

Reference-Limited Compositional Zero-Shot Learning

Siteng Huang Zhejiang University Hangzhou, China huangsiteng@westlake.edu.cn Qiyao Wei University of Cambridge Cambridge, UK qw281@cam.ac.uk

ABSTRACT

Compositional zero-shot learning (CZSL) refers to recognizing unseen compositions of known visual primitives, which is an essential ability for artificial intelligence systems to learn and understand the world. While considerable progress has been made on existing benchmarks, we suspect whether popular CZSL methods can address the challenges of few-shot and few referential compositions, which is common when learning in real-world unseen environments. To this end, we study the challenging reference-limited compositional zero-shot learning (RL-CZSL) problem in this paper, i.e., given limited seen compositions that contain only a few samples as reference, unseen compositions of observed primitives should be identified. We propose a novel Meta Compositional Graph Learner (MetaCGL) that can efficiently learn the compositionality from insufficient referential information and generalize to unseen compositions. Besides, we build a benchmark with two new large-scale datasets that consist of natural images with diverse compositional labels, providing more realistic environments for RL-CZSL. Extensive experiments in the benchmarks show that our method achieves state-of-the-art performance in recognizing unseen compositions when reference is limited for compositional learning.

CCS CONCEPTS

• Computing methodologies → Learning settings; Supervised learning by classification; Neural networks.

KEYWORDS

Zero-shot Learning, Meta Learning, Compositional Zero-shot Learning, Few-shot Learning

ACM Reference Format:

Siteng Huang, Qiyao Wei, and Donglin Wang. 2023. Reference-Limited Compositional Zero-Shot Learning. In *International Conference on Multimedia Retrieval (ICMR '23), June 12–15, 2023, Thessaloniki, Greece.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3591106.3592225

1 INTRODUCTION

Different from standard systems that are limited to a fixed set of categories at a time, humans generalize to a large, essentially

*Corresponding author.

This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

ICMR '23, June 12–15, 2023, Thessaloniki, Greece © 2023 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0178-8/23/06. https://doi.org/10.1145/3591106.3592225



Figure 1: Comparison among compositional zero-shot learning (CZSL), few-shot learning (FSL), and our proposed reference-limited compositional zero-shot learning (RL-CZSL). Different colors indicate different categories of entities or primitives.

"unbounded" concept space by reasoning in a compositional manner [1, 32]. This method of identifying novel complex concepts by composing known components (which we call the "primitives" in this paper) is called compositional generalization, representing the essential ability of human intelligence to make "infinite use of finite means" [4, 10]. For example, based on familiarity with tomatoes and other red objects, people can recognize a red tomato when they first encounter it. Similarly, it is easy to understand the behavior of cutting a pizza after cutting a cake and knowing what a pizza is. In the widely studied compositional zero-shot learning (CZSL) problem, with the side information (e.g., attributes, textual descriptions, and label embeddings), visual perception models are also expected to identify unseen compositional concepts, that is, they learn the compositionality of primitives from sufficient training samples and are tasked with generalizing to unseen combinations of these primitives [21, 22, 25].

While these efforts have contributed to a more comprehensive perception of the world, we argue that the existing setup seems idealistic and inappropriate to simulate natural human learning, and two core challenges should be considered when evaluating compositional learners. (1) Few-shot: Humans have an inherent ability to learn the compositionality of complex concepts with only a few examples and transfer the learned knowledge to different situations. However, AI systems will suffer from severer generalization issues if training samples are insufficient, as the empirical risk is far from being a good approximation for expected risk [34]. Although an increasing number of models have tried to alleviate potential overfitting [3, 6, 28], they still treat every class as an independent entity and require referential data for any novel concept. Hence we would like to investigate whether compositional learning can be performed with restricted sample size, in other words, whether few-shot learners can generalize to unseen label compositions. (2) Few referential compositions: Unlike recent methods [17, 27, 39] that have to refer to multiple combinations with the same primitive

to extract semantic invariants from them, humans can discover potential primitives from a few combinations, or even only one, based on prior knowledge. This contributes to the adaptation of humans to the long-tailed distribution of various compositional concepts in the real world, *i.e.*, there exist a few common primitives and many more composition-scarce primitives, making collecting all possible scenarios for each primitive in advance expensive and time-consuming. Therefore, few referential compositions should also be a natural constraint for human-level compositional learning.

In this paper, we introduce an untouched problem, referencelimited compositional zero-shot learning (RL-CZSL), to approximate real-world situations that would be encountered when compositional learning is required. The term "reference-limited" is adopted to indicate that when the model performs compositional learning, the combinations that can be used as references are limited in terms of the number of both categories and labeled samples, as shown in Figure 1. Therefore, RL-CZSL requires the learner to incorporate appropriate priors into learning, so that it can quickly learn the compositionality clues without superfluous references. Furthermore, due to the lack of rich categories and types of primitives and compositions, the existing datasets can not afford to create a large amount of testing environments for a comprehensive measurement. To address the limitations of the datasets and provide suitable conditions for systematic comparisons on RL-CZSL task, we build two benchmark datasets that consist of over 99k and 30k natural images covering sufficient attribute-object and action-object compositional labels, supporting us to sample realistic episodes to simulate partially observable worlds.

