

CORL: Compositional Representation Learning for Few-Shot Classification

Ju He¹ Adam Kortylewski^{1,2,3} Alan Yuille¹

¹Johns Hopkins University ²Max Planck Institute for Informatics ³University of Freiburg

Abstract

Few-shot image classification consists of two consecutive learning processes: 1) In the meta-learning stage, the model acquires a knowledge base from a set of training classes. 2) During meta-testing, the acquired knowledge is used to recognize unseen classes from very few examples. Inspired by the compositional representation of objects in humans, we train a neural network architecture that explicitly represents objects as a dictionary of shared components and their spatial composition. In particular, during meta-learning, we train a knowledge base that consists of a dictionary of component representations and a dictionary of component activation maps that encode common spatial activation patterns of components. The elements of both dictionaries are shared among the training classes. During meta-testing, the representation of unseen classes is learned using the component representations and the component activation maps from the knowledge base. Finally, an attention mechanism is used to strengthen those components that are most important for each category. We demonstrate the value of our interpretable compositional learning framework for a few-shot classification using miniImageNet, tieredImageNet, CIFAR-FS, and FC100, where we achieve comparable performance.

1. Introduction

Advances in the architecture design of deep convolutional neural networks (DCNNs) [18, 31, 11] increased the performance of computer vision systems at image classification enormously. However, in practice, their performance is usually limited when not enough labeled data is available. Few-shot classification is concerned with the problem of learning from a small number of samples. In particular, it consists of two consecutive learning processes: 1) In the meta-learning stage, the model acquires a knowledge base from a set of training classes. 2) During meta-testing, the acquired knowledge is used to recognize unseen classes from very few examples. Hence, few-shot classification wants to emulate human learning efficiency [19, 2, 6, 13] by requiring to transfer the knowledge gained through training

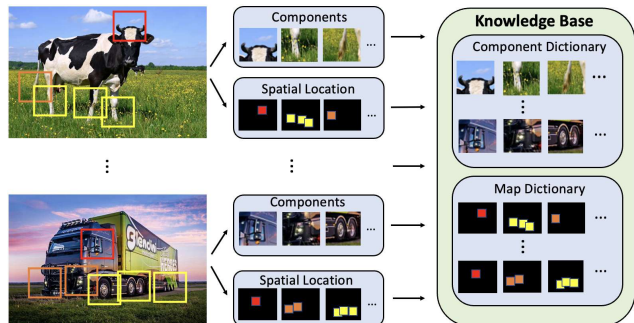


Figure 1: Intuitive illustration of how our model acquires knowledge during meta-learning. In particular, it learns a dictionary of component representations, which resemble individual object components. Some of these can be shared across different classes, e.g., car tires. In addition, it learns a map dictionary that contains common spatial activation patterns of components. During meta-testing, the knowledge base facilitates the learning of novel classes by re-using the already learned components and spatial activation patterns.

on a large number of base classes to enhance the learning of new classes from just a few classes.

Various approaches to few-shot classification were proposed in the past that take different perspectives. While these methods try to share the common knowledge among base classes and novel classes, since few-shot datasets do not include the attribute or component annotations like in zero-shot datasets, they do not explicitly consider that objects can have similar components and shapes that can be reused.

In this paper, we introduce a novel approach to few-shot classification that explicitly exploits that object components and their spatial activation patterns can be shared among different object classes. For example, the spatial structure of the class “horse” can be used for learning efficiently about the class “donkey”. We implement such a compositional representation sharing by train a knowledge base during meta-learning that consists of a dictionary of component representations and a dictionary of component activation maps that encode common spatial activation patterns of components (Figure 1). We start by extracting the fea-

ture representations of an image up to the last convolution layer of a standard backbone architecture, such as ResNet [11]. Following recent work on unsupervised component detection [21, 44, 45], the component dictionary is learned by clustering the individual feature vectors from the feature encoding of the training images. Moreover, we extract *component activation maps* by computing the spatial activation pattern of components in the training images. The component activation maps are clustered to learn a dictionary of prototypical maps that encode the most common spatial activation patterns of components. In practice, the elements of the map dictionary are optimized to be distinct from each other to avoid redundancies. During meta-testing, our model learns representations of objects by composing them from the components and component activation maps of the knowledge base. We use an attention layer to increase the weight of the components that are most discriminative for an object class. Finally, the learned object representations are fed into a classifier to predict the class label. During meta-training, the full model pipeline is trained end-to-end. During meta-testing, we observed that it is sufficient to train the classification head only, while freezing the learned backbone and knowledge base. This is different from the majority of other meta-learning methods and highlights the strong generalization performance induced by integrating compositional representation sharing into neural networks.

We evaluate our model on four popular few-shot classification datasets and achieve comparable performance on all datasets. In summary, we make several important contributions in this work:

1. To the best of our knowledge, we are the first to study and demonstrate the effectiveness and interpretability of compositional representation learning on few-shot classification.
2. We introduce CORL, a novel neural architecture for few-shot classification that implements the inductive prior of compositional representation sharing. It learns a knowledge base with component representations and their common spatial activation patterns, and re-uses this knowledge to learn efficiently about novel classes.
3. We achieve comparable performance on several standard benchmarks, outperforming many recent complex optimization methods.

