objective, as illustrated in Figure 2. To apply contrastive learning, we create different views of the same function  $k$  by splitting each sample set  $O^k$  into  $J$  subsets of size  $N/J$ . Defining  $t_j := \{(j-1)N/J + 1, \dots, jN/J\}$ , we obtain a split of  $O^k$  into  $J$  disjoint subsets of equal size:

$$O^k = \cup_{j=1}^J O_{t_j}^k \text{ with } O_{t_i}^k \cap O_{t_j}^k = \emptyset \text{ if } i \neq j \quad (3)$$

where each subset  $O_{t_j}^k$  is a partial view of the underlying function  $k$ . Here,  $J$  is a hyper-parameter and can vary in the range of  $[2, N]$  and must divide  $N$ , i.e.  $N \bmod J = 0$ . Its value is empirically chosen based on the data domain: for 1D and 2D regression problems (Sections 4.1), the number of examples per view ( $N/J$ ) is relatively large as a single context point does not provide much information about the underlying function; whereas in scenes (Section 4.3), a few (or even one) images per view (partial observation) may provide enough information. We expand on the appropriate choice of  $J$  in our experiments and the respective ablations. For brevity of notation, we define  $v_j^k := g_{(\phi, \Phi)}(O_{t_j}^k)$ . We now formulate the contrastive learning objective as follows:

$$\mathcal{L} = \sum_{k=1}^K \sum_{1 \leq i < j \leq J} \left[ \log \frac{\exp(\text{sim}(v_j^k, v_i^k) / \tau)}{\sum_{m=1}^K \exp(\text{sim}(v_j^k, v_i^m) / \tau)} \right] \quad (4)$$

where  $\text{sim}(a, b) := \frac{a^\top b}{\|a\| \|b\|}$  is the cosine similarity measure.

Intuitively, the objective function in Equation (4) encourages that the similarity measure  $\text{sim}(v_{(\cdot)}^p, v_{(\cdot)}^q)$  acts as a discriminatory function, yielding a large value if  $v_{(\cdot)}^p$  and  $v_{(\cdot)}^q$  are representations of sets of samples drawn from the same function, i.e. if  $p = q$  (*positives*), and a small value otherwise, i.e. if  $p \neq q$  (*negatives*). The second summation in Equation (4) over  $1 \leq i \leq j \leq J$  is over available pairs of positives, and  $\tau$  is a temperature parameter which scales the scores returned by the similarity measure. Similar to SimCLR (Chen et al., 2020), we find that temperature adjustment is important for learning good representations and treat it as a hyperparameter (ablation study in Appendix A.1).

We note that the learning objective effectively balances two goals. The first is that of avoiding overfitting (i.e., regularization). It encourages that any two independent samples from the same distribution get mapped to similar points. This is akin to the method of “symmetrization by a ghost sample” which is a standard trick in proving learning theory bounds (Vapnik, 1995). Essentially, if two means on different samples are close, then they will also be close to their expectation, i.e., they will not overfit to the data.<sup>1</sup> In spirit, this is close also to the idea of regularization by enforcing stability (i.e., weak dependence on sampling points) (Bousquet & Elisseeff, 2002). The second goal is to preserve contrastive information, ensuring that samples from different distributions get mapped to different points. Both goals are intricately linked in our setting, where the aggregation function is being learnt, since the second component is necessary to prevent the system from trivially meeting the first goal by, say, mapping everything to 0.

### 3.1. Application to Downstream Tasks

Once representation learning using FCRL is concluded,  $h_\Phi$  is fixed and can now be used for few-shot downstream prediction tasks  $\mathcal{T}$  defined on the underlying data-generating functions. To solve a particular downstream task, one may optimize a parametric decoder  $p_\psi(\cdot|r)$  conditioned on the learned representation  $r$ . Specifically, the decoder maps the representations learned in the previous step to the variable of interest in the given task. Depending on the nature of the downstream task, the conditional distributions and the associated objectives can be defined in different ways.