To address the new challenges, we propose a novel method Meta Compositional Graph Learner (MetaCGL). MetaCGL constructs a compositional graph to learn the dependencies between primitive and composition representations, and learns better semantic embeddings by aggregating information of neighbor nodes. With the updated semantic embeddings, MetaCGL generates a prior correlation map to estimate which features are related to the prediction target. Moreover, MetaCGL applies an effective bi-level optimization strategy during training, which contains a simple data augmentation method named Compositional Mixup to enhance the generalization ability. We compare our MetaCGL with representative CZSL methods on the proposed datasets, and the results show that MetaCGL significantly outperforms prior methods in recognizing unseen compositions. By shedding light on the limitations of existing settings and approaches, we hope to spur future work to develop human-level compositional generalization ability for intelligent systems. In summary, our contributions are as follows:

1. We introduce a new problem named reference-limited compositional zero-shot learning (RL-CZSL), where given only a few samples of limited compositions, the model is required to generalize to recognize unseen compositions. This offers a more realistic and challenging environment for evaluating compositional learners.

2. We establish two benchmark datasets with diverse compositional labels and well-designed data splits, providing the required platform for systematically assessing progress on the task.

3. We propose a novel method, Meta Compositional Graph Learner (MetaCGL), for the challenging RL-CZSL problem. Experimental results show that MetaCGL consistently outperforms popular base-lines on recognizing unseen compositions.

2 RELATED WORK

Compositional Zero-shot Learning (CZSL) aims to recognize unseen attribute-object compositions at test time while each constituent exists in training samples, given side information that describes novel composition pairs, e.g., word embeddings, attribute annotations, or text descriptions. As the early Visual Product (Vis-Prod) [21] baseline just computes the outputs of individual primitive classifiers as the predicted probability of the corresponding composition, some notable works [18, 19, 21, 22] argue that compositionality requires learning a joint compatibility function between the image, the attribute, and the object. And recently, state-of-theart methods [17, 27, 39] rely on comparing different compositions of the same primitive to learn to disentangle visual features. However, all of these works assume that there exist sufficient referential compositions and samples for learning in a compositional manner. In this paper, we propose the reference-limited compositional zero-shot learning (RL-CZSL) problem that removes this assumption to be more close to the real-world unseen environments. And our experimental results show that the state-of-the-art CZSL methods also struggle with this new challenge. As our introduced RL-CZSL benchmark datasets also contain compositional labels of action-object pairs, we also compare our topic with human-object interaction (HOI) detection [8, 12], which aims at detecting all the human-interaction-object triplets in an image. This task can also be viewed as a compositional zero-shot learning problem, as test images may contain interaction-object pairs that do not appear in the training data. However, HOI detection methods rely on a pre-trained object detector to localize the human and objects for further processing, which is not available for novel primitives in unseen environments.

Few-shot Learning (FSL) requires learning new tasks with few labeled examples. Recent FSL advances can be roughly categorized into the following three groups: (1) metric-based methods [28, 29, 33] learn a generalizable embedding model to transform all samples into a common metric space, where simple classifiers can be executed directly. (2) initialization-based methods [6, 26] learn a good set of initial parameters for the whole model or part of it, so that the model can quickly adapt to novel classes in a small number of gradient update steps. (3) pretraining-based methods [3, 5, 30] train a feature extractor with all the training data, and fix it during the meta-test phase whilst learning new classifiers for novel classes. Recently, several FSL works [9, 31, 40] have aimed to improve the generalization performance with compositional representations. Limited by the traditional FSL setting on which they are based, these methods only consider feature compositionality and have not explored how to generalize to new label compositions.

3 REFERENCE-LIMITED COMPOSITIONAL ZERO-SHOT LEARNING

3.1 **Problem Formulation**

The ultimate goal of the RL-CZSL task is to recognize unseen visual pair compositions, whose primitives have only appeared in limited seen compositions containing only a few samples. In this paper, we follow the FSL setting to use the sampled episodes as a simulation of independent test environments, which refer to the data for learning as *support* and the data for inference as *query*. In addition, we

Table 1: Basic statistics of our proposed datasets. The symbol # is used as an abbreviation for "number of".

Dataset	RL-CZSL-ATTR	RL-CZSL-ACT
Composition type	attribute-object	action-object
Total #c	1,768	574
#c in train / val / test	1,076 / 136 / 556	214 / 22 / 338
Total #p ¹	190	185
p^1 in train / val / test	105 / 33 / 52	52 / 10 / 123
Total #p ²	488	154
# p^2 in train / val / test	281 / 12 / 195	59 / 11 / 84
Total #samples	99,771	30,420
#samples in train / val / test	51,928 / 29,922 / 17,921	20,604 / 1,207 / 8,609

apply an **open world** setting that while all compositions contained in the *support classes* are *seen* ones, the *query classes* include not only *unseen compositions*, but also *seen compositions*. At the same time, no constraint on the test time search space is imposed. Allowing predictions to come from all possible pairs in the current episode, the setting is more close to the unseen environments that are likely to arise in real-world deployments, and thus leads to a more comprehensive study on achieving a balanced and promising performance of both seen and unseen compositions.