2. Related Work

In this section, we review existing work on few-shot classification and compositional models.

2.1. Few-shot learning

Few-shot learning has received a lot of attention over the last years. Related work can be roughly classified

into two branches. The first branch focuses on finding a quick adaptation for the classifier when meeting new tasks. MAML [9] proposed a general optimization algorithm that can get improvements on a new task with a few gradient steps. MetaOptNet [20] replaced the linear predictor with an SVM in a MAML framework and introduced a differentiable quadratic programming solver to allow end-to-end training. FEAT [41] proposed set-to-set functions for a quick adaptation between instance and target embeddings. MELR [5] exploited inter-episode relationships to improve the model robustness to poorly-sampled shots.

Another line of work focuses on how to learn more generalizable feature embeddings and design simple yet efficient metrics for classification. Previous methods mainly learned the representations in a meta-learning framework where the training data was organized in the form of many meta-tasks. Matching Networks [36] employed two networks for support and query samples, respectively, followed by an LSTM with read-attention to encode the full embedding. Recently, large-training-corpus methods have become the new trend which argue that training a base network on the whole training set directly is also feasible. For example, Dynamic Few-shot [10] extended object recognition systems with an attention weight generator and redesigned the classifier module as the cosine similarity function. RFS [35] simply trained the embedding function on the combined meta-training sets followed by exploiting knowledge distillation to further improve the performance. It proves that learning a good representation through a proxy task, such as image classification, can give state-of-the-art performances. BML [47] further improves the learned global feature through an episode-wise local feature learning.

Though all these methods improve few-shot learning in different ways, they do not explicitly take into account that objects can have similar parts and shapes which can be reused. Our method follows the large-training-corpus thought and shows that by explicitly taking into account that objects can have similar components and shapes, we can reuse the information among them and further improve the performance.

2.2. Compositional models

A rich literature on compositional models for image classification exists. However, with the exception of very few works [8, 48], most methods use part annotations for training and do not share components among object classes. By contrast, our work shows that sharing parts and their spatial distributions among classes without labor-intensive part annotations enables efficient representation learning. Moreover, many traditional works [8, 48, 3, 7, 40] learn the model parameters directly from image pixels. The major challenge for these approaches is that they need to explicitly account for nuisances such as illumination and deformation. Several

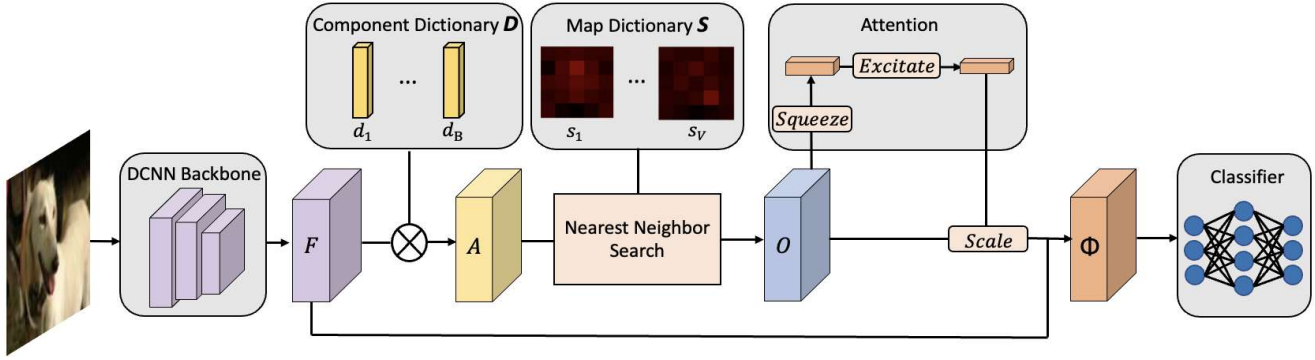


Figure 2: Feed-forward inference with CORL. A DCNN backbone is used to extract the feature map F . The items of the component dictionary D are used as kernels to compute a component activation map A . We then compare each channel in the component activation map A_b to the spatial patterns in the map dictionary S and multiply it element-wisely with the most similar one to compute the output O . An attention mechanism is used to further strengthen components that are most discriminative for an object class. The attention-weighted output is denoted as Φ . We concatenate Φ with average-pooled F and forward it to the classifier module to compute the final classification result.

recent works proposed to learn compositional models from the features of higher layers of deep convolutional neural networks, since these features have shown to be robust to nuisances and have some semantic meanings:

Liao et al. [21] proposed to integrate compositionality into DCNNs by regularizing the feature representations of DCNNs to cluster during learning. Their qualitative results show that the resulting feature clusters resemble detectors of different parts. Zhang et al. [44] demonstrated that component detectors emerge in DCNNs by restricting the activations in feature maps to have a localized spatial distribution. Kortylewski et al. [17] proposed to learn generative dictionary-based compositional models from the features of a DCNN. They use their compositional model as “backup” to an independently trained DCNN if the DCNNs classification score falls below a certain threshold. In follow-up work, Kortylewski et al. [15, 16] further proposed a fully differentiable compositional model for image classification that shows strong robustness to occlusion scenes. Sun et al. [33] demonstrated that these methods could be extended to combine image classification and amodal segmentation by leveraging compositional shape priors.