<sup>1</sup>This is an example of the more general phenomenon of concentration of measure, applicable not just to means but also to other functions that aggregate samples. For a simple argument, let  $\mathbb{E}_{O^i}$  denote the expectation w.r.t. drawing the sample  $O^i$ , and  $g$  be the mapping function applied to two independent samples  $O^1, O^2$  from the same distribution. Then we have  $|g(O^1) - \mathbb{E}_{O^1}[g(O^1)]| = |g(O^1) - \mathbb{E}_{O^2}[g(O^2)]| = |\mathbb{E}_{O^2}[g(O^1) - g(O^2)]| \leq \mathbb{E}_{O^2}[|g(O^1) - g(O^2)|]$ . The second equality uses independence of the samples  $O^1$  and  $O^2$ , and the last step uses Jensen’s inequality. This shows that if in expectation the embeddings of two samples are close (r.h.s.), then each embedding is close to its expectation.

## 4. Experiments

To illustrate the benefits of learning function representations without an explicit decoder, we consider four different experimental settings. In all experiments that follow, we first learn the encoder, and then keep it fixed. Subsequently, we optimize decoders for the specific downstream problems at hand, while keeping the meta-representations from the encoder detached from the computational graph.

**Baselines.** We compare the downstream predictive performance of FCRL based representations with the representations learned by the closest task-oriented, meta-learning baselines. For a fair comparison, all the baselines and FCRL have the same encoding architecture. For instance, for 1D and 2D regression functions, we consider CNPs and NPs as baselines. We share with these methods an identical way of mapping the context set to its representation, but unlike us, they optimize directly for the performance of the decoder  $p(y|x)$  jointly with the said representation. For scene datasets, we use GQN (a variant of NPs) as the baseline, one that explicitly learns to reconstruct scenes using a limited number of context samples (pairs of camera viewpoints and the corresponding images).

### 4.1. 1D and 2D Functions

In the first set of experiments, we consider two different distributions of functions i.e., a distribution over 1D sinusoidal waves, proposed by (Finn et al., 2017), and a relatively harder distribution where images are modelled as 2D functions (Garnelo et al., 2018a;b; Gordon et al., 2020; Kim et al., 2019). The representation learning stage for both datasets is similar, however the downstream tasks differ.

**1D Sinusoidal Functions:** We consider a dataset of 20,000 training, 1000 validation and 1000 test sinusoidal functions. The amplitude and the phase of the functions are sampled uniformly from  $[0.1, 0.5]$  and  $[0, \pi]$  respectively. For each function  $f$ , the  $x$ -coordinates are uniformly sampled from  $[-5.0, 5.0]$  and then  $f$  is applied to obtain the  $y$ -coordinates.

**Modeling Images as 2D Functions:** In this setting, each image is regarded as a function mapping from 2D pixel coordinates (comprising function input  $x_i$ ) to the pixel intensities at the corresponding pixel coordinate (comprising function output  $y_i$ ). We consider images of MNIST digits (LeCun et al., 1998), where  $x_i \in [0, 1]^2$  are the normalized pixel coordinates and  $y_i \in [0, 1]$  is a grayscale pixel intensity. The training and validation datasets consists of 60,000 MNIST training and 10,000 test samples, respectively.

#### 4.1.1. REPRESENTATION LEARNING STAGE

We first describe the representation learning stage for both datasets, and then provide results on their respective downstream tasks. For training the encoder  $g_{(\phi, \Phi)}$ , we have a dataset  $O = \{O^k\}_{k=1}^K$  at our disposition, where each  $k$




Figure 3. Qualitative comparison of different methods on 5-shot regression task on three different sinusoids. The decoder trained with FCRL representation predicts the correct form of the sinusoids.

corresponds to a function  $f_k$  which has been sampled as described above. Each individual sample  $O^k = \{(x_i^k, y_i^k)\}_{i=1}^N$  from the dataset is itself a set, comprising  $N$  input-output pairs from that particular function  $f_k$ , i.e.  $y_i^k = f_k(x_i^k)$ . For sinusoidal functions, we fix the maximum number of context points to 20 and the number of examples  $N$  is chosen randomly in  $[2, 20]$  for each  $k$ . For MNIST digits as 2D functions, we allow a maximum of 200 samples per context set, and  $N$  is sampled randomly from  $[2, 200]$  for each  $k$ . The encoder  $g_{(\phi, \Phi)}$  is then trained by splitting each context set  $O^k$  into  $J$  disjoint views. We set  $J = 2$  for the sinusoidal functions and  $J = 10$  for the 2D functions. An ablation study for the choice of  $J$  is presented in Appendix A.1.