More formally, we consider the visual recognition setting where each image x is associated with a complex concept c that is a pair composition of two primitives p^1 and p^2 , *i.e.*, $c = (p^1, p^2)$. For example, p^1 can represent a state like "cooked" or an action like "cut", while p^2 can refer to an object such as "chicken" or "pizza". When testing, the model are evaluated on episodes that are sampled from a set of *novel* data $\mathcal{D}_n = \{(x_n^{(i)}, c_n^{(i)})\}$ with label space $\mathcal{C}_n \subset \mathcal{P}_n^1 \times \mathcal{P}_n^2 = \{(p_n^1, p_n^2) | p_n^1 \in \mathcal{P}_n^1, p_n^2 \in \mathcal{P}_n^2\}$. C_n denotes the novel composition of tion set, $\mathcal{P}_n^1, \mathcal{P}_n^2$ are the two corresponding novel primitive sets with different primitive types (i.e., attributes, actions, or objects), and each primitive set contains N^p primitive categories. Each episode contains a *support* set $S = \{(x_s^{(i)}, c_s^{(i)}) | i = 1, 2, ..., N_s^c \times K_s^c\}$ that consists of N_s^c support classes with K_s^c labeled samples per class, and a query set $Q = \{(x_q^{(i)}, c_q^{(i)}) | i = 1, 2, ..., N_q^c \times K_q^c\}$ that consists of N_q^c query classes with K_q^c samples per class. The query classes not only contain N_s^c seen compositions that are all in the support classes, but also comprise $(N_a^c - N_s^c)$ unseen compositions that do not overlap with seen compositions. However, seen and unseen compositions in the same episode share the same two primitive sets sampled from \mathcal{P}_n^1 and \mathcal{P}_n^2 , providing the possibility for unseen compositions to be recognized. Following the open world setting, the prediction space of the model contains $N^p \times N^p$ compositions including seen, unseen and unfeasible ones. And the goal of the model is to correctly predict the compositional labels of samples in Q with the access to S.

To extract the prior knowledge for learning to rapidly separate primitive features from images, in the training phase, the model possesses the access to a set of *base* data $\mathcal{D}_b = \{(x_b^{(i)}, c_b^{(i)})\}$ with label space $C_b \subset \mathcal{P}_b^1 \times \mathcal{P}_b^2$. Note that the base and novel primitive sets do not overlap, *i.e.*, $\mathcal{P}_b^1 \cap \mathcal{P}_n^1 = \emptyset$ and $\mathcal{P}_b^2 \cap \mathcal{P}_n^2 = \emptyset$, and thus $C_b \cap C_n = \emptyset$ also holds. We would like to mention that RL-CZSL does not require a specific procedure for learning from the training data, that is, episodic and non-episodic methods are both allowed.



(d) RL-CZSL-ACT: train

(f) RL-CZSL-ACT: test

Figure 2: Examples taken from RL-CZSL-ATTR (the top row) and RL-CZSL-ACT (the bottom row) benchmark datasets.

(e) RL-CZSL-ACT: val

3.2 Proposed Benchmarks

Although several CZSL datasets like MIT-States [11], UT-Zap50K [36], and C-GQA [22], have been proposed, there exist limitations that prevent them from becoming appropriate benchmarks for measuring human-level compositional generalization ability: (1) They contain images with only state-object pair labels, failing to cover the most frequent types of compositions in the real world. We argue that another common composition type, action-object pairs, should also be considered to examine the compositional reasoning ability of models. (2) More importantly, we found that when trying to divide them into three splits, existing datasets did not have enough categories of compositions and primitives to ensure that no primitives in the splits of different phases overlapped. This made it impossible to create a sufficient number of various episodes to simulate diverse unseen environments during the test phase, ultimately reducing the validity of the evaluation.

Therefore, we refine public datasets to form two suitable benchmark datasets for RL-CZSL, named **RL-CZSL-ATTR** and **RL-CZSL-ACT**. RL-CZSL-ATTR contains over 99k images attached with 1,768 attribute-object pair labels from C-GQA [22], UT-Zap50K [36], and MIT-States [11]. And RL-CZSL-ACT consists of over 30k images with 574 action-object pair labels from HICO [2], Visual Genome [15], and imSitu [35]. The statistics are summarized in Table 1, and we show some sample images in Figure 2.

Organizing Data. To ensure that the two datasets we obtained meet the needs of the RL-CZSL task, we filtered the data in all splits from the perspective of labels. Specifically, (1) compositions with fewer than 10 samples were screened out to ensure that enough same-class support and query samples can be simultaneously sampled without duplicates, (2) for primitives that appeared in multiple splits, we kept them in at most one split, (3) size-related attribute primitives that cannot be accurately depicted in the images, such as "small", "large", "long" and "short", were also filtered out from our datasets. Specially, images from the Visual Genome dataset are densely annotated with numerous attributes and objects, lacking a description of the focus of the content. Therefore, we kept the attribute-object compositional label of the largest bounding box in each image, which most likely corresponds to the main content, and removed other annotations.

Episode Sampling Strategy. In this section, we outline the sampling strategy that creates more realistic episodes for the RL-CZSL task, and the corresponding pseudocode is illustrated in Algorithm 1. In each experiment, the value of N^p is fixed for all sampled episodes in the same phase. However, we allow episodes to have a different number of seen and unseen compositions, *i.e.*, the values of N_s^c and N_q^c may vary from episode to episode. As N_q^c actually corresponds to the number of all potential compositions that can be obtained by pairing all primitives in the episode and also exist in the dataset, N_s^c is randomly sampled within a certain range of $[N^p + 1, N_q^c - 1]$, making the episode closer to reality. The maximum value of this range guarantees that there exist unseen compositions in the episode, and the minimum value implies that each primitive has the opportunity to appear in more than one seen composition without being bound to another primitive all the time.