These recent advances inspire our work in integrating compositional models and deep neural networks. In this work, we propose to generate part information with compositional model and share it among different classes. In particular, our model for few-shot classification learns component representations and how to compose them together spatially into a whole object representation. We exploit that components and their spatial activation patterns can be shared among different classes, which enables our model to learn efficiently from very few examples.

3. Method

We first briefly review the framework of few-shot classification. Then we present how we learn the component dictionary module followed by a discussion on how to learn the map dictionary module and how to integrate these modules into a pipeline for few-shot classification. Lastly, we discuss how to train our model in an end-to-end manner.

3.1. Few-Shot Classification

Few-shot image classification consists of two consecutive learning processes: 1) In the meta-learning stage, the model acquires a knowledge base from a set of training classes. 2) During meta-testing, the acquired knowledge is used to recognize unseen classes from very few examples. The meta-training set T and meta-testing set S can be both organized as a collection of meta tasks where each meta-task is a N-way-K-shot classification problem. In this paper, we train our model on the combined meta-training set T followed by directly testing on meta-testing set S without fine-tuning the model parameters.

3.2. Learning a component dictionary via clustering

Formulation. We denote a feature map $F^l \in \mathbb{R}^{H \times W \times C}$ as the output of a layer l in a deep convolutional neural network, with C being the number of channels. A feature vector $f_p^l \in \mathbb{R}^C$ is the vector of features in F^l at position p on the 2D lattice \mathcal{P} of the feature map. In the remainder of this section, we omit the superscript l for notational clarity because this is fixed a-priori.

Learning component representations. A number of prior works [21, 44, 17, 15] on learning compositional representations showed that when clustering feature vectors f_p , the

cluster centers resemble image patterns that frequently occur in the training images. These patterns often share semantic meanings and therefore resemble part-like detectors. Motivated by these results, we aim at constructing a *component dictionary* $D = \{d_1, \dots, d_B\}$, in which the items $d_b \in \mathbb{R}^C$ are cluster centers of the feature vectors f_p from the training images. To achieve this, we integrate an additional clustering loss which will be introduced later into the overall loss function when training the network. Intuitively, this will encourage the dictionary items d_b to learn component representations from the intermediate layer l of a DCNN, and hence to capture the mid-level semantics of objects. Figure 3 illustrates examples of the dictionary items d_b after the meta-learning stage by showing image patches that activate each item the most. Note how the component representations indeed respond to semantically meaningful image patterns, such as the head of a dog.

3.3. Compositional Representation Sharing for Few-Shot Classification

Computing the spatial activation maps of components.

Given the component dictionary D , we compute the activation of a component representation d_b at a position p in the feature map F by computing the cosine similarity between d_b and the feature vector f_p . We implement this module as a convolution layer, which we call *component detection layer*. The convolutional kernels of the component detection layer are the items of the component dictionary D , and their kernel size is 1×1 . At every forward time, the kernels and input feature maps are $L2$ -normalized before computing the cosine similarity. The output of the detection layer is a component activation tensor $A \in \mathbb{R}^{H \times W \times B}$, where B is the number of items in dictionary D . Each channel in this tensor $A_b \in \mathbb{R}^{H \times W}$ is referred to as *component activation map*.

Learning dictionaries of spatial activation patterns. Our goal is to enable the model to share component activation patterns among different classes. This is inspired by the idea that components of different objects can have similar spatial activation patterns and that this natural redundancy should be exploited (e.g. the spatial structure of the class "dog" can be used for learning efficiently about the class "wolf"). We achieve this by learning a *map dictionary* $\mathbf{S} = \{S_1, \dots, S_V\}$, which contains the most common component activation patterns in the training data. We integrate the dictionary items $S_v \in \mathbb{R}^{H \times W}$ into the feed-forward stage by comparing them to the individual component activation maps A_b using the cosine similarity. We then select the closest item $\bar{v} = \arg \max_v \cos(S_v, A_b)$ and compute the output channel as point-wise multiplication between A_b and $S_{\bar{v}}$. After repeating this operation for all spatial distribution maps, we get the activated spatial distribution output denoted as $O \in \mathbb{R}^{H \times W \times B}$. In this way, each component activation

map A_b is encouraged to learn information from the most similar stored spatial activation pattern S_v .