#### 4.1.2. DOWNSTREAM TASKS ON 1D SINUSOIDS

After training the encoder  $g_{(\phi, \Phi)}$ , we discard the projection head  $\rho_\phi$  and use the trained encoder  $h_\Phi$  to extract the representations. For 1D sinusoids, we define two downstream tasks on the learned representation:  $\mathcal{T}_{1D} = \{T_{f_{sr}}, T_{f_{spi}}\}$ , where  $T_{f_{sr}}$  and  $T_{f_{spi}}$  are few-shot regression and few-shot parameter identification tasks, respectively. The decoders for the downstream tasks are trained as follows:

**Few-shot Regression (FSR).** FSR for 1D functions is a well-studied problem in meta-learning (Garnelo et al., 2018a; Finn et al., 2017; Kim et al., 2019; Xu et al., 2019). For each sampled function  $f_k$ , we are given a context set  $O^k = \{(x_i^k, y_i^k = f_k(x_i^k))\}_{i=1}^N$  of size  $N$ , which can be utilized to infer the meta-representation  $r^k$  of  $f_k$  via the pre-trained encoder  $h_\Phi$ . We are then provided with  $M$  additional samples from  $f_k$  (not seen by the encoder  $h_\Phi$ ). The goal for a downstream decoder is to predict  $y_i^k$ , given  $x_i^k$  and the meta-representation  $r^k$ . In other words, the downstream decoder with parameters  $\psi$  models the distribution  $p_\psi(y_i^k | x_i^k, r^k)$ . Where  $D^k = \{(x_i^k, y_i^k)\}_{i=1}^{N+M}$ , the decoder is therefore trained to solve the following problem:

$$\max_{\psi} \mathbb{E}_{f_k \sim p(f)} \left[ \mathbb{E}_{(x_i^k, y_i^k) \sim D^k} [\log p_\psi(y_i^k | x_i^k, r^k)] \right] \quad (5)$$

Here, the value of  $M$  is sampled randomly from the interval  $[0, 20 - N]$ . The decoder  $p_\psi$  is an MLP with two hidden

Models	Few-shot Regression (FSR)	
	5-shot	20-shot
NP	0.310 $\pm$ 0.05	0.218 $\pm$ 0.02
CNP	0.265 $\pm$ 0.03	0.149 $\pm$ 0.02
FCRL	<b>0.172 <math>\pm</math> 0.04</b>	<b>0.100 <math>\pm</math> 0.02</b>

Models	Few-shot Parameter Identification	
	5-shot	20-shot
NP	0.0087 $\pm$ 0.0007	0.0037 $\pm$ 0.0005
CNP	0.0096 $\pm$ 0.0007	0.0049 $\pm$ 0.0011
FCRL	<b>0.0078 <math>\pm</math> 0.0004</b>	<b>0.0032 <math>\pm</math> 0.0002</b>

Table 1. Mean squared error (MSE) for all the target points in 5 and 20 shot regression and parameter identification tasks on test sinusoid functions. The reported values are the mean and standard deviation of three independent runs. FCRL performs slightly better than CNP and NP on both tasks.

layers and it is trained with the same training functions as the encoder  $h_\Phi$ . In addition to the Gaussian mean of  $y_i^k$ , it also outputs the variance in order to quantify the uncertainty in the point estimates. The qualitative results on test functions as shown in Figure 3 demonstrate that our model is able to quickly adapt with as few as 5 context points. In Table 1, we compare our method with CNP and NP quantitatively, and show that the predictions of our method are closer to the groundtruth, even though the encoder and decoder in both CNP and NP are explicitly trained to directly maximize the log likelihood to fit the sinusoid.

**Few-shot Parameter Identification (FSPI).** The goal here is to predict the amplitude ( $y_{amp}^k$ ) and phase ( $y_{phase}^k$ ) of the sampled sine wave  $f_k$ , given a context set  $O^k = \{(x_i^k, y_i^k = f_k(x_i^k))\}_{i=1}^N$  of  $N$  samples. Having encoded the context set  $O^k$  to meta-representation  $r^k$  via the pre-trained encoder  $h_\Phi$  (following Equation (2)), we train a linear decoder  $p_\psi$  on top of the said representation by maximizing the likelihood of the sine wave parameters. The predictive distribution is  $p_\psi(y_{amp}^k, y_{phase}^k | r^k)$  and the objective is:

$$\max_{\psi} \mathbb{E}_{f_k \sim p(f)} [\log p_\psi(y_{amp}^k, y_{phase}^k | r^k)] \quad (6)$$

Similar to FSR, we use the same training functions to train  $p_\psi$  as we did to train the encoder  $h_\Phi$ . In Table 1, we report the mean squared error for three independent runs, averaged across all the test tasks for 5-shots and 20-shots FSR and FSPI. In both prediction tasks, the decoders trained on FCRL representations outperform CNP and NP. More details on the experiment are given in Appendix F.

