
Learning a Universal Template for Few-shot Dataset Generalization

Eleni Triantafillou^{1 2 3} Hugo Larochelle² Richard Zemel¹ Vincent Dumoulin²

Abstract

Few-shot dataset generalization is a challenging variant of the well-studied few-shot classification problem where a diverse training set of several datasets is given, for the purpose of training an adaptable model that can then learn classes from *new datasets* using only a few examples. To this end, we propose to utilize the diverse training set to construct a *universal template*: a partial model that can define a wide array of dataset-specialized models, by plugging in appropriate components. For each new few-shot classification problem, our approach therefore only requires inferring a small number of parameters to insert into the universal template. We design a separate network that produces an initialization of those parameters for each given task, and we then fine-tune its proposed initialization via a few steps of gradient descent. Our approach is more parameter-efficient, scalable and adaptable compared to previous methods, and achieves the state-of-the-art on the challenging Meta-Dataset benchmark.

1. Introduction

While deep learning approaches have recently driven remarkable progress on many important applications, their ability to rapidly learn new concepts from small datasets is significantly lacking. This observation has inspired research towards creating more flexible and adaptable methods. A well-studied problem in this direction is few-shot classification: the problem of utilizing a (possibly large) labeled training set to create an *adaptable* model that is then capable of learning new classes from few examples. Concretely, a few-shot learning model is evaluated on test *tasks*, each of which poses a classification problem between previously-unseen classes, given only a few examples of each.

Departing from early benchmarks for this problem, recent

¹University of Toronto, Vector Institute ²Google Research, Brain Team ³Work done at Google.. Correspondence to: Eleni Triantafillou <eleni@cs.toronto.edu>.

work is shifting towards a more challenging instance of few-shot classification *across datasets* (Chen et al., 2019; Tseng et al., 2020; Triantafillou et al., 2020; Requeima et al., 2019; Bateni et al., 2020; Bronskill et al., 2020; Saikia et al., 2020; Dvornik et al., 2020; Liu et al., 2021). Similarly, we consider the challenging problem of *few-shot dataset generalization*, where the training set is comprised of classes originating from multiple diverse datasets, like ImageNet, Omniglot, Aircraft, and so on, and the aim is to utilize this diverse data towards building a model that can solve new classification tasks between classes of previously-unseen datasets at test time, using only a few examples.

Few-shot dataset generalization exhibits the same challenge of data scarcity as its traditional (single-dataset) few-shot classification counterpart, but it also presents two additional difficulties. Firstly, the heterogeneous nature of the training dataset is reminiscent of *multi-task learning* (Caruana, 1997), where care must be taken to avoid interference when training jointly for different objectives, or different datasets in this case. Secondly, the fact that the test-time tasks are composed of classes from previously-unseen datasets breaks the i.i.d. assumption, resembling the problem of *domain generalization* (Gulrajani & Lopez-Paz, 2020). In a nutshell then, to succeed in few-shot dataset generalization, a method must 1) ingest diverse information without interference, and 2) define a mechanism to appropriately re-purpose that acquired knowledge in order to accommodate extensively different data at test time, using few examples.

In this work, we propose to tackle this problem by learning a *universal template*: a partially-parameterized model trained on multiple datasets in parallel that can be used to define a wide array of dataset-specialized models by providing values for the remaining parameters. We posit that this design incorporates a useful inductive bias for few-shot dataset generalization: in addition to benefiting from knowledge acquired on diverse data, it also reflects the requirement that the template should be flexible enough to serve as the basis to parameterize a diverse set of task-solving networks.

To this end, we propose to train a feature extractor jointly across diverse datasets using FiLM (Perez et al., 2018) in the form of conditional batch normalization: we share the parameters of the convolutional layers across datasets, but allocate a separate set of batch normalization parameters for

each. This joint training regime is key for creating a universal template: by forcing the batch normalization parameters to be fully responsible for dataset specialization, our training objective ensures that the convolutional parameters are *general*, and thus indeed act as a universal template that is able to support generalization to vastly different datasets. To then tackle test-time tasks from unseen datasets, we propose to use the few available labeled examples of each given task to *learn* a new set of batch normalization parameters for that task, starting from a task-dependent initialization: a combination of the per-dataset trained sets of parameters whose co-efficients are proposed by a separate learned network. We refer to this as Few-shot Learning with a Universal TEMplate (FLUTE), illustrated in Figure 1.

We experimentally evaluate FLUTE on few-shot dataset generalization using the recent Meta-Dataset benchmark (Triantafyllou et al., 2020) that is comprised of 10 diverse datasets, 8 of which can be used for training, with the remaining 2 reserved for evaluation. To obtain a richer set of evaluation tasks, we incorporate 3 additional evaluation-only datasets, following Requeima et al. (2019). FLUTE significantly outperforms the state-of-the-art on few-shot dataset generalization on Meta-Dataset (by $> 5\%$), despite having significantly fewer parameters (approx. 8 times fewer).

Aside from its strong performance, FLUTE has several advantages over previous approaches. Firstly, it is more scalable: it’s more parameter efficient, produces more parsimonious representations, and has favorable test-time computational complexity, as summarized in Table 1. It is also very expressive since, for each test task, it blends the per-dataset trained parameters in *all* representation levels (in the batch normalization layers throughout the network), and it is easily adaptable due to its efficient training of a small subset of the feature extractor’s parameters in each test task.

2. Background

Few-shot Classification Let $\mathcal{D}^{tr} = \{(x_i, y_i)\}_{i=1}^{|\mathcal{D}^{tr}|}$ denote a (possibly large) training set, comprised of examples x_i and their corresponding class labels y_i , with each $y_i \in \mathcal{C}^{tr}$. That is, all of the training examples originate from a dedicated training set of classes \mathcal{C}^{tr} , while a disjoint validation set of classes \mathcal{C}^{val} is also available for model selection. The aim is to use \mathcal{D}^{tr} to obtain an *adaptable* model that supports learning new classes from few examples.