Concretely, for each episode, we first randomly sample N^p seen compositions to obtain two primitive sets $\mathcal{P}_n^1, \mathcal{P}_n^2$ without duplicate primitives. Then, we check if \mathcal{P}_n^1 and \mathcal{P}_n^2 can be paired to get enough existing compositions for being divided into seen and unseen ones. An episode that can achieve the required number of composition pairs will be regarded as a valid episode, and the remaining compositions will be randomly assigned to seen and unseen groups on the premise of satisfying the restriction of N_s^c . Therefore, we have N_s^c seen compositions for the support classes and N_q^c compositions including seen and unseen ones for the query classes. Next, we randomly sample K_s^c support samples for each seen composition and K_q^c samples for each composition in the query classes. Note that for seen compositions that exist in both support and query classes, the samples assigned to the two sets are not duplicated.

Evaluation Metrics. We use three metrics to evaluate how well the learned model recognizes both unseen and seen composition pairs, consistent with the adopted open world setting: (1) **Unseen accuracy (UA)**: The average of the accuracy computed on query samples from unseen compositions on all test episodes. (2) **Seen accuracy (SA)**: The average of the accuracy computed on query samples from seen compositions on all test episodes. (3) **Harmonic mean (HM)**: A metric that quantifies the overall performance of both seen and unseen accuracy based on the results of all test episodes, defined as: HM = 2 (SA * UA) / (SA + UA).

4 META COMPOSITIONAL GRAPH LEARNER

For RL-CZSL, we propose a novel Meta Compositional Graph Learner (MetaCGL), whose architecture is illustrated in Figure 3. Following the popular CZSL framework [19, 22], MetaCGL embeds both images and composition representations into a shared embedding space and computes a compatibility score to measure the similarity between visual and composition embeddings. Also, MetaCGL follows the episodic training paradigm to keep the spirit of matching training and test conditions, *i.e.*, MetaCGL is trained on episodes sampled from \mathcal{D}_b using the same algorithm as used for test episodes. In the following, we introduce how MetaCGL learns composition and visual embeddings, calculates compatibility scores, and is trained with a designed optimization process.

Algorithm 1: Episode Sampling in RL-CZSL **Require:** \mathcal{D} with label space *C* according to the requested class split of the given dataset, N^p , K^c_s , K^c_a **Output:** the sampled episode (S, Q)**Step 1.** Sample primitive sets \mathcal{P}^1 , \mathcal{P}^2 $\mathcal{P}^1 = \{\}, \mathcal{P}^2 = \{\}, \text{ seen compositions } C_{\text{seen}} = \{\}$ while $|C_{seen}| < N^p$ do Randomly sample a composition $c = (p^1, p^2)$ from *C* **if** $(p^1 \text{ not in } \mathcal{P}^1)$ and $(p^2 \text{ not in } \mathcal{P}^2)$ **then** | Add *c* into C_{seen} , p^1 into \mathcal{P}^1 , and p^2 into \mathcal{P}^2 **Step 2.** Sample the support set S $S = \{\}$, candidate compositions $C_{\text{candidate}} = \{\}$, unseen compositions $C_{unseen} = \{\}$ for $c \in \mathcal{P}^1 \times \mathcal{P}^2$ do if (c in C) and $(c \text{ not in } C_{seen})$ then Add c into $C_{\text{candidate}}$ if $|C_{candidate}| < 2$ then // Seen and unseen compositions are insufficient Jump back to Step 1 Randomly assign the first two compositions in $C_{\text{candidate}}$ to each of Cseen and Cunseen, and the remaining ones in $C_{\text{candidate}}$ are randomly assigned to either C_{seen} or C_{unseen} each time for $c \in C_{seen}$ do Randomly sample K_s^c samples of c into S **Step 3.** Sample the query set *Q* $Q = \{\}$ for $c \in C_{seen}$ do Randomly sample K_q^c samples of c from those do not overlap with S into Qfor $c \in C_{unseen}$ do Randomly sample K_a^c samples of c into Q

4.1 Learning Composition Embeddings

We choose to learn semantic embeddings for compositions and primitives with a pre-defined compositional graph, as the graph structure has been proved to be effective in modeling dependency relationships between them [22]. Specifically, we construct a graph that contains $H = |\mathcal{P}^1| + |\mathcal{P}^2| + |\mathcal{C}|$ nodes in each episode, where \mathcal{P}^1 and \mathcal{P}^2 correspond to the primitive sets in the episode, and Cis the set of all potential pairs composed of primitives in \mathcal{P}^1 and \mathcal{P}^2 , including existent and nonexistent compositions in the dataset. The features of primitive nodes are initialized with the pre-trained word embeddings to utilize the prior knowledge extracted from large corpora, and the features of composition nodes are initialized by averaging the word embeddings of associated primitives. We represent the initial features of all nodes as $V^{(0)} \in \mathbb{R}^{H \times d_w}$, where d_w denotes the dimension of the word embeddings. The nodes of p^1 , p^2 and *c* are connected one by one for each $c = (p^1, p^2) \in C$, and a self-loop is added to each node. Thus, a symmetric adjacency matrix $A \in \mathbb{R}^{H \times H}$ can be obtained by setting $A_{ij} = 1$ if nodes *i* and *j* are connected, otherwise $A_{ij} = 0$. By applying a multi-layer graph convolutional network [14] \mathcal{G} , the node features can be updated as

Reference-Limited Compositional Zero-Shot Learning



Figure 3: The architecture of our proposed Meta Compositional Graph Learner (MetaCGL). Best viewed in color.

$$V^{(l+1)} = \sigma(D^{-\frac{1}{2}}AD^{-\frac{1}{2}}V^{(l)}\theta_{\mathcal{G}}).$$
 (1)

Here the non-linear activation function σ is ReLU, $V^{(l+1)}$ is the output of the l^{th} layer, $\theta_{\mathcal{G}}$ is the trainable weight matrix. $D \in \mathbb{R}^{H \times H}$ is a diagonal node degree matrix with $D_{ii} = \sum_j A_{ij}$, which preserves the scale of feature vectors to avoid gradient vanishing or explosion.