#### 4.1.3. DOWNSTREAM TASKS ON 2D FUNCTIONS

Similar to the tasks above, after training the model  $g_{(\phi, \Phi)}$ , we discard the projection head  $\rho_\phi$  and use the trained en-




Figure 4. (left) Quantitative evaluation of the models in terms of digit classification from the fixed number of context points (varying along the x-axis). The error bands show the standard deviation over three runs. FCRL achieves substantially higher accuracy than both baselines for all evaluated numbers of context points. (right) Quantitative comparison for robustness to noise on MNIST content classification downstream task. The representations learned with FCRL are much more robust to noise than with CNPs and NPs.

coder  $h_\Phi$  to extract the representations. For MNIST digits as functions, we formulate two downstream prediction tasks on the learned representations:  $\mathcal{T}_{2D} = \{T_{fsic}, T_{fscc}\}$ , where  $T_{fsic}$  corresponds to few-shot image completion and  $T_{fscc}$  corresponds to few-shot content classification task. The decoders for the downstream tasks are trained as following.

**Few-shot Content Classification (FSCC).** To evaluate how much semantic information is captured by the meta-representations, we propose the task of few-shot content classification (FSCC). The goal here is to predict the class of each MNIST image given a context set  $O^k = \{(x_i^k, y_i^k)\}_{i=1}^N$  comprising a few randomly sampled pixel coordinates  $x_i^k$  and the corresponding grayscale intensities  $y_i^k$ . The lack of explicit spatial structure in the context points makes it a challenging problem. We use the pre-trained encoder  $h_\Phi$  to encode  $O^k$  to its representation  $r^k$ , and train a linear decoder on top to classify the class label  $y_{one.hot}^k$  corresponding to the MNIST image from which  $O^k$  is sampled. The decoder  $p_\psi$  therefore solves the following classification problem:

$$\max_{\psi} \mathbb{E}_{f_k \sim p(f)} [\log p_\psi(y_{one.hot}^k | r^k)] \quad (7)$$

We train the decoder with the same functions (images) that were used for training the encoder  $h_\Phi$ , and subsequently evaluate them on unseen functions from the validation set. Figure 6 shows the performance of decoders applied to representations obtained from FCRL, CNP and NP for varying size of the context set  $O^k$ . We find that FCRL significantly outperforms the baselines at any given number of context points, suggesting that the encoder  $h_\Phi$  is able to efficiently extract semantic information in an unsupervised manner. We also observe that it is able to generalize to larger number of context points than encountered during training (200).

**Few-shot Image Completion (FSIC).** This setting is identical to that of Few-Shot Regression (FSR) described in Section 4.1.2, except  $M$  is sampled randomly from  $[0, 200 - N]$  and the decoder is a two-layer MLP with two input units (to account for the fact that the input  $x_i^k$  is now two dimensional). Qualitative results of FSIC on test images are shown

Models	Few-shot Image Completion (FSIC)	
	50-shot	200-shot
NP	$0.0531 \pm 0.0002$	$0.0424 \pm 0.0002$
CNP	$0.0477 \pm 0.0006$	$0.0347 \pm 0.0011$
FCRL	$0.0481 \pm 0.0001$	$0.0355 \pm 0.0001$

Table 2. MSE for all the target points in 50 and 200 shot image completion task on MNIST as functions. The reported values are the mean and standard deviation of three independent runs.

in Figure 6. It can be seen that the decoder trained on FCRL representations is able to predict the pixel intensities reasonably well, even when the number of context points is as low as 50, or approximately 6% of the image. We compare its performance against CNP, which uses the same parameterization of both the encoder and the decoder. We however note a crucial distinction: in FCRL, the meta-representation (resulting from the encoder) is not optimized for the image completion task. In particular, no gradient flows from the decoder to the encoder, and the former is trained independently of the latter. On the contrary, CNP jointly optimizes both encoder and decoder parameters to solve the image completion task (i.e. to predict the pixel values). Despite the fact, it appears that the quality of reconstructions from the FCRL decoder matches that from the CNP decoder.