Then, evaluation is carried out over a series of classification *tasks*, or *episodes*, designed to measure the trained model’s ability to learn new classes from few examples. Specifically, a k -shot N -way test task presents N classes from a set \mathcal{C}^{test} that is disjoint from \mathcal{C}^{tr} and \mathcal{C}^{val} . We denote each task as a tuple $\mathcal{T} = (\mathcal{S}_{\mathcal{T}}, \mathcal{Q}_{\mathcal{T}})$ of a ‘support set’ and a ‘query set’, representing small ‘within-

task’ training and test sets, respectively. The support set contains k labeled examples of each of the N classes: $\mathcal{S}_{\mathcal{T}} = \{(x_1, y_1), \dots, (x_{kN}, y_{kN})\}$ with each $y_i \in \{1 \dots N\}$, while the query set $\mathcal{Q}_{\mathcal{T}} = \{(x_1^*, y_1^*), \dots\}$ contains different (unlabeled) examples of the same N classes.

Assuming a model with parameters θ , the predicted label \hat{y}^* of a query example x^* is:

$$\hat{y}^* = \operatorname{argmax}_{y^* \in \{1 \dots N\}} \log p_{\theta}(y^* | x^*, \mathcal{S}_{\mathcal{T}})$$

The evaluation metric is then the average query set accuracy over multiple few-shot classification tasks.

Few-shot Dataset Generalization In this challenging instance of few-shot classification, the training and validation sets are defined as the union over M different datasets:

$$\begin{aligned} \mathcal{D}^{tr} &= \mathcal{D}_1^{tr} \cup \mathcal{D}_2^{tr} \dots \cup \mathcal{D}_M^{tr} \\ \mathcal{D}^{val} &= \mathcal{D}_1^{val} \cup \mathcal{D}_2^{val} \dots \cup \mathcal{D}_M^{val} \end{aligned}$$

where \mathcal{D}_m^{tr} and \mathcal{D}_m^{val} contain data from the dedicated training and validation sets of classes of dataset m , respectively. More compactly, we can view the training set (and analogously the validation set too) as $\mathcal{D}^{tr} = \{(x_i, y_i, d_i)\}_{i=1}^{|\mathcal{D}^{tr}|}$, where the x_i ’s and y_i ’s denote examples and their corresponding class labels, as before, and d_i denotes the dataset from which the associated example originates from.

Then, in keeping with the spirit of few-shot classification, the aim at evaluation time is to learn previously-unseen classes from few examples. However, in this more challenging setup, the new classes originate from *new datasets* $\mathcal{D}_{M+1} \dots$, thus inducing a large generalization gap between the training set \mathcal{D}^{train} and the test tasks. The evaluation metric now is the query set accuracy averaged over multiple tasks sampled from multiple held-out datasets.

Feature-wise Linear Modulation (FiLM) FiLM (Dumoulin et al., 2018) is a general-purpose conditioning mechanism that has been used extensively across diverse applications, like question answering, drawing a painting in the style of another, and so on. It performs an affine feature-wise transformation of its input x based on a condition d :

$$\text{FiLM}(x) = \gamma(d) \odot x + \beta(d).$$

The dependency of γ and β on d is handled by maintaining distinct values for each setting of d and selecting the appropriate one for the forward pass of each example.

3. Introducing FLUTE

At a high-level, the core idea behind FLUTE is to use multiple diverse training datasets to construct a *universal template*




Figure 1. (Left) FLUTE trains the shared convolutional weights comprising the universal template, and the M dataset-specific sets of FiLM parameters ($M = 2$ in this illustration) using a multi-task classification loss. (Right) The FiLM parameter values for a test episode are learned through gradient descent using a nearest-centroid classifier (NCC) fitted on the support set as the output layer. The *Blender network* is used to initialize the new FiLM parameters as a convex combination of the trained sets of FiLM parameters.

that can represent a wide array of dataset-specific feature extractors by filling in the shared template with appropriate parameters. During training, FLUTE learns M such sets of dataset-specific parameters, one per dataset in \mathcal{D}^{train} .

We propose to use FiLM on batch normalization layers (also known as conditional batch normalization (De Vries et al., 2017)) as a concrete instantiation of this framework. The shared template therefore consists of the convolutional kernels of a residual network (He et al., 2016), and the dataset-specific parameters to be filled-in are the batch normalization layers’ scaling and shifting coefficients, which we refer to as the *FiLM parameters* from now on.

Our reason for choosing FiLM for this purpose is two-fold. First, while simple, it is very expressive and flexible, as evidenced by its wide success in vastly different problem settings (Dumoulin et al., 2018). This is very important for our goal of generalizing to extensively different data. Second, FiLM layers contain a small fraction of a network’s overall parameters, which is crucial for our few-shot setting where limited data is available for learning how to specialize our template to test tasks of different datasets.

At evaluation time, for each test task, our goal is to produce a new set of FiLM parameters that yields a good feature extractor for that task when plugged into the universal template. We propose to *learn* such parameters through gradient descent on the given task’s support set. In doing so, we also need to specify how to initialize the new FiLM parameters. In practice, we find that initialization really matters, and that simply performing more gradient steps is insufficient to recover from a bad initialization. We propose a simple heuristic that we find works well in practice: initialize a new task’s FiLM parameters from a convex combination of the training datasets’ FiLM parameters where the combination weights are a function of how “compatible” each training dataset is with the new task. We define this “compatibility” to be the output of a dataset classifier that discriminates between the M training datasets. We refer to this approach in what follows using the term *Blender network*.

In what follows, we first describe how we train our universal template, FiLM parameters, and Blender network, and then present our algorithm for tackling few-shot classification tasks from previously-unseen datasets.

3.1. Training the universal and per-dataset parameters

We define our universal template as the convolutional layer parameters of our ResNet-18 feature extractor, and the dataset-specific components take the form of FiLM parameters via different settings for the batch normalization layers.