4.2 Learning Visual Embeddings

To map the image features $F \in \mathbb{R}^{c \times h \times w}$ extracted by the backbone to the shared embedding space, an additional embedding function is required. While prior methods [19, 22] use simple multi-layer perceptrons (MLPs) to directly embed image features into a vector, we argue that features semantically unrelated to the target composition, such as the background, are also encoded into the embedding. To avoid such noise affecting the learning of the shared embedding space and thus damaging the accuracy of recognition, we propose to make a prior prediction about the features that might be relevant to the recognition target. Specifically, we transform the learned semantic embeddings into a prior correlation map with a correlation map generating network \mathcal{M} . Such a correlation map tells the probability of the features belonging to the target composition. For the *i*-th composition $c^{(i)} = (p^{1;(i)}, p^{2;(i)})$, we obtain the corresponding composition and primitive embeddings $v_{c^{(i)}}, v_{p^{1;(i)}}, v_{p^{2;(i)}} \in \mathbb{R}^d$ from the outputs of \mathcal{G} , where d is the dimension of the output vectors. In each episode, more images contain the same primitive than those that are associated with the same composition, i.e., more samples can be used to learn the visual invariants for the primitive embeddings. Therefore, we leverage primitive embeddings instead of composition embeddings to produce the prior correlation map. For each primitive embedding v_p , we firstly apply global average pooling (GAP) on the feature map F, and concatenate GAP(F) with v_p to obtain a new vector, which is then forwarded to \mathcal{M} . In practice, ${\cal M}$ consists of two convolutions with kernel size 1 and can also be viewed as two linear transformations with a ReLU activation in between. By exploiting the relationships between each dimension of the semantic vector and each channel of visual vector, \mathcal{M} generates a correlation map $m \in \mathbb{R}^{c}$. We mention that the parameters $\theta_{\mathcal{M}}$ are not shared between different primitive types, and we apply an element-wise summation to the outputs of different $\mathcal M$ to produce the final correlation map $m' \in \mathbb{R}^c$, which is used to highlight the relevant channel. In summary, we have

Algorithm 2: Training and Inference of MetaCGL
Training
Require: base data \mathcal{D}_b
Output: trained model with parameters θ
Randomly initialize θ
while not done do
Randomly sample a train episode $(\mathcal{S}, \mathcal{Q})$ from \mathcal{D}_b
Calculate $\mathcal{L}_{\mathcal{S}}(f_{\theta})$ using Eq. (4)
Calculate updated parameters $\theta^{'}$ using Eq. (7)
Generate $ ilde{m{Q}}$ using Eq. (5) and Eq. (6)
Calculate $\mathcal{L}_{\tilde{Q}}(f_{\theta'})$ using Eq. (4)
Update θ using Eq. (8)
Inference
Require: trained model with parameters θ
Calculate $\mathcal{L}_{\mathcal{S}}(f_{\theta})$ using Eq. (4) with \mathcal{S} in the test episode

Calculate $\mathcal{L}_{\mathcal{S}}(f_{\theta})$ using Eq. (4) with \mathcal{S} in the **test** episod $(\mathcal{S}, \mathcal{Q})$ Calculate updated parameters θ' using Eq. (7)

Use θ' for the inference of samples in Q

$$m^{\prime(i)} = \sigma(\mathcal{M}^{1}([\text{GAP}(F); v_{p^{1};(i)}])$$
⁽²⁾

$$+\mathcal{M}^{2}([GAP(F); v_{p^{2};(i)}])),$$
 (3)

where the activation function σ is sigmoid, and [;] denotes the operation of concatenating. Therefore, we have $F'^{(i)} = m'^{(i)} \otimes F$, where channel-wise correlation map values are broadcasted along the spatial dimension, and \otimes denotes element-wise multiplication. The reason we focus on the channel dimension of the feature map rather than the spatial dimension is that, the semantics of different primitives are inherently entangled in the spatial dimension, *i.e.*, it would be difficult to determine which primitive is represented by a single pixel. However, as each channel of a high-level feature map can be considered as a feature detector [37], the correlations between channels and primitive semantics can be more easily learned. Another embedding function \mathcal{E} , implemented by a MLP with parameters $\theta_{\mathcal{E}}$, is then used to encoder $F'^{(i)}$ into a vector $e^{(i)}$.