We note that the gap between CNP and FCRL is reduced if the training scheme aligns with the downstream task. In FSIC, the downstream task is to obtain a generative model which is exactly what CNPs are trained for, therefore CNPs tend to perform on par with FCRL as shown in Table 2.

## 4.2. Robustness to Noise Corruptions

In our experiments so far, we have considered the functions to be deterministic. However in real-world settings, data-generating functions are corrupted with noise. In this section, we assume that they take the form:

$$y = f(x, \xi) = f(x) + \xi; \text{ where } \xi \in \mathcal{N}(0, \sigma) \quad (8)$$

where  $\mathcal{N}(0, \sigma)$  is the standard Gaussian distribution with standard deviation  $\sigma$ . We now investigate the robustness of FCRL and the baselines as  $\sigma$  is varied. To this end, we train all the models on the noisy data and evaluate the quality of the learned representation on the Few-Shot Content Classification downstream task, as defined above. We find that the representations learned by FCRL to be significantly more robust to increasing noise strength ( $\sigma$ ) than the baselines, as illustrated in Figure 4. One possible explanation for the susceptibility of CNP and NP to noise is the fact that the representations are learned by reconstructing the outputs, where signal to noise ratio is low. On the other hand, FCRL learns by contrasting the set of examples, extracting only




Figure 5. tSNE projections of the meta-representations learned by FCRL on MPI3D dataset (treating scene as functions). (top row) Each plot shows the latent structure corresponding to the factors mentioned. (bottom row) Shows the latent structure corresponding to the factor of robot arm rotation along 1<sup>st</sup> DOF. Each embedding corresponds to the aggregated representation of three views of the same scene, denoted as  $f^{angle^0}$ . It can be seen that by learning to distinguish between functions, FCRL captures the semantic underlying structure of the functions’ distribution. *Note:* Each plot is generated by varying the factor of interest and keeping the rest of the factors fixed, except the factors of the first degree of freedom and the second degree of freedom.

the invariant features and discarding non-correlated noise in the input. Similar results on scene understanding datasets are presented in Appendix B.




Figure 6. Qualitative comparison of CNP and FCRL based few-shot image completion. Here, each digit corresponds to one function. The context is shown in the second row where the target pixels are colored blue. FCRL is comparable with CNPs at predicting the correct form of a digit.

### 4.3. Representing Scenes as Functions

Like Eslami et al. (2018), we represent scenes as deterministic functions which map camera viewpoints to images. Precisely, each such scene is represented by a function  $f_k$ , and we consider context sets  $O^k = \{(x_i^k, y_i^k)\}_{i=1}^N$ , where  $x_i$  is the 3D camera viewpoint and  $y_i$  is the corresponding image taken from that viewpoint. Given this set of viewpoint-image pairs, we apply the proposed method on  $O^k$  to obtain a representation of the scene,  $r^k$ . The use-

fulness of this representation is then evaluated for three downstream tasks: scene understanding, scene reconstruction and reinforcement learning (RL).

For the first task, our goal is to determine whether the representation  $r^k$  contains enough information to infer the underlying factors of variation (Bengio et al., 2013) of a given scene. To this end, we use MPI3D (Gondal et al., 2019), a real-world robotics dataset comprising pairs of images from three camera viewpoints and the corresponding factors of variation (including the position, orientation, size and color of an object in the scene). For the second task, we analyze whether the learned representation  $r^k$  can be used to reconstruct the scene from an unseen viewpoint. This objective is similar to what GQNs (Eslami et al., 2018) are originally trained for, serving as an ideal baseline for this task. For the last task, our objective is to determine whether  $r^k$  contains enough useful information to guide a RL agent towards maximizing its reward. To this end, we create RLScenes, a multi-view robotics dataset based on an open-source physics simulation engine (details in Appendix D). Having trained the encoder on RLScenes, we feed the representation  $r^k$  of the scene as input to a control policy rewarded for solving the considered RL task.

**Scenes’ Representation Learning Stage.** We use the same setting for learning the representations on both datasets. We fix the maximum number of context sets ( $J$ ) to 3 in MPI3D dataset and 20 in RLScenes. The number of tuples drawn,  $N$ , is then chosen randomly in  $[2, 3]$  and  $[2, 20]$  respectively.