We train our feature extractor *jointly* over all M training datasets, and we maintain M sets of FiLM parameters for distinguishing the forward passes of examples of different datasets. Formally, let f denote our feature extractor function, and let Φ and Ψ be the shared universal template parameters (the convolutional layer parameters of f) and the dataset-specific parameters, respectively, with Ψ_m (row vector of Ψ) denoting the FiLM parameters for dataset m . Finally, let Ω denote the parameters of per-dataset classification readout heads r_1, \dots, r_M , where the readout layer r_m for dataset m has $|\mathcal{C}_m^{tr}|$ units and is parameterized by Ω_m .

Then, each update to the model is computed as follows: we sample a batch \mathcal{B} of data from the training set $\mathcal{B} = \{(x_i, y_i, d_i)\}_{i=1}^{|\mathcal{B}|}$, with each element i of the batch being a tuple that contains an image x_i , its one-hot-encoded class label y_i , and its one-hot encoded dataset label d_i . We compute the network’s prediction of the class label \hat{y}_i as follows:

$$\hat{y}_i = r_{d_i}(f(x_i; \Phi, \Psi_{d_i}); \Omega_{d_i})$$

We then train Φ , Ψ , and Ω to minimize the cross-entropy loss:

$$-\sum_{i=1}^{|\mathcal{B}|} y_i \log \hat{y}_i$$

Notice that the loss of an example x_i influences only the FiLM parameters corresponding to its originating dataset d_i , but *all* examples influence the shared parameters Φ . We posit that this design is essential for rendering Φ a *universal*

template that is general enough to enable dataset specialization to a wide range of diverse datasets by merely changing the values of the parameter-light batch normalization layers.

After this phase, we no longer require the per-dataset readout heads r_1, \dots, r_M , so we discard them at this point.

3.2. Training the Blender network

The purpose of this network is to create (an initialization for) the FiLM parameters for an unseen dataset d^* via appropriately blending the M sets of FiLM parameters $\{\Psi_m\}_{m=1}^M$. The aim, in particular, is to yield a good feature extractor for d^* by in-filling the universal template with these new FiLM parameters. We define the blending operation as a convex combination whose co-efficient for a particular dataset m is given by an estimate of compatibility between m and d^* .

Towards estimating that compatibility, we propose to train a *dataset classifier*: a network that reads a batch of data and predicts which of the training datasets it was sampled from. The idea is then to interpret this dataset classifier’s probabilities as a proxy for the compatibility between the data it ingested and each training dataset.

Formally, let g denote a network that takes as input a *set* of examples \mathcal{B} and creates a vector representation of that set. We use a permutation-invariant ‘set encoder’ (Zaheer et al., 2017) for g that we describe in the Appendix. Also let l denote a M -dimensional linear readout layer.

Each update to the parameters of g and l is performed by first sampling a training dataset m (an integer between 1 and M), and then sampling \mathcal{B} from \mathcal{D}_m^{tr} . We compute the dataset classifier’s logits as $\hat{d} = l(g(\mathcal{B}))$, and we train the parameters of g and l to minimize the cross-entropy loss for the dataset classification problem.

Once the dataset classifier is trained, we define the blender function as

$$\text{BLEND}(\mathcal{B}) = \text{softmax}(l(g(\mathcal{B})))^T \Psi$$

In words, this is a convex combination of the trained sets of FiLM parameters (the rows of Ψ), where the combination co-efficients are the probabilities of the given batch of data \mathcal{B} of unknown origin belonging to each of the training datasets.

3.3. Tackling few-shot test tasks

Let $\mathcal{T} = \{\mathcal{S}_{\mathcal{T}}, \mathcal{Q}_{\mathcal{T}}\}$ denote a few-shot test task coming from an unseen dataset d^* . Our approach to solving \mathcal{T} can be broken down into two steps. First, we define a feature extractor for \mathcal{T} via an appropriate in-filling of the universal template, and second, we predict a class label for each query example via a simple algorithm that operates on top of the constructed feature extractor.

Step 1: Defining the feature extractor In this step, we seek a set of FiLM parameters Ψ_{d^*} for the given task of unseen dataset d^* that yield a good feature extractor $f(\cdot; \Phi, \Psi_{d^*})$ when plugged into the universal template.

We initialize Ψ_{d^*} to the convex combination proposed by the blender network based on the task’s support set:

$$\Psi_{d^*}^{init} = \text{BLEND}(\mathcal{S}_{\mathcal{T}})$$

Then, we use gradient descent to train Ψ_{d^*} from the initialization $\Psi_{d^*}^{init}$. We next describe the training objective that we use for that purpose, using a Nearest-Centroid Classifier.

Nearest-Centroid Classifier (NCC) The NCC (Mensink et al., 2013; Snell et al., 2017) uses the support set to define a centroid c_j for each class j by averaging the features of the support examples belonging to class j :

$$c_j^{\Psi_{d^*}} = \frac{1}{|\mathcal{S}_{\mathcal{T}}^j|} \sum_{x \in \mathcal{S}_{\mathcal{T}}^j} f(x; \Phi, \Psi_{d^*})$$

where $\mathcal{S}_{\mathcal{T}}^j = \{(x, y) \in \mathcal{S}_{\mathcal{T}} : y = j\}$ is the subset of the support set for class j . The Ψ_{d^*} superscript in $c_j^{\Psi_{d^*}}$ serves to emphasize that the centroid was created from features obtained by plugging in Ψ_{d^*} into the shared template.