4.3 Calculating Compatibility Score

For each $c^{(i)} \in C$ in the current episode, we now can calculate a compatibility score $s^{(i)} = e^{(i)} \odot v_{c^{(i)}}$, where \odot represents element-wise dot product. Naturally, if $c^{(i)}$ is the ground-truth compositional label of the current image x, $s^{(i)}$ is expected to be larger and can be optimized with a cross-entropy loss function:

$$\mathcal{L}_{\boldsymbol{x}}(f_{\theta}) = -\log\left(\frac{\exp(s^{(i)})}{\sum_{i}^{|C|}\exp(s^{(j)})}\right),\tag{4}$$

where $\theta = \{\theta_{\mathcal{G}}, \theta_{\mathcal{M}}, \theta_{\mathcal{E}}\}$ represents all the trainable parameters in MetaCGL, and f_{θ} represents the whole model with parameters θ . And during inference, the prediction can be made by applying a argmax operation on all compatibility scores calculated for *x*.

4.4 Training Strategy

As RL-CZSL requires the model to generalize to unseen compositions with insufficient samples from only a few referential compositions, the training process also requires a special design to avoid overfitting. We apply a bi-level optimization for MetaCGL, as such an optimization strategy has been proven to be effective in helping to quickly adapt to new environments on a range of problems [6, 16]. Moreover, we propose a new data augmentation method Compositional Mixup to enhance the generalization and robustness. Based on the popular data augmentation method Mixup [38], Compositional Mixup improves the process of generating labels to satisfy the needs of compositional learning.

Specifically, in a train episode (S, Q) sampled from base data \mathcal{D}_b , we leverage Q to construct a set of augmented query samples $\tilde{Q} = \{(\tilde{x}_q^{(i)}, \tilde{c}_q^{(i)}) | i = 1, 2, ..., N_q^c \times K_q^c\}$. For each query sample $(x_q^{(i)}, c_q^{(i)})$ in Q, we randomly sample another query sample $(x_q^{(j)}, c_q^{(j)})$ from Q. A new example of image can be formed by a weighted linear interpolation of $x_q^{(i)}$ and $x_q^{(j)}$:

$$\tilde{x}_{q}^{(i)} = \lambda x_{q}^{(i)} + (1 - \lambda) x_{q}^{(j)}, \tag{5}$$

where $\lambda \in [0, 1]$ is a random value drawn from $\text{Beta}(\alpha, \alpha)$ distribution, and the hyper-parameter α is set to 1.0 in our experiments. Specially, considering that $c_q^{(i)} = (p^{1;(i)}, p^{2;(i)})$ and $c_q^{(j)} = (p^{1;(j)}, p^{2;(j)})$ are compositional labels, we form the new augmented label as

$$\begin{split} \tilde{c}_{q}^{(i)} &= \lambda^{2} c_{q}^{(i)} + \lambda (1 - \lambda) c_{q}^{(ij)} \\ &+ \lambda (1 - \lambda) c_{q}^{(ji)} + (1 - \lambda)^{2} c_{q}^{(j)}, \end{split} \tag{6}$$

where $c_q^{(ij)} = (p^{1;(i)}, p^{2;(j)}), c_q^{(ji)} = (p^{1;(j)}, p^{2;(i)})$. An intuitive explanation of our Compositional Mixup is that it implicitly introduces new compositions $c_q^{(ij)}$ and $c_q^{(ji)}$ that may not appear in the episode or even in the dataset. And the bi-level optimization process pushes the model to generalize to these augmented query samples well after even one gradient update on S, reducing the number of undesirable oscillations when predicting outside the samples from few referential compositions. Formally, θ are updated as

$$\theta' = \theta - \epsilon \nabla_{\theta} \mathcal{L}_{\mathcal{S}}(f_{\theta}), \tag{7}$$

$$\theta \leftarrow \theta - \gamma \nabla_{\theta} \mathcal{L}_{\tilde{Q}}(f_{\theta'}), \tag{8}$$

where hyper-parameters ϵ and γ are the step size and the meta step size. The experimental results in Table 3 show that the combination of our Compositional Mixup and bi-level optimization can effectively handle the challenges of few-shot and few referential compositions, and thus improves the accuracy of recognizing unseen compositions. We illustrate the pseudocode of the training and inference in Algorithm 2 for reference.

5 EXPERIMENTS

5.1 Experimental Setup

For a fair comparison, the basic experiments are conducted with a four-layer convolution backbone (Conv-4) as in [3] for all implemented methods, and the backbone is fixed after training on the base data. We also show the results on deeper backbones including ResNet-10 and ResNet-18 [7] in Section 5.5 with the reason why we do not take them as the default choice. If not specified, K_s^c , K_q^c , and N^p are all set by default to 5 while N_s^c and N_q^c are dynamic

and randomly sampled in each episode. For methods using side information, we initialize the word embeddings with pre-trained 300dimensional word2vec [20] vectors. And the best model is selected with the HM performance on the validation set. The reported results are obtained by averaging 3 random experiments. Code and datasets will be available at https://github.com/bighuang624/RL-CZSL.

Pretraining Settings. To pretrain the backbone, we use an Adam [13] optimizer for the Conv-4 network and the stochastic gradient descent (SGD) optimizer for other backbone networks. The backbone network, appended with a softmax layer, is trained with base data \mathcal{D}_b to classify all compositions in C_b using the cross-entropy loss. Standard data augmentation including random crop, left-right flip, and color jitter, is applied. The pretraining lasts for a maximum of 500 epochs with a batch size of 128. And the initial learning rate is set to 10^{-3} with a L2 penalty of 5×10^{-4} .