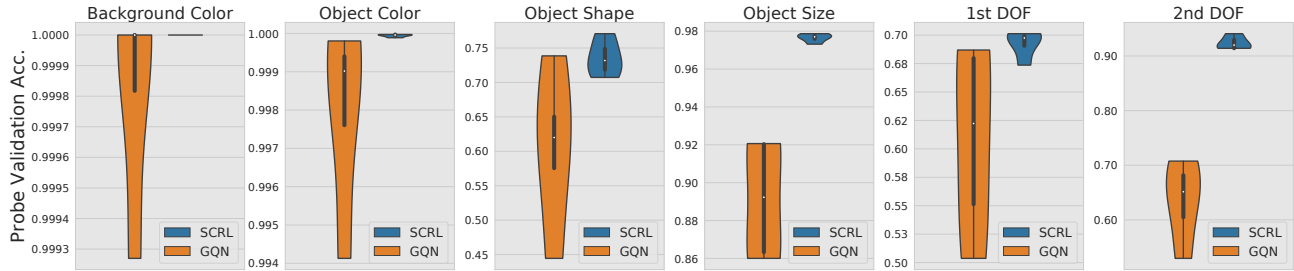


Figure 7. Quantitative Comparison of FCRL and GQN on MPI3D downstream classification tasks. The classifiers trained with FCRL’s representation outperform classifiers based on GQN’s representations on all the tasks.

#### 4.3.1. DOWNSTREAM TASKS ON SCENES

After training the encoder  $g_{(\phi, \Phi)}$ , we discard the projection head  $\rho_{\phi}$  and use the trained encoder  $h_{\Phi}$  to extract the representations  $r^k$  and train decoders for downstream problems.

**Scene Understanding on MPI3D Dataset.** In MPI3D, each scene is identified by 6 factors of variations. This allows us to define a set of 6 tasks  $\mathcal{T}_{mpi3D} = \{T_v^k\}_{v=1}^6$ , where the task  $T_v^k$  maps the scene to a discretized factor of variation  $y_v^k$ . For each task, we train a linear decoder using the objective in Equation (7), using a single image to infer the representation  $r^k$ . Figure 7 shows the linear probes validation performance for six independently trained models on both GQN learned representations and FCRL learned representations and we see that the representations learned by FCRL consistently outperform GQN for identifying all the factors of variations in scenes.

**Scene Reconstruction on MPI3D Dataset.** Similar to the FSIC task in Section 4.1.3, we train a separate decoder to reconstruct the scenes corresponding to an unseen (query) viewpoint  $x_q^k$ . Conditioning on the inferred representation  $r^k$  and the query viewpoint  $x_q^k$ , the decoder reconstructs the corresponding view of the scene  $y_q^k$ . The qualitative comparison in Figure 8 shows that the decoder trained with FCRL representation is capable of preserving the information required to reconstruct the subtle details in the scene.

**Reinforcement Learning on RLScenes Dataset.** In RLScenes, the goal for the agent (a robotic finger) is to locate the object in the arena, reach it, and stay close to it for the remainder of the episode. We use the Soft-Actor-Critic (SAC) algorithm (Haarnoja et al., 2018) to learn a MLP policy for all the joints of the robot, where the policy takes as input the representation  $r^k$  (inferred from a single image) and outputs an action. As the baseline, we use a RL policy trained with GQN representation as input. Figure 9 shows the mean rewards and standard deviations over five runs achieved by both FCRL and GQN-based policies. We find that the FCRL agent clearly outperforms the baseline




Figure 8. Qualitative Comparison of FCRL and GQN on scene reconstruction task. The decoder trained with FCRL’s representation performs on par with GQN in terms of reconstructing a scene from unseen viewpoints.

GQN-agent, both in terms of the final performance and sample-efficiency. In particular, the FCRL agent obtains convergence level control performance with approximately 2 times fewer interactions with the environment.