Re-weighting important components with attention. To further augment components that are most important for representing a particular object, we adopt an attention mechanism to calculate different weights for the spatial distributions. We follow the design of SENet [12] with small changes. In particular, we first squeeze the global spatial information of O into a channel descriptor by using a learned filter $R \in \mathbb{R}^{H \times W \times B}$. Formally, a summary vector $z \in \mathbb{R}^B$ is generated by shrinking O through its spatial dimensions $H \times W$, such that the b -th entry of the vector z is calculated by:

$$z_b = \sum_{h=1}^H \sum_{w=1}^W R_b(h, w) O_b(h, w) \in \mathbb{R}. \quad (1)$$

To fully exploit the squeezed information, we then use the same gating mechanism as SENet which contains a bottleneck with two fully-connected layers and non-linearity activation. It can be represented as

$$l = \sigma(W_2 \delta(W_1 z)) \in \mathbb{R}^B \quad (2)$$

where σ refers to the Sigmoid activation and W_1, W_2 are the weights of the fully-connected layers. With the computed activation l , the final output is obtained by re-weighting the input O with l :

$$\Phi_b = l_b \cdot O_b \in \mathbb{R}^{H \times W} \quad (3)$$

where \cdot refers to channel-wise multiplication between the scalar l_b and the channel output O_b . Finally, we normalize feature vectors along the channel dimension in Φ to have unit norm and concatenate it with average-pooled F then forward it into the classifier to obtain a final prediction.

3.4. End-to-end Training of the model

During training, we use a two-layer fully-connected structure as a classifier to predict the classification results. Our model is fully differentiable and can be trained end-to-end using back-propagation. The trainable parameters of our model are $\Theta = \{\Omega, D, S\}$, where Ω are the parameters of the backbone used for feature extraction, e.g., ResNet-12. D is the component dictionary, and S is the dictionary of component activation maps. We optimize these parameters jointly using stochastic gradient descent. Our loss function contains three terms:

$$\mathcal{L}(y, y') = \mathcal{L}_{class}(y, y') + \gamma_1 \mathcal{L}_{cluster}(D) + \gamma_2 \mathcal{L}_{sparse}(S) \quad (4)$$

$\mathcal{L}_{class}(y, y')$ is the cross-entropy loss between the predicted label y' and the ground-truth label y . The second

Table 1: **Comparison to prior work on miniImageNet and tieredImageNet.** Average few-shot classification accuracies(%) with 95% confidence intervals on the meta-testing sets of miniImageNet and tieredImageNet. a-b-c-d denotes a 4-layer convolutional network with a, b, c, d filters in each layer.

model	backbone	miniImageNet 5-way		tieredImageNet 5-way	
		1-shot	5-shot	1-shot	5-shot
MAML [9]	32-32-32-32	48.70 ± 1.84	63.11 ± 0.92	51.67 ± 1.81	70.30 ± 1.75
Matching Networks [36]	64-64-64-64	43.56 ± 0.84	55.31 ± 0.73	-	-
Prototypical Networks [32]	64-64-64-64	49.42 ± 0.78	68.20 ± 0.66	53.31 ± 0.89	72.69 ± 0.74
Dynamic Few-shot [10]	64-64-128-128	56.20 ± 0.86	73.00 ± 0.64	-	-
Relation Networks [34]	64-96-128-256	50.44 ± 0.82	65.32 ± 0.70	54.48 ± 0.93	71.32 ± 0.78
TDADM [25]	ResNet-12	58.50 ± 0.30	76.70 ± 0.30	-	-
MetaOptNet [20]	ResNet-12	62.64 ± 0.61	78.63 ± 0.46	65.99 ± 0.72	81.56 ± 0.53
FEAT [41]	WRN-28-10	65.10 ± 0.20	81.11 ± 0.14	70.41 ± 0.23	84.38 ± 0.16
RFS [35]	ResNet-12	64.82 ± 0.60	82.14 ± 0.43	71.52 ± 0.69	86.03 ± 0.49
Neg-Cosine [22]	ResNet-12	63.85 ± 0.81	81.57 ± 0.56	-	-
MABAS [14]	ResNet-12	65.08 ± 0.86	82.70 ± 0.54	74.40 ± 0.68	86.61 ± 0.59
MELR [5]	ResNet-12	67.40 ± 0.43	83.40 ± 0.28	72.14 ± 0.51	87.01 ± 0.35
BML [47]	ResNet-12	67.04 ± 0.63	83.63 ± 0.29	68.99 ± 0.50	85.49 ± 0.34
DMF [38]	ResNet-12	67.76 ± 0.46	82.71 ± 0.31	71.89 ± 0.52	85.96 ± 0.35
infoPatch [23]	ResNet-12	67.67 ± 0.45	82.44 ± 0.31	71.51 ± 0.52	85.44 ± 0.35
IEPT [43]	ResNet-12	67.05 ± 0.44	82.90 ± 0.30	72.24 ± 0.50	86.73 ± 0.34
Meta DeepBDC [37]	ResNet-12	67.34 ± 0.43	84.46 ± 0.28	72.34 ± 0.49	87.31 ± 0.32
Ours	ResNet-12	65.74 ± 0.53	83.03 ± 0.33	73.82 ± 0.58	86.76 ± 0.52