Training Details. Referring to [3], for methods that require training parameters in test episodes, we use the entire support set to train for 100 iterations with a batch size of 4. All CZSL methods are trained with a SGD optimizer with an initial learning rate 10^{-2} and a L2 penalty of 10^{-3} . And we keep the other settings of hyperparameters in the public source code. For our MetaCGL, we train at most 60,000 episodes with an initial learning rate 10^{-3} and a L2 penalty of 5×10^{-4} . And the step size ϵ is set to 0.4. Standard data augmentation is also applied when training all methods. We implement our experiments in PyTorch [24] and use a Nvidia V100 GPU to train all models.

Unchosen Baselines. We here discuss several existing CZSL methods that we did not include in our experiments. Attributes-asoperators (ATTOP) [23] views attributes as linear operators in the embedding space of object label embeddings. As ATTOP is a method that often participates in comparison, we found it impossible to adapt to RL-CZSL, and the main reason is the benchmark datasets, especially RL-CZSL-ACT, fail to provide antonyms for the method to calculate a loss term that operates over pairs of antonym attributes. Besides, the commutative loss and the inverse loss cannot be calculated under the constraint of few referential compositions. Some recent CZSL methods including OADis [27], SCEN [17] and IVR [39] rely on simultaneously sampling images with same object and different attribute (or same attribute and different object) when updating the model, which is also unavailable in RL-CZSL as one primitive may only appear in one composition in the episode. Another compositional learning method proposed by [12] was also considered at first, which is a rare HOI learning method that does not employ a pre-trained detector. However, we found the authors did not release the source code, making us unable to implement.

5.2 Comparing with Baselines

We compare MetaCGL and CZSL baseline methods on the two RL-CZSL benchmark datasets. As shown in Table 2, MetaCGL attains the best HA and UA. Although TMN and CompCos achieve better SA, the poor UA of all CZSL methods reveals that they overfit to seen compositions and fail to generalize to unseen ones when the number of referential compositions and samples is limited.

5.3 Error Analysis

To gain a further intuition of the performance of all methods, we analyze their errors when recognizing samples of unseen pairs.

Table 2: Comparison with the baselines on two proposed benchmarks. Detailed results (%) are reported with 95% confidence intervals. Best results are displayed in boldface.