Under this classifier, the probability of an example x belonging to class j is proportional to the exponential of the cosine similarity between its features and that class centroid:

$$p^{\Psi_{d^*}}(j|x, \{c_j^{\Psi_{d^*}}\}_{j=1}^N) \propto \exp \frac{f(x; \Phi, \Psi_{d^*}) \cdot c_j^{\Psi_{d^*}}}{\|f(x; \Phi, \Psi_{d^*})\| \cdot \|c_j^{\Psi_{d^*}}\|}$$

Learning the per-task parameters We learn Ψ_{d^*} via gradient descent to minimize the NCC loss on the support set:

$$\Psi_{d^*} = \Psi_{d^*}^{init} - \epsilon \frac{\partial}{\partial \Psi_{d^*}} \left(\sum_{(x,y) \in \mathcal{S}_{\mathcal{T}}} \mathcal{L}^{\Psi_{d^*}}(x, y, \{c_j^{\Psi_{d^*}}\}_{j=1}^N) \right)$$

where ϵ is the learning rate and $\mathcal{L}^{\Psi_{d^*}}(x, y, \{c_j^{\Psi_{d^*}}\}_{j=1}^N)$ denotes the cross-entropy loss for example (x, y) under the NCC. In practice, we might take more than a one step of gradient descent; we treat that as a hyperparameter. Note that all parameters aside from Ψ_{d^*} are frozen at this time.

We also emphasize that we learn Ψ_{d^*} on a per-task basis at test time, using the task’s support set. That is, Ψ_{d^*} is not re-used across the test tasks that are sampled from d^* .

Step 2: Predict a class label for each query example Finally, having defined a feature extractor $f(\cdot; \Phi, \Psi_{d^*})$ for the given test task, we can utilize the NCC to classify each query example. Concretely, we classify a query x^* as:

$$\hat{y}^* = \underset{j}{\text{argmax}} p^{\Psi_{d^*}}(j|x^*, \{c_j^{\Psi_{d^*}}\}_{j=1}^N)$$

4. Related Work

Few-shot dataset generalization is related to various well-studied areas. The aim in multi-task learning (Crawshaw, 2020), for example, is to learn multiple tasks simultaneously with a shared model, similar to FLUTE’s approach for training on multiple diverse datasets jointly. However, contrary to that problem, our goal is to generalize to *new datasets*. Domain generalization (Gulrajani & Lopez-Paz, 2020) also relates to our problem setting due to the shared goal of learning a model that generalizes to vastly different data distributions, but it typically assumes that the label space is shared across datasets, contrary to our setting.

Another related area is transfer learning (Zhuang et al., 2020), aiming to improve a learner’s performance on target domains via leveraging knowledge obtained from related source domains. A representative approach that is related to FLUTE is that of Puigcerver et al. (2021): they first train a set of diverse ‘experts’ on different data and then, for each downstream task, they select the most appropriate expert and fine-tune it on that task. FLUTE differs in that, thanks to its universal template that can be re-used across tasks of different datasets, it only needs to fine-tune a small number of per-task parameters at test time (the FiLM parameters); a design that allows to effectively leverage previous knowledge and avoid overfitting in our data-scarce setting.

Moreover, FLUTE is related to several few-shot learning methods. Many popular ones are metric-based (Vinyals et al., 2016; Snell et al., 2017; Sung et al., 2018), or learning-to-learn approaches (Ravi & Larochelle, 2017; Finn et al., 2017). A particular variant of the influential Model Agnostic Meta Learning (MAML) model that is especially related to FLUTE is CAVIA (Zintgraf et al., 2019) that also divides the feature extractor’s parameters into task-general and task-specific ones. At test time, only the task-specific ones are adapted for each new episode. Moreover, (Fleenerhag et al., 2019) propose to interleave meta-learned ‘warp-layers’ between the layers of a task learner, in the context of gradient-based meta-learning. Reminiscent of FLUTE’s template, their warp-layers are shared across tasks, with the aim in that case being to efficiently parameterize a preconditioning matrix that facilitates within-task learning. Albeit related, these previous approaches do not target our across-dataset problem setting of interest, and differ algorithmically as they utilize a meta-learning approach, whereas FLUTE is trained across datasets, with the simple training objective of joint dataset classification. Keeping with the spirit of separating shared from specific weights, (Wang et al., 2019) design a model for few-shot object detection based on the insight of disentangling category-specific and category-agnostic parameters, in a similar spirit to FLUTE’s dataset-specific and dataset-agnostic parameters. However, their approach differs significantly on an architectural level, as well as in terms

of their training objective and overall goal and application area of the work.

FiLM has also been used in several few-shot learning approaches, typically for the purpose of conditioning on each episode’s support set to construct a set of features that are tailored to the task at hand (Oreshkin et al., 2018; Requeima et al., 2019; Bateni et al., 2020). FLUTE differs from these works in that, consistent with its goal of constructing a dataset-general template, it conditions on the dataset label, instead of conditioning on a continuous representation of support-set features as those previous works. This idea relates to (Rebuffi et al., 2017)’s work that introduce ‘residual adapters’, a generalization of FiLM, to condition on different datasets. However, that work does not consider the problem that is our focus: generalization to new datasets at test time.

The idea of using a dataset classifier in the Blender network to guide the initialization of new sets of FiLM parameters is reminiscent of (Torralba & Efros, 2011)’s earlier work, where they noticed that it is possible to accurately classify datasets, suggesting that each has a unique “signature”.

Finally, perhaps the most closely-related family of work is that of building a *universal representation*; a term coined in Bilen & Vedaldi (2017) to describe a set of rich features that enable strong performance on several different datasets in a multi-task learning setup. Building on this idea, the approach of Selecting from a Universal Representation (SUR (Dvornik et al., 2020)) achieved strong performance by training a separate feature extractor on each training dataset (for a total of M extractors). To tackle each test task, they obtain the M sets of extracted features for the task’s data, and concatenate them to form a universal representation of size MD , where D is the representation size of each feature extractor. They then utilize a selection mechanism that weights these MD features based on relevance to the task at hand. Liu et al. (2021) later improved upon SUR by replacing its feature weighting mechanism by a meta-learned attention layer referred to as the Universal Representation Transformer layer (URT). To alleviate the cost of training M separate feature extractors, SUR and URT offer ‘parametric family’ variants, SUR-pf and URT-pf, respectively, where their M feature extractors share all but their batch normalization weights; a design that is architecturally similar to FLUTE’s. Unfortunately, though, the compactness of SUR-pf/URT-pf results in a significant drop in accuracy.