Figure 3: Illustration of the interpretability of our elements in the component dictionary. Each row visualizes image patterns from the miniImageNet dataset that activate a dictionary item the most.

term $\mathcal{L}_{cluster}(D)$ is used to add additional regularization for the dictionary of components:

$$\mathcal{L}_{cluster}(D) = \sum_p \min_b (1 - \cos(D_b | f_p)) \quad (5)$$

where f_p refers to the feature vector at position p in the feature map F and $\cos(\cdot, \cdot)$ refers to the cosine similarity. Intuitively, this loss encourages the dictionary’s items to become similar to the feature vectors f_p . Thus the dictionary is forced to learn component representations that frequently

occur in the training data.

To regularize the map dictionary, we add a sparse loss on the dictionary S :

$$\mathcal{L}_{sparse} = \sum_{v=1}^V \arg \max_{v'} \cos(S_v, S_{v'})^2 \quad (6)$$

where $\cos(S_v, S_{v'})$ is the cosine similarity between two dictionary elements of S . This regularizer encourages the map dictionary elements to be sparse, thus avoiding that the elements become too similar to each other. We find that exploiting the second-order information of the cosine similarity avoids that the sparse loss will dominate the direction of the gradient at later stages of the training and thus helps the model to converge.

3.5. Replacing the classifier during meta-testing

At meta-testing time, unlike many other methods, we do not further fine-tune our model based on the support sets D_j^{train} in the meta-testing stage. Instead, we replace the fully-connected classification head with a simpler classifier to avoid overfitting. We tested different classifier, such as nearest neighbor based on different distance metrics, logistic regression classifier, linear support vector machine. We found that the logistic regression gives the best results. In summary, for a task $(D_j^{train}, D_j^{test})$ sampled from meta-testing set S , we forward D_j^{train} through the whole embedding function to get the attentioned component activation

map Φ contacted with the average-pooled F , and train the logistic regression classifier on this representation.

4. Experiment

In this section, we conduct extensive experiments that prove the effectiveness of our model. We first describe our detailed setup, which includes datasets, model structure, and hyper-parameters. Then we evaluate our model and make comparisons to related work on four few-shot classification benchmark datasets: miniImageNet [36], tieredImageNet [29], CIFAR-FS [1], Fewshot-CIFAR100 (FC100) [25]. The concrete performance on ImageNet and CIFAR derivatives are discussed respectively. We further conduct ablation studies to study the effects of the individual modules in our CORL pipeline. In the end, we show the receptive field of items in our component dictionary.

4.1. Experimental Setups

Architecture. Following previous work [24, 25, 28, 4], we use a ResNet12 as our feature extraction network which contains 4 residual blocks, where each of them contains 3 convolution layers. We drop the last average-pooling layer and use feature maps before the pooling for later computation. Dropblock is used in our model as a regularizer. The number of items in the component dictionary D is 512 and the number of items in the map dictionary S is 2048.

Implementation details. The loss coefficients in Eq. 4 are set to $\gamma_1 = 1$ and $\gamma_2 = 0.5$ respectively. We use the SGD optimizer with a momentum of 0.9 and a weight decay of $5e^{-4}$. Our batch size is set to 64, and the base learning rate is 0.05. We initialize the component dictionary D via K-means clustering on the feature vectors f_p and fine-tune it at the meta-training stage. We found that a random initialization of the component dictionary would not reduce the final performance, but the K-means initialization helps our model to converge faster as the cluster loss is lower at the start of training. On miniImageNet and tieredImageNet, we train our model 100 epochs and for CIFAR derivatives, the total epochs for training are 90. We adopt cosine annealing as the learning rate scheduler. During training, we adopt regular data augmentation schemes such as random flipping. When handling CIFAR derivatives datasets, we resize the input image to 84×84 pixels in order to have enough spatial resolution. Following common experimental setups, we report our performance based on an average of 600 meta-tasks, where each of them contains 15 test instances per class. For fair comparison, we only train our model on the training set of each dataset and do not perform any test-time training.

4.2. Experiments on ImageNet derivatives

The miniImageNet dataset is the most classic few-shot classification benchmark proposed by Matching Networks

[36]. It consists of 100 randomly sampled different classes, and each class contains 600 images of size 84×84 pixels. We follow the widely-used splitting protocol proposed by Ravi et al. [27], which uses 64 classes for meta-training, 16 classes for meta-validation, and 20 classes for meta-testing.

The tieredImageNet dataset is a larger subset of ImageNet, composed of 608 classes grouped into 34 high-level categories. They are further divided into 20 categories for training, 6 categories for validation, and 8 categories for testing, which corresponds to 351, 97, and 160 classes for meta-training, meta-validation, and meta-testing, respectively. This splitting method, which considers high-level categories, is applied to minimize the semantic overlap between the splits. Images are of size 84×84 .