## 5. Related Work

**Meta-Learning.** Supervised meta-learning can be broadly classified into two main categories. The first category considers the learning algorithm to be an optimizer and meta-learning is about optimizing that optimizer, for e.g., gradient-based methods (Ravi & Larochelle, 2016; Finn et al., 2017; Li et al., 2017; Lee et al., 2019) and metric-learning based methods (Vinyals et al., 2016; Snell et al., 2017; Sung et al., 2018; Allen et al., 2019; Qiao et al., 2019). The second






Figure 9. Comparison between GQN and FCRL on learning a data-efficient control policy for an object reaching downstream task. FCRL based policy clearly outperforms GQN based policy.

category is the family of Neural Processes (NP) (Garnelo et al., 2018a;b; Kim et al., 2019; Eslami et al., 2018) which draw inspirations from Gaussian Processes (GPs). These methods use data-specific priors in order to adapt to a new task at test time while using only a simple encoder-decoder architecture. However, they approximate the distribution over tasks in terms of their predictive distributions which does not incentivize NP to fully utilize the information in the data-specific priors. Our method draws inspiration from this simple, yet elegant framework. However, our proposed method extracts the maximum information from the context which is shown to be useful for solving not just one task, but multiple downstream tasks.

**Self-Supervised Learning.** Self-supervised learning methods aim to learn the meaningful representations of the data by performing some pretext learning tasks (Zhang et al., 2017; Doersch et al., 2015). These methods have recently received huge attention (Oord et al., 2018; Tian et al., 2019; Hjelm et al., 2018; Bachman et al., 2019; Chen et al., 2020; He et al., 2019) mainly owing their success to noise contrastive learning objectives (Gutmann & Hyvärinen, 2010). At the same time, different explanations have recently come out to explain the success of such methods for e.g. from both empirical perspective (Tian et al., 2020; Tschannen et al., 2019) and theoretical perspective (Wang & Isola, 2020; Arora et al., 2019). The goal of these methods has mostly been to extract useful, low-dimensional representation of the data while using downstream performance as a proxy to evaluate the quality of the representation. For example, CPC (Oord et al., 2018) proposes an auto-regressive model to obtain a representation of a sample at time  $t$  that is then matched with that at time  $t + k$ , making it specialized for sequence-valued inputs. On the other hand, (Tian et al., 2019; Chen et al., 2020; He et al., 2019) match the representation of a sample with the representation of its randomly augmented view. In this work, we take inspiration

from these methods and propose a self-supervised learning method which meta-learn the representation of the functions. However, instead of requiring randomly augmented views or sequential ordered data points, our self-supervised loss uses partially observed, unordered views, sampled from the underlying functions.

In a similar spirit of enriching NPs (Garnelo et al., 2018b) with better approximation capability, (Ton et al., 2019) proposed to replace the conditional expectations  $\mathbb{E}(y|x)$  in NPs with more expressive conditional densities  $p(y|x)$  estimated via NCE (Gutmann & Hyvärinen, 2010). In contrast to FCRL, it directly estimates the conditional distribution  $p(y|x)$  and uses a binary classifier for NCE. However, this estimation is done in a small data regime where the standard conditional density estimation does not work well. Therefore, their method is practically limited to low dimensional problems. Recently, (Zhang et al., 2020; Srinivas et al., 2020) has shown the benefits of using self-supervised representations, learned without reconstruction, for reinforcement learning tasks. In this work, we explore the utility of such representations for the reinforcement learning tasks defined on scenes (functions).

## 6. Conclusion

In this work, we proposed a novel self-supervised representations learning algorithm for few-shot learning problems. We deviate from the commonly-used, task-specific training routines in meta-learning frameworks and propose to learn the representations of the relevant functions independently of the prediction task. Experiments on various datasets and the related set of downstream few-shot prediction tasks show the effectiveness of our method. The flexibility to reuse the same representation for different task distributions defined over functions brings us one step closer towards learning a generic meta-learning framework. Using a shared generic representation of the data-generating process, we plan to adapt the proposed framework in order to tackle multiple challenging few-shot problems such as object detection, segmentation, visual question answering.

## ACKNOWLEDGMENTS

The authors would like to thank Krikamol Muandet, Luigi Gresele, Ilya Tolstikhin and Simon Buchholz for the helpful discussions and feedback. We thank CIFAR for their support. This work was supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B, and by the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645.