Compared to those works, FLUTE embodies a different inductive bias: we don’t seek a set of universal features from which we select dataset-relevant ones. Instead, we seek a universal template that defines several dataset-specialized feature extractors, each of which only produces dataset-relevant features. An additional difference is that SUR-pf/URT-pf train the shared parameters only on ImageNet,

whereas we employ a joint training phase that permits all datasets to influence those parameters; a choice that is crucial for turning the shared parameters into a general-purpose universal template. Finally, FLUTE scales better with the number of training datasets, as summarized in Table 1.

Table 1. Comparing FLUTE to universal representation methods in terms of scalability with the number M of training datasets. P is the number of feature extractor parameters, D the size of representation that the feature extractor outputs (e.g. 512 for a ResNet-18), H the number of URT’s attention heads and T the number of FLUTE’s within-task training steps to learn Ψ_h during evaluation. ‘Computation’ measures the number of passes through the feature extractor(s) for each task. SUR(-pf) and URT(-pf) forward-pass through their M extractors (some of which share parameters in the pf variants) whereas we only maintain a single feature extractor. We perform a forward and a backward pass on the support set for each of our T steps, and a final forward pass on the query set ($2T + 1$ in total); thus FLUTE’s amount of test-time computation does *not* grow with the number of training datasets as for SUR(-pf) and URT(-pf), and T can be small in practice.

	Num params	Size of repr.	Computation
SUR	MP	MD	M
URT	MP	HD	M
SUR-pf	P	MD	M
URT-pf	P	HD	M
FLUTE	P	D	$2T + 1$

5. Experimental Setup

We conduct our experimental evaluation on the recent Meta-Dataset benchmark, which lends itself well to studying few-shot dataset generalization, as its designated training set is indeed comprised of diverse datasets, and its test set includes classes from previously-unseen datasets too.

Meta-Dataset In more detail, the training set contains classes from ImageNet, Omniglot, Aircraft, Birds, Flowers, Quickdraw, Fungi, and Textures, some of which contain natural images (e.g. ImageNet, Birds, and Flowers), while others differ significantly in appearance. For example, Omniglot and Quickdraw depict hand drawings on a plain black background and Textures’s images present perceptual characteristics of varying structure and don’t feature a single object in the center of the image. The test set introduces two new datasets: Traffic Signs and MSCOCO, with the former differing significantly thematically from the training set (it depicts images of different traffic signs), and the latter, albeit thematically similar to ImageNet, has lower resolution and exhibits inclusion. Since the aim of this work is to study few-shot dataset generalization, we incorporate the 3 additional test datasets introduced in Requeima et al. (2019),

to obtain a larger and more varied test set for this problem. These datasets are MNIST, CIFAR-10, and CIFAR-100.

Strong versus weak generalization The problem of few-shot dataset generalization requires *strong generalization* as there is a large generalization gap between the training and test distributions. Another popular, albeit easier, setting creates the test tasks from held-out classes of the datasets that were encountered at training time. This still falls within the umbrella of few-shot classification, as the test classes were not used for training, but is an easier problem since the generalization gap is inherently smaller. We thus refer to this setting as *weak generalization*. Although our focus is few-shot dataset generalization in this work, we also examine FLUTE’s performance on the weak generalization tasks for completeness and compatibility with previous work.

We apply FLUTE in exactly the same way for this setting as described in Section 3. In particular, despite the test classes originating from seen datasets in this setting, we are not told which of the M datasets each task is sampled from. Therefore, we still require our Blender network to propose a blending of the M sets of FiLM parameters as the initialization of Ψ_{d^*} , although we expect that it should almost exclusively select the correct dataset from which the task originates, assuming the dataset classifier is accurate.

Implementation details Following the prior work, we use a ResNet-18 as the feature extractor. During training of our universal template, we treat each dataset-specific read-out head r_m as a cosine classifier (as in Chen et al. (2019; 2020)): a linear layer without a bias, where both the layer inputs and the rows of the weight matrix are l2-normalized. We use stochastic gradient descent with momentum as the optimizer for this phase, with a cosine decay with restarts schedule for the learning rate. On the other hand, the second phase of training for our dataset classifier network converges quite quickly (around 14K steps). The design of the dataset classifier involves a deep set encoder, comprised by 5 convolutional layers, each followed by batch normalization and ReLU, with a standard linear readout on top. We use Adam for this phase. We report all details in the Appendix. Our code is publicly available and has been incorporated into the Meta-Dataset codebase¹.

6. Results

Performance on Meta-Dataset As our first experiment, we evaluate FLUTE on Meta-Dataset, in order to assess its performance on the challenging problem of few-shot dataset generalization, as well as on the weaker generalization setup. We compare against previous approaches that have the same number of parameters as FLUTE, as well as the state-of-

¹<https://github.com/google-research/meta-dataset>

Table 2. Comparison of FLUTE to the previous state-of-the-art approaches on Meta-Dataset. The few-shot dataset generalization performance is shown in the second group of rows (Traffic Signs - CIFAR-100), representing unseen datasets that require Strong Generalization (SG). For completeness, we also present results on the easier problem of Weak Generalization (WG) in the first 8 rows (ImageNet - Flower). We compare against CNAPs (Requeima et al., 2019), TaskNorm (Bronskill et al., 2020)³, SimpleCNAPs (Bateni et al., 2020), SUR(-pf) (Dvornik et al., 2020) and URT(-pf) (Liu et al., 2021). We emphasize that SUR and URT have 8 times more parameters compared to all other competitors (denoted by x8). FLUTE outperforms these previous methods by a significant margin for the problem of few-shot dataset generalization (> 5% on Average SG) while still outperforming the state-of-the-art overall. Each number represents the average query set accuracy over 600 test tasks, and its 95% confidence interval.