Results. Table 1 summarizes the results on the 5-way miniImageNet and tieredImageNet. Our method achieves comparable performance on the miniImageNet benchmark for both 5-way-1-shot and 5-way-5-shot tasks. On tieredImageNet, we also achieve the best performance on the 5-way-5-shot task and comparable performance on the 5-way-1-shot task. Note that related works use very complex training schemes to improve their performance. For example, LEO [30] used an encoder and relation network in addition to the WRN-28-10 backbone network to produce sample-dependent initialization of the gradient descent. FEAT [41] and LEO [30] pre-train the WRN-28-10 backbone to classify 64 meta-training set of miniImageNet and then continue meta-training. FEAT [41] and MABAS [14] require additional fine-tuning on meta-testing sets. In contrast to all those approaches, our model just needs to train the embedding function through standard classification without further fine-tuning. This strategy allows us to clearly demonstrate the effect of a good embedding function by achieving stronger performance with an arguably simpler training.

4.3. Experiments on CIFAR derivatives

The CIFAR-FS dataset is a recently proposed few-shot image classification benchmark derived from CIFAR. It consists of all 100 classes and is further randomly split into 64 training classes, 16 validation classes, and 20 testing classes. Each class contains 600 images of size 32×32 .

The FC100 dataset is another few-shot classification dataset based on CIFAR. Its main idea is very similar to tieredImageNet, where the whole 100 classes are grouped into 20 superclasses. Each superclass is composed of standard 5 classes. These superclasses are divided into 12, 4, 4 for training, validation, testing correspondingly.

Results. Table 2 summarizes the performance on the 5-way CIFAR-FS and FC100. Our model achieves comparable performance on all tasks in both CIFAR-FS and FC100 benchmark. We observe that the relative improvement rate on the CIFAR-FS dataset is larger compared to the FC100 dataset which is similar to generalization pattern on the Im-

Table 2: **Comparison to prior work on CIFAR-FS and FC100.** Average few-shot classification accuracies(%) with 95% confidence intervals on the meta-testing sets of CIFAR-FS and FC100. a-b-c-d denotes a 4-layer convolutional network with a, b, c, d filters in each layer.

model	backbone	CIFAR-FS 5-way		FC100 5-way	
		1-shot	5-shot	1-shot	5-shot
MAML [9]	32-32-32-32	58.90 ± 1.90	71.50 ± 1.00	-	-
Prototypical Networks [32]	64-64-64-64	55.50 ± 0.70	72.00 ± 0.60	35.30 ± 0.60	48.60 ± 0.60
Relation Networks [34]	64-96-128-256	55.00 ± 1.00	69.30 ± 0.80	-	-
TADAM [25]	ResNet-12	-	-	40.10 ± 0.40	56.10 ± 0.40
Shot-Free [28]	ResNet-12	69.20 ± n/a	84.70 ± n/a	-	-
TEWAM [26]	ResNet-12	70.40 ± n/a	81.30 ± n/a	-	-
Prototypical Networks [32]	ResNet-12	72.20 ± 0.70	83.50 ± 0.50	37.50 ± 0.60	52.50 ± 0.60
MetaOptNet [20]	ResNet-12	72.60 ± 0.70	84.30 ± 0.50	41.10 ± 0.60	55.50 ± 0.60
DeepEMD [42]	ResNet-12	-	-	46.47 ± 0.78	63.22 ± 0.71
RFS [35]	ResNet-12	73.90 ± 0.80	86.90 ± 0.50	44.60 ± 0.70	60.90 ± 0.60
MABAS [14]	ResNet-12	73.51 ± 0.96	85.49 ± 0.68	42.31 ± 0.75	57.56 ± 0.78
ConstellationNet [39]	ResNet-12	75.40 ± 0.20	86.80 ± 0.20	43.80 ± 0.20	59.70 ± 0.20
BML [47]	ResNet-12	73.45 ± 0.47	88.04 ± 0.33	-	-
Ours	ResNet-12	74.13 ± 0.71	87.54 ± 0.51	44.82 ± 0.73	61.31 ± 0.54

Table 3: **Ablation study.** Performance of our ablated models on four few-shot classification benchmarks. The model is fixed when we conduct experiments about ablation for loss terms. The metric is average few-shot classification accuracies(%).

Map Dictionary	Attention	Cluster Loss	Sparse Loss	miniImageNet		tieredImageNet		CIFAR-FS		FC100	
				1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
				61.42	77.43	69.67	82.33	70.23	83.21	40.28	57.02
✓				61.71	78.65	70.53	82.67	70.89	83.72	40.62	57.72
✓	✓			62.03	79.44	71.10	83.42	71.67	84.59	41.13	58.25
✓	✓	✓		64.24	81.74	72.32	85.93	73.07	85.54	43.24	60.35
✓	✓	✓	✓	65.74	83.03	73.82	86.76	74.13	87.54	44.82	61.31

ageNet derivatives. Namely, the performance on the benchmark with semantic gaps between the meta-training set and meta-testing set benefits less from our method. We expect to alleviate this problem by finding a good way to fine-tune our model at meta-testing stage in future work.