## References

- Allen, K. R., Shelhamer, E., Shin, H., and Tenenbaum, J. B. Infinite mixture prototypes for few-shot learning. *arXiv preprint arXiv:1902.04552*, 2019.
- Anand, A., Racah, E., Ozair, S., Bengio, Y., Côté, M.-A., and Hjelm, R. D. Unsupervised state representation learning in atari. In *Advances in Neural Information Processing Systems*, pp. 8766–8779, 2019.
- Arora, S., Khandeparkar, H., Khodak, M., Plevrakis, O., and Saunshi, N. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.
- Bachman, P., Hjelm, R. D., and Buchwalter, W. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pp. 15509–15519, 2019.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8): 1798–1828, 2013.
- Bousquet, O. and Elisseeff, A. Stability and generalization. *JMLR*, 2:499–526, 2002.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- Doersch, C., Gupta, A., and Efros, A. A. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pp. 1422–1430, 2015.
- Doersch, C., Gupta, A., and Zisserman, A. Crosstransformers: spatially-aware few-shot transfer. *arXiv preprint arXiv:2007.11498*, 2020.
- Eslami, S. A., Rezende, D. J., Besse, F., Viola, F., Morcos, A. S., Garnelo, M., Ruderman, A., Rusu, A. A., Danihelka, I., Gregor, K., et al. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1126–1135. JMLR. org, 2017.
- Garnelo, M., Rosenbaum, D., Maddison, C. J., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D. J., and Eslami, S. Conditional neural processes. *arXiv preprint arXiv:1807.01613*, 2018a.
- Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S., and Teh, Y. W. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018b.
- Gondal, M. W., Wuthrich, M., Miladinovic, D., Locatello, F., Breidt, M., Volchkov, V., Akpo, J., Bachem, O., Schölkopf, B., and Bauer, S. On the transfer of inductive bias from simulation to the real world: a new disentanglement dataset. In *Advances in Neural Information Processing Systems*, pp. 15740–15751, 2019.
- Gordon, J., Bruinsma, W., Foong, A. Y., Requeima, J., Dubois, Y., and Turner, R. E. Convolutional conditional neural processes. 2020.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 297–304, 2010.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018.
- Joshi, S., Widmaier, F., Agrawal, V., and Wüthrich, M. [https://github.com/open-dynamic-robot-initiative/trifinger\\_simulation](https://github.com/open-dynamic-robot-initiative/trifinger_simulation), 2020.
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. Attentive neural processes. *arXiv preprint arXiv:1901.05761*, 2019.
- Kipf, T., van der Pol, E., and Welling, M. Contrastive learning of structured world models. *arXiv preprint arXiv:1911.12247*, 2019.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lee, K., Maji, S., Ravichandran, A., and Soatto, S. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10657–10665, 2019.
- Li, Z., Zhou, F., Chen, F., and Li, H. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.

- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Qiao, L., Shi, Y., Li, J., Wang, Y., Huang, T., and Tian, Y. Transductive episodic-wise adaptive metric for few-shot learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3603–3612, 2019.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. 2016.
- Snell, J., Swersky, K., and Zemel, R. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pp. 4077–4087, 2017.
- Srinivas, A., Laskin, M., and Abbeel, P. Curl: Contrastive unsupervised representations for reinforcement learning. *arXiv preprint arXiv:2004.04136*, 2020.
- Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P. H., and Hospedales, T. M. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1199–1208, 2018.
- Tian, Y., Krishnan, D., and Isola, P. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., and Isola, P. What makes for good views for contrastive learning. *arXiv preprint arXiv:2005.10243*, 2020.
- Ton, J.-F., Chan, L., Teh, Y. W., and Sejdinovic, D. Noise contrastive meta-learning for conditional density estimation using kernel mean embeddings. *arXiv preprint arXiv:1906.02236*, 2019.
- Tschannen, M., Djolonga, J., Rubenstein, P. K., Gelly, S., and Lucic, M. On mutual information maximization for representation learning. *arXiv preprint arXiv:1907.13625*, 2019.
- Vapnik, V. *The Nature of Statistical Learning Theory*. Springer, NY, 1995.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pp. 3630–3638, 2016.
- Wang, T. and Isola, P. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. *arXiv preprint arXiv:2005.10242*, 2020.
- Xu, J., Ton, J.-F., Kim, H., Kosiorek, A. R., and Teh, Y. W. Metafun: Meta-learning with iterative functional updates. *arXiv preprint arXiv:1912.02738*, 2019.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *Advances in neural information processing systems*, pp. 3391–3401, 2017.
- Zhang, A., McAllister, R., Calandra, R., Gal, Y., and Levine, S. Learning invariant representations for reinforcement learning without reconstruction. *arXiv preprint arXiv:2006.10742*, 2020.
- Zhang, R., Isola, P., and Efros, A. A. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1058–1067, 2017.