Dataset	CNAPs	TaskNorm	SimpleCNAPs	SUR-pf	URT-pf	SUR (x8)	URT (x8)	FLUTE
ImageNet	52.3 ± 1.0	50.6 ± 1.1	58.6 ± 1.1	56.4 ± 1.2	55.5 ± 1.1	56.3 ± 1.1	55.7 ± 1.1	51.8 ± 1.1
Omniglot	88.4 ± 0.7	90.7 ± 0.6	91.7 ± 0.6	88.5 ± 0.8	90.2 ± 0.6	93.1 ± 0.5	94.4 ± 0.4	93.2 ± 0.5
Aircraft	80.5 ± 0.6	83.8 ± 0.6	82.4 ± 0.7	79.5 ± 0.8	79.8 ± 0.7	85.4 ± 0.7	85.8 ± 0.6	87.2 ± 0.5
Birds	72.2 ± 0.9	74.6 ± 0.8	74.9 ± 0.8	76.4 ± 0.9	77.5 ± 0.8	71.4 ± 1.0	76.3 ± 0.8	79.2 ± 0.8
Textures	58.3 ± 0.7	62.1 ± 0.7	67.8 ± 0.8	73.1 ± 0.7	73.5 ± 0.7	71.5 ± 0.8	71.8 ± 0.7	68.8 ± 0.8
Quickdraw	72.5 ± 0.8	74.8 ± 0.7	77.7 ± 0.7	75.7 ± 0.7	75.8 ± 0.7	81.3 ± 0.6	82.5 ± 0.6	79.5 ± 0.7
Fungi	47.4 ± 1.0	48.7 ± 1.0	46.9 ± 1.0	48.2 ± 0.9	48.1 ± 0.9	63.1 ± 1.0	63.5 ± 1.0	58.1 ± 1.1
Flower	86.0 ± 0.5	89.6 ± 0.6	90.7 ± 0.5	90.6 ± 0.5	91.9 ± 0.5	82.8 ± 0.7	88.2 ± 0.6	91.6 ± 0.6
Traffic Signs	56.5 ± 1.1	-	59.2 ± 1.0	52.2 ± 0.8	52.0 ± 1.4	53.4 ± 1.0	51.1 ± 1.1	58.4 ± 1.1
MSCOCO	42.6 ± 1.1	43.4 ± 1.0	46.2 ± 1.1	52.1 ± 1.0	52.1 ± 1.0	52.4 ± 1.1	52.2 ± 1.1	50.0 ± 1.0
MNIST	92.7 ± 0.4	92.3 ± 0.4	93.9 ± 0.4	93.2 ± 0.4	93.9 ± 0.4	94.3 ± 0.4	94.8 ± 0.4	95.6 ± 0.5
CIFAR-10	61.5 ± 0.7	69.3 ± 0.8	74.3 ± 0.7	66.4 ± 0.8	66.1 ± 0.8	66.8 ± 0.9	67.3 ± 0.8	78.6 ± 0.7
CIFAR-100	50.1 ± 1.0	54.6 ± 1.1	60.5 ± 1.0	57.1 ± 1.0	57.3 ± 1.0	56.6 ± 1.0	56.9 ± 1.0	67.1 ± 1.0
Average WG	69.7	71.9	73.8	73.6	74.0	75.6	77.3	76.2
Average SG	60.7	-	66.8	64.2	64.3	64.7	64.5	69.9
Average all	66.2	-	71.1	70.0	70.3	71.4	72.3	73.8

the-art models on Meta-Dataset that use roughly 8 times more parameters than FLUTE. We present these results in Table 2. We observe that FLUTE significantly outperforms all approaches on few-shot dataset generalization. Overall, it sets the new state-of-the-art on Meta-Dataset, despite FLUTE’s compactness and parameter efficiency.

Having established FLUTE’s effectiveness, we next take a closer look at the effect of its various design choices.

The importance of the Blender network Recall that the role of this network is to propose an initialization for the FiLM parameters for each test task. We now explore two alternative ways of initializing those parameters: ‘from scratch’ and ‘from ImageNet’ (Ψ_{IN}). For each layer i , the former sets the additive terms β_i to vectors of 0s, and the multiplicative terms γ_i to vectors of 1s, while the latter sets the FiLM parameters throughout the network to those learned specifically for the ImageNet dataset, which we denote by Ψ_{IN} . We singled out ImageNet for this since it is often treated as a good source for transfer.

The results for this experiment are shown in the first three columns of Table 3. We find that initializing from Ψ_{IN} yields impressive results on several datasets (e.g. ImageNet (unsurprisingly), but also Textures, CIFAR-10, CIFAR-100)

³We don’t report Traffic Signs results for this method due to an inconsistency in how evaluation on this dataset was performed (described in the Appendix)

but very poor results on datasets that are vastly different from ImageNet (e.g. Omniglot, Quickdraw, and Aircraft). On the other hand, initializing from scratch does not recover that strong performance on the ImageNet-like datasets, but evidently constitutes a much better solution for various other datasets (e.g. Omniglot, Quickdraw and Aircraft). Therefore, the initialization of Ψ_{d^*} really matters, with different techniques working best for different evaluation datasets. Our Blender network thus provides a good solution by tailoring the initialization to each task at hand, via the predicted compatibility with each training dataset.

The importance of diverse training data Next, we empirically examine our hypothesis that training our feature extractor on *all* training datasets is an important choice for encouraging (the convolutional layers of) our feature extractor to act as a universal template. In other words, had we not trained it on diverse data, we hypothesize it would be less general and thus less capable of supporting generalization to new datasets via merely changing the values of the batch normalization parameters. To investigate this, we trained a feature extractor (of identical architecture to FLUTE’s) on ImageNet only. In this case, we only learn one set of FiLM parameters during training, since there is only one training dataset. We then attempt to tackle the test tasks similarly to FLUTE, by learning a new set of FiLM parameters for each, via one of the three initialization schemes proposed in the previous round of experiments.




Figure 2. The support and query accuracy over 600 test episodes of unseen datasets as a function of the fine-tuning steps for Ψ_{d^*} .