4.4. Ablation Experiments

In this section, we conduct ablation studies on our CORL pipeline to analyze how its variants affect the few-shot classification result. We study the following three parts of our method: (a) The map dictionary; (b) The attention module on activated spatial distribution maps; (c) The cluster loss of the component dictionary; (d) The sparse loss of the map dictionary. In addition, we also analyze the result of the number of items in the component dictionary D , map dictionary S .

Table 3 shows the result of our ablation studies on mini-ImageNet, tieredImageNet, CIFAR-FS and FC100. We can see that when introducing the map dictionary, the model goes beyond the pure bag-of-words model and achieves on average 0.7% performance gain even without further using

loss to restrain it. This clearly shows that considering the spatial relationship between components helps the model. Besides, the attention mechanism for augmenting important components and their relationship makes the average performance improve on average around 0.6% on all datasets. With our cluster loss that regularizes the items in the component dictionary D , we gain on average about 2.3%. In addition, this loss increases the interpretability of our model as it makes the image patches detected by these component detectors more semantically meaningful. Our sparse loss regularizer improves the performance by another 1.5%, which demonstrates the benefit of making the items in the map dictionary distinct from each other.

Table 4 shows the influence of the size B of the component dictionary D on the performance of our model on miniImageNet. With too less items in the dictionary, our model do not contain enough information for modeling the component-whole based relationships of the objects. However, if the size B becomes too large, it harms each component representation to accurately capture the corresponding features and many items might focus on meaningless back-




Figure 4: **Visualization of elements in the component dictionary D .** Each row visualizes the activation of two dictionary components. Note that the learned components activate to semantically meaningful image patterns such as tires, animal heads, or parts of a bottle, even though no part annotations are used during training process.

Table 4: **Test accuracies(%) on meta-testing set of mini-ImageNet with a varying number of items in the component dictionary.** Either too much or too less items harm the performance of the model.

Component Dictionary Size B	miniImageNet	
	1-shot	5-shot
256	63.82	81.13
512	65.74	83.03
1024	65.12	82.45

ground thus enlarge the learning difficulty.

Figure 5 illustrates the influence of the number of items in the map dictionary S on the performance of our model on four benchmarks. The performance improves at first when the number of items increases but saturates as the dictionaries become larger. The performance keeps at the same level and even shows a tendency to drop. These results suggest that when the capacity of the dictionary is small, our model cannot store all necessary information. However, if the capacity becomes too large, the model starts overfitting.

4.5. Visualization of the components

In Figure 4 we visualize the activation of the elements in the component dictionary D following the method proposed by [46]. Specifically, we threshold the response map of the dictionary elements on images from the miniImageNet dataset and scale them up to the original image size. The visualization results show that the dictionary components respond to semantically meaningful image patterns such as tires, animal heads, or parts of a bottle. In summary, the strong discriminative abilities of our model (Tables 1 & 2) and the qualitative visualization in Figure 4 suggests that our compositional representation learning approach enables the model to learn semantically meaningful local image pat-




Figure 5: **Test accuracies(%) on meta-testing sets with a varying number of items in the map dictionary.** The performance of our model increases at first and saturates at some point with a slight tendency to drop further.

terns with no part-level supervisions during training.

5. Conclusion

In this work, we study the problem of few-shot image classification. Inspired by the compositional representation of objects in humans, we introduce CORL, a novel neural architecture for few-shot classification that learns through compositional representation sharing. In particular, CORL learns a knowledge base that contains a dictionary of component representations and a dictionary of component activation maps that encode frequent spatial activation patterns of components. During meta-testing, this knowledge is reused to learn about unseen classes from very few samples. Our extensive experiments demonstrate the effectiveness of our method, which achieves comparable performance on four popular few-shot classification benchmarks.

Acknowledgements. The authors gratefully acknowledge supports from ONR N00014-21-1-2690. AK acknowledges support via his Emmy Noether Research Group funded by the German Science Foundation (DFG) under Grant No. 468670075.