Table 3. The effect of training on different data (‘All’, as in FLUTE, or ‘ImageNet-only’), and alternative initialization schemes for Ψ_{d^*} : from scratch (‘Scratch’), from ImageNet’s FiLM parameters (‘ Ψ_{IN} ’) and from Blender. The column corresponding to the setting “All, Blender” is our proposed FLUTE model. Each number represents the average query set accuracy over 600 test tasks, and its 95% confidence interval. We include confidence intervals for this table in the Appendix.

Training data Init scheme	All			ImageNet only		
	Scratch	Ψ_{IN}	Blender	Scratch	Ψ_{IN}	Blender
ImageNet	47.7	53.9	51.8	34.2	46.9	46.9
Omniglot	91.2	75.6	93.2	57.2	61.6	77.9
Aircraft	80.6	66.0	87.2	35.3	48.5	67.3
Birds	72.0	73.2	79.2	28.8	47.9	60.0
Textures	69.8	71.9	68.8	55.5	63.8	61.2
Quickdraw	78.1	69.3	79.5	50.6	57.5	60.9
Fungi	51.6	46.0	58.1	23.0	31.8	34.3
Flower	91.0	89.2	91.6	65.9	80.1	73.8
Traffic Signs	56.2	54.8	58.4	36.4	46.5	43.0
MSCOCO	42.5	50.6	50.0	29.3	41.4	41.3
MNIST	95.5	83.4	95.6	78.9	80.8	86.9
CIFAR-10	69.6	76.9	78.6	47.1	65.4	65.4
CIFAR-100	58.0	67.5	67.1	34.5	52.7	52.2
Average WG	72.8	68.1	76.2	43.8	54.8	60.3
Average SG	64.4	66.7	69.9	45.2	57.4	57.8
Average all	69.5	67.6	73.8	44.4	55.8	59.3

In order to be able to use the Blender heuristic with this feature extractor too, we require a set of FiLM parameters for each of the training datasets. For the ImageNet dataset, we already have this: the FiLM parameters correspond to the batch normalization parameters obtained when training this feature extractor on ImageNet. For each of the remaining training datasets, we launch a training phase to learn a new set of batch normalization parameters, keeping the convolutional parameters frozen to those learned by this ImageNet-only feature extractor. This is reminiscent of SUR-pf’s procedure for learning batch normalization parameters for the different training datasets, for use with an ImageNet-trained feature extractor. Once we have obtained the FiLM parameters for each of the training datasets for this ImageNet feature extractor, we use our dataset classifier within Blender to obtain the initialization of Ψ_{d^*} as usual.

The performance of this ImageNet-only feature extractor, with different initialization heuristics for Ψ_{d^*} , are shown

in the last three columns of Table 3. We find that, for any given initialization heuristic, the feature extractor trained on all datasets performs significantly better across the board compared to that trained on ImageNet only, validating our hypothesis. Notably, the ‘Scratch’ heuristic no longer performs well for datasets like Omniglot and Quickdraw, as it did when using FLUTE’s feature extractor. This speaks to the inferior generalizability of this feature extractor that was trained on less diverse data. Finally, we observe that even when evaluating on test classes from ImageNet (first row), FLUTE’s feature extractor evidently constitutes a better template compared to the ImageNet-trained one, indicating positive transfer during FLUTE’s joint training phase.

The effect of the number of steps to train Ψ_{d^*} Figure 2 plots the support and query accuracy during the fine-tuning of the per-task parameters Ψ_{d^*} on the support set, averaged over 600 test tasks of each unseen domain (strong generalization setting). We also show the corresponding plots for weak generalization in the Appendix. We used 6 steps for this, with a learning rate of 0.005, which are the values we used for FLUTE’s results in Table 2, chosen based on the validation set. We observe that some datasets benefit more than others from this fine-tuning. In fact, in some cases (CIFAR-100, MSCOCO) the support accuracy even decreases. This indicates that perhaps those hyperparameters aren’t ideal for all test tasks. Moreover, the validation set may not be a good proxy for the test set, since it is comprised of different datasets, making hyperparameter selection hard. Tailoring the learning rate or number of steps to each task would be an interesting avenue for future work.

Inspecting the Blender’s proposals We visualize the combination coefficients that the Blender produces for test tasks in Figure 3. We notice that for test episodes of datasets seen during training (weak generalization setting), the Blender very strongly selects the dataset from which the task originated. For new datasets that did not appear in the training set (last 5 columns), the Blender seems to make reasonable choices (e.g. Quickdraw for MNIST, ImageNet for MSCOCO) but the probability distribution may be less peaky (e.g. for Traffic Signs and CIFAR-100). In the Appendix we also investigate the variance of these predictions across different test tasks within each dataset.




Figure 3. Visualization of the coefficients of the convex combination produced by the Blender network for various test datasets. Each column represents the average coefficients produced for 600 test tasks of a particular dataset.

The effect of taking a convex combination Having inspected the combination weights that Blender produces, we noticed that it often relies almost exclusively on a single dataset’s FiLM parameters to form the initialization of the FiLM parameters of a new task. A natural question to ask, therefore, is whether the convex combination is required. An alternative design choice is a ‘Hard Blender’ that selects only the FiLM parameters of the most compatible training dataset (as assessed by the dataset classifier), instead of combining them all in a weighted manner. We provide a detailed comparison with this variant in the Appendix. Indeed, the two initialization schemes perform similarly, especially on weak generalization tasks, where the Blender’s co-efficients already resemble a hard selection. We observe, though, that the combination approach may yield small gains on strong generalization tasks. Overall, since Blender is a more general approach and it may improve the performance on new datasets, we adopt this initialization scheme as our default one for FLUTE.

Discussion As alluded to earlier, an open challenge is defining an appropriate model selection mechanism for the difficult scenario when the provided validation set is not a good proxy for the test set. The problem that we study in this work, similarly to domain generalization (Gulrajani & Lopez-Paz, 2020), suffers from underspecification: there may be several solutions that all perform well on the training and validation sets but behave differently on the test set. With this in mind, perhaps we should adjust our evaluation metrics to account for this potential source of variance. For instance, in the context of FLUTE, we noticed in Table 3 that the initialization of Ψ_{d^*} really matters, therefore perhaps the results are sensitive to the chosen checkpoint of the dataset classifier. We report in the Appendix the performance across more runs of that component and observe some variance on

held-out datasets, but our conclusions hold across runs.

7. Conclusion

To conclude, we proposed FLUTE, a method for few-shot dataset generalization that creates a universal template, i.e. a partially-parameterized model, that can serve to readily specify a good feature extractor for a held-out dataset, by simply plugging in an appropriate set of FiLM parameters. We propose an efficient approach for estimating those parameters, using a few steps of gradient descent starting from a task-dependent initialization. FLUTE sets the new state-of-the-art on Meta-Dataset, significantly outperforming previous methods on few-shot dataset generalization.

References

- Batani, P., Goyal, R., Masrani, V., Wood, F., and Sigal, L. Improved few-shot visual classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14493–14502, 2020.
- Bilen, H. and Vedaldi, A. Universal representations: The missing link between faces, text, planktons, and cat breeds. *arXiv preprint arXiv:1701.07275*, 2017.
- Bronskill, J., Gordon, J., Requeima, J., Nowozin, S., and Turner, R. E. Tasknorm: Rethinking batch normalization for meta-learning. In *Proceedings of the International Conference on Machine Learning*, 2020.
- Caruana, R. Multitask learning. *Machine learning*, 28(1): 41–75, 1997.
- Chen, W.-Y., Liu, Y.-C., Kira, Z., Wang, Y.-C. F., and Huang, J.-B. A closer look at few-shot classification. In *Proceedings of the International Conference on Learning Representations*, 2019.
- Chen, Y., Wang, X., Liu, Z., Xu, H., and Darrell, T. A new meta-baseline for few-shot learning. *arXiv preprint arXiv:2003.04390*, 2020.
- Crawshaw, M. Multi-task learning with deep neural networks: A survey. *arXiv preprint arXiv:2009.09796*, 2020.
- De Vries, H., Strub, F., Mary, J., Larochelle, H., Pietquin, O., and Courville, A. C. Modulating early visual processing by language. In *Advances in Neural Information Processing Systems*, pp. 6594–6604, 2017.
- Dumoulin, V., Perez, E., Schucher, N., Strub, F., Vries, H. d., Courville, A., and Bengio, Y. Feature-wise transformations. *Distill*, 2018.
- Dvornik, N., Schmid, C., and Mairal, J. Selecting relevant features from a multi-domain representation for few-shot

- classification. In *European Conference on Computer Vision*, pp. 769–786, 2020.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the International Conference on Machine Learning*, pp. 1126–1135, 2017.
- Flennerhag, S., Rusu, A. A., Pascanu, R., Visin, F., Yin, H., and Hadsell, R. Meta-learning with warped gradient descent. *arXiv preprint arXiv:1909.00025*, 2019.
- Gulrajani, I. and Lopez-Paz, D. In search of lost domain generalization. In *(to appear) Proceedings of the International Conference on Learning Representations*, 2020.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Liu, L., Hamilton, W., Long, G., Jiang, J., and Larochelle, H. A universal representation transformer layer for few-shot image classification. In *(to appear) Proceedings of the International Conference on Learning Representations*, 2021.
- Mensink, T., Verbeek, J., Perronnin, F., and Csurka, G. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2624–2637, 2013.
- Oreshkin, B., López, P. R., and Lacoste, A. Tadam: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 721–731, 2018.
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., and Courville, A. FiLM: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- Puigcerver, J., Riquelme, C., Mustafa, B., Renggli, C., Pinto, A. S., Gelly, S., Keysers, D., and Houlsby, N. Scalable transfer learning with expert models. In *(to appear) Proceedings of the International Conference on Learning Representations*, 2021.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. In *Proceedings of the International Conference on Learning Representations*, 2017.
- Rebuffi, S.-A., Bilen, H., and Vedaldi, A. Learning multiple visual domains with residual adapters. *arXiv preprint arXiv:1705.08045*, 2017.
- Requeima, J., Gordon, J., Bronskill, J., Nowozin, S., and Turner, R. E. Fast and flexible multi-task classification using conditional neural adaptive processes. In *Advances in Neural Information Processing Systems*, pp. 7957–7968, 2019.
- Saikia, T., Brox, T., and Schmid, C. Optimized generic feature learning for few-shot classification across domains. *arXiv preprint arXiv:2001.07926*, 2020.
- Snell, J., Swersky, K., and Zemel, R. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pp. 4077–4087, 2017.
- 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.
- Torralba, A. and Efros, A. A. Unbiased look at dataset bias. In *CVPR 2011*, pp. 1521–1528. IEEE, 2011.
- Triantafillou, E., Zhu, T., Dumoulin, V., Lamblin, P., Evci, U., Xu, K., Goroshin, R., Gelada, C., Swersky, K., Manzagol, P.-A., et al. Meta-Dataset: A dataset of datasets for learning to learn from few examples. In *Proceedings of the International Conference on Learning Representations*, 2020.
- Tseng, H.-Y., Lee, H.-Y., Huang, J.-B., and Yang, M.-H. Cross-domain few-shot classification via learned feature-wise transformation. In *Proceedings of the International Conference on Learning Representations*, 2020.
- 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, Y.-X., Ramanan, D., and Hebert, M. Meta-learning to detect rare objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9925–9934, 2019.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R., and Smola, A. Deep sets. In *Advances in neural information processing systems*, 2017.
- Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., and He, Q. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020.
- Zintgraf, L., Shiarli, K., Kurin, V., Hofmann, K., and Whiteson, S. Fast context adaptation via meta-learning. In *International Conference on Machine Learning*, pp. 7693–7702. PMLR, 2019.