References

- [1] Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018.
- [2] Irving Biederman. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2):115, 1987.
- [3] Jifeng Dai, Yi Hong, Wenze Hu, Song-Chun Zhu, and Ying Nian Wu. Unsupervised learning of dictionaries of hierarchical compositional models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [4] Guneet S Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. *arXiv preprint arXiv:1909.02729*, 2019.
- [5] Nanyi Fei, Zhiwu Lu, Tao Xiang, and Songfang Huang. Melr: Meta-learning via modeling episode-level relationships for few-shot learning. In *International Conference on Learning Representations*, 2020.
- [6] Jacob Feldman. The structure of perceptual categories. *Journal of mathematical psychology*, 41(2):145–170, 1997.
- [7] Sanja Fidler, Marko Boben, and Ales Leonardis. Learning a hierarchical compositional shape vocabulary for multi-class object representation. *arXiv preprint arXiv:1408.5516*, 2014.
- [8] Sanja Fidler and Ales Leonardis. Towards scalable representations of object categories: Learning a hierarchy of parts. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007.
- [9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.
- [10] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [12] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [13] Alan Jern and Charles Kemp. A probabilistic account of exemplar and category generation. *Cognitive psychology*, 66(1):85–125, 2013.
- [14] Jaekyeom Kim, Hyoungseok Kim, and Gunhee Kim. Model-agnostic boundary-adversarial sampling for test-time generalization in few-shot learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 599–617. Springer, 2020.
- [15] Adam Kortylewski, Ju He, Qing Liu, and Alan L. Yuille. Compositional convolutional neural networks: A deep architecture with innate robustness to partial occlusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [16] Adam Kortylewski, Qing Liu, Angtian Wang, Yihong Sun, and Alan Yuille. Compositional convolutional neural networks: A robust and interpretable model for object recognition under occlusion. *International Journal of Computer Vision*, pages 1–25, 2020.
- [17] Adam Kortylewski, Qing Liu, Huiyu Wang, Zhishuai Zhang, and Alan Yuille. Combining compositional models and deep networks for robust object classification under occlusion. *arXiv preprint arXiv:1905.11826*, 2019.
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [19] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [20] Kwonjoon Lee, Subhansu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [21] Renjie Liao, Alex Schwing, Richard Zemel, and Raquel Urtasun. Learning deep parsimonious representations. In *Advances in Neural Information Processing Systems*, pages 5076–5084, 2016.
- [22] Bin Liu, Yue Cao, Yutong Lin, Qi Li, Zheng Zhang, Mingsheng Long, and Han Hu. Negative margin matters: Understanding margin in few-shot classification. In *European Conference on Computer Vision*, pages 438–455. Springer, 2020.
- [23] Chen Liu, Yanwei Fu, Chengming Xu, Siqian Yang, Jilin Li, Chengjie Wang, and Li Zhang. Learning a few-shot embedding model with contrastive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 8635–8643, 2021.
- [24] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141*, 2017.
- [25] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 721–731. Curran Associates, Inc., 2018.
- [26] Limeng Qiao, Yemin Shi, Jia Li, Yaowei Wang, Tiejun Huang, and Yonghong Tian. Transductive episodic-wise adaptive metric for few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [27] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. 2016.
- [28] Avinash Ravichandran, Rahul Bhotika, and Stefano Soatto. Few-shot learning with embedded class models and shot-free meta training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.

- [29] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676*, 2018.
- [30] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*, 2018.
- [31] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [32] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 4077–4087. Curran Associates, Inc., 2017.
- [33] Yihong Sun, Adam Kortylewski, and Alan Yuille. Weakly-supervised amodal instance segmentation with compositional priors, 2020.
- [34] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H.S. Torr, and Timothy M. Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [35] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola. Rethinking few-shot image classification: a good embedding is all you need? In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 266–282. Springer, 2020.
- [36] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29:3630–3638, 2016.
- [37] Jiangtao Xie, Fei Long, Jiaming Lv, Qilong Wang, and Peihua Li. Joint distribution matters: Deep brownian distance covariance for few-shot classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7972–7981, 2022.
- [38] Chengming Xu, Yanwei Fu, Chen Liu, Chengjie Wang, Jilin Li, Feiyue Huang, Li Zhang, and Xiangyang Xue. Learning dynamic alignment via meta-filter for few-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5182–5191, 2021.
- [39] Weijian Xu, Huaijin Wang, Zhuowen Tu, et al. Attentional constellation nets for few-shot learning. In *International Conference on Learning Representations*, 2020.
- [40] Ya Jin and S. Geman. Context and hierarchy in a probabilistic image model. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2145–2152, 2006.
- [41] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [42] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [43] Manli Zhang, Jianhong Zhang, Zhiwu Lu, Tao Xiang, Mingyu Ding, and Songfang Huang. Tept: Instance-level and episode-level pretext tasks for few-shot learning. In *International Conference on Learning Representations*, 2020.
- [44] Quanshi Zhang, Ying Nian Wu, and Song-Chun Zhu. Interpretable convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [45] Zhishuai Zhang, Cihang Xie, Jianyu Wang, Lingxi Xie, and Alan L Yuille. Deepvoting: A robust and explainable deep network for semantic part detection under partial occlusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1372–1380, 2018.
- [46] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Object detectors emerge in deep scene cnns. *arXiv preprint arXiv:1412.6856*, 2014.
- [47] Ziqi Zhou, Xi Qiu, Jiangtao Xie, Jianan Wu, and Chi Zhang. Binocular mutual learning for improving few-shot classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8402–8411, 2021.
- [48] Long Zhu, Yuanhao Chen, Antonio Torralba, William Freeman, and Alan Yuille. Part and appearance sharing: Recursive compositional models for multi-view. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1919–1926. IEEE, 2010.