		F	L-CZSL-AT	TR		RL-CZSL-ACT					
Method UA SA HM Prim 1 HM Prir		Prim 2 HM	UA	SA	HM	Prim 1 HM	Prim 2 HM				
		$K_s^c = 1$									
VisProd [21]	$1.24{\scriptstyle\pm0.43}$	$20.99{\scriptstyle\pm0.18}$	$2.34{\scriptstyle\pm0.76}$	$25.44{\scriptstyle\pm2.15}$	$27.41{\scriptstyle\pm1.79}$	0.88 ± 0.19	$21.97{\scriptstyle~\pm 0.56}$	$1.68{\scriptstyle\pm 0.36}$	$26.12{\scriptstyle\pm1.78}$	$25.43{\scriptstyle\pm1.15}$	
LE [21]	1.01 ± 0.96	$14.98{\scriptstyle\pm0.52}$	$1.89{\scriptstyle\pm1.67}$	$22.25{\scriptstyle\pm0.90}$	$22.45{\scriptstyle~\pm 0.53}$	$1.27{\scriptstyle\pm0.85}$	$14.02{\scriptstyle\pm0.18}$	$2.32{\scriptstyle\pm1.44}$	$17.90{\scriptstyle\pm3.28}$	$22.45{\scriptstyle~\pm4.56}$	
TMN [25]	0.52 ± 0.50	$28.31 {\scriptstyle \pm 0.39}$	$1.02{\scriptstyle\pm0.97}$	$24.86{\scriptstyle\pm1.59}$	$32.41{\scriptstyle\pm0.50}$	$0.62{\scriptstyle\pm0.34}$	28.85 ± 0.10	$1.21{\scriptstyle\pm0.65}$	$31.76 {\scriptstyle \pm 0.37}$	$26.03{\scriptstyle\pm1.55}$	
SymNet [18]	$1.94{\scriptstyle\pm0.08}$	$17.34{\scriptstyle\pm0.80}$	$3.48{\scriptstyle\pm0.12}$	$27.01{\scriptstyle\pm1.05}$	$23.95{\scriptstyle~\pm 2.87}$	$2.28{\scriptstyle\pm1.71}$	$17.90{\scriptstyle\pm0.56}$	$4.01{\scriptstyle\pm2.72}$	$27.35{\scriptstyle\pm1.59}$	$23.02{\scriptstyle\pm2.82}$	
CompCos [19]	2.57 ± 0.55	$25.14{\scriptstyle\pm0.70}$	$4.66{\scriptstyle\pm0.93}$	$26.84{\scriptstyle\pm0.72}$	$33.53{\scriptstyle\pm2.39}$	$3.02{\scriptstyle\pm0.34}$	$28.19{\scriptstyle\pm0.55}$	$5.45{\scriptstyle~\pm0.56}$	$\textbf{32.07} {\scriptstyle \pm 2.46}$	28.51±2.57	
CGE [22]	$4.65 {\scriptstyle \pm 1.12}$	$15.40{\scriptstyle\pm0.54}$	$7.13{\scriptstyle\pm1.29}$	$25.97{\scriptstyle\pm3.30}$	$31.56{\scriptstyle\pm1.26}$	$4.05{\scriptstyle\pm0.78}$	$15.51{\scriptstyle\pm0.91}$	$6.41{\scriptstyle\pm0.91}$	$28.56{\scriptstyle\pm2.09}$	$26.39{\scriptstyle\pm1.50}$	
MetaCGL (Ours)	10.44 ± 0.42	$19.01{\scriptstyle\pm1.78}$	$\boldsymbol{13.47} \pm 0.77$	$\textbf{30.22} {\pm} 1.58$	$\textbf{38.37} \pm 2.29$	7.76±0.31	$15.95{\scriptstyle\pm1.24}$	$\boldsymbol{10.44} {\scriptstyle\pm 0.41}$	$31.19{\scriptstyle\pm1.38}$	$26.68{\scriptstyle\pm2.76}$	
$K_s^c = 5$											
VisProd [21]	$0.55{\scriptstyle\pm0.45}$	$15.83{\scriptstyle\pm0.60}$	$1.07{\scriptstyle\pm 0.83}$	$22.80{\scriptstyle\pm2.59}$	$22.90{\scriptstyle\pm1.31}$	0.18±0.14	$16.19{\scriptstyle\pm 0.32}$	$0.35{\scriptstyle\pm 0.28}$	$13.72{\scriptstyle\pm1.80}$	$21.88{\scriptstyle\pm2.73}$	
LE [21]	0.72 ± 0.49	$13.79{\scriptstyle\pm 0.02}$	$1.37{\scriptstyle\pm0.89}$	$21.14{\scriptstyle\pm0.89}$	$20.45{\scriptstyle~\pm 2.64}$	$1.23{\scriptstyle\pm1.01}$	$12.67{\scriptstyle\pm0.62}$	$2.23{\scriptstyle\pm1.66}$	$18.18{\scriptstyle\pm2.41}$	$18.88{\scriptstyle\pm2.24}$	
TMN [25]	0.27 ± 0.16	$\textbf{32.02} \pm 0.51$	$0.54{\scriptstyle\pm0.32}$	$26.53{\scriptstyle\pm0.60}$	$35.22{\scriptstyle\pm1.78}$	0.30 ± 0.29	$31.28{\scriptstyle\pm0.37}$	$0.59{\scriptstyle\pm0.57}$	34.03 ± 1.07	$27.35{\scriptstyle\pm1.43}$	
SymNet [18]	1.96 ± 0.95	$18.47{\scriptstyle\pm0.68}$	$3.54{\scriptstyle\pm1.54}$	$27.24{\scriptstyle\pm2.03}$	$24.47{\scriptstyle~\pm 2.57}$	$2.28{\scriptstyle\pm1.71}$	$17.90{\scriptstyle\pm0.56}$	$4.01{\scriptstyle\pm2.72}$	$27.35{\scriptstyle\pm1.59}$	$23.02{\scriptstyle\pm2.82}$	
CompCos [19]	1.05 ± 0.23	$31.62{\scriptstyle\pm0.54}$	$2.03{\scriptstyle\pm0.44}$	$27.91{\scriptstyle\pm2.68}$	$34.71{\scriptstyle\pm2.69}$	$1.16{\scriptstyle\pm0.55}$	$\textbf{34.32} {\pm} 0.64$	$2.25{\scriptstyle\pm1.02}$	$32.59{\scriptstyle\pm0.37}$	$26.66{\scriptstyle\pm2.82}$	
CGE [22]	4.10 ± 1.09	$17.03{\scriptstyle\pm0.13}$	$6.61{\scriptstyle\pm1.43}$	$25.06{\scriptstyle\pm0.20}$	$31.04{\scriptstyle\pm0.84}$	$2.73{\scriptstyle\pm0.78}$	$19.12{\scriptstyle\pm0.65}$	$4.78{\scriptstyle\pm1.17}$	$25.57{\scriptstyle\pm2.48}$	$23.05{\scriptstyle\pm1.11}$	
MetaCGL (Ours)	11.85 ± 2.55	$20.70{\scriptstyle\pm1.21}$	$15.05{\scriptstyle\pm1.81}$	$31.88{\scriptstyle\pm2.68}$	$\textbf{40.41} {\pm} 1.25$	8.01±0.23	$17.48{\scriptstyle\pm0.98}$	$\boldsymbol{10.99} {\scriptstyle\pm 0.23}$	$32.93{\scriptstyle\pm1.73}$	28.01±1.34	



Figure 4: The ratios of whether error cases from unseen compositions are confused for seen pairs or incorrect unseen pairs. The lower the value, the smaller the trend of overfitting. Best viewed in color.

Specifically, we calculate the metric of the prediction ratio of error cases from unseen compositions as $\frac{U \rightarrow S}{U \rightarrow U}$, where *S* is the number of error cases that are confused for seen pairs, and *U* represents the number of error cases that are confused for incorrect unseen pairs. Different from the evaluation metrics used in the previously reported results, $\frac{U \rightarrow S}{U \rightarrow U}$ reflects the tendency of the model to predict seen compositions when encountering samples belonging to unseen compositions, *i.e.*, it shows the degree of overfitting to seen compositions. Therefore, although $\frac{U \rightarrow S}{U \rightarrow U}$ can not be directly used to represent the performance of the model, we expect that for a

Table 3: Ablation study with various model configurations of MetaCGL. G: the compositional graph. M: the correlation map generating network. BO: the bi-level optimization strategy. CM: the Compositional Mixup data augmentation.

	RL-	CZSL-A	ГTR	RL-CZSL-ACT				
Model	UA	SA	HM	UA	SA	HM		
w/o G	8.64	11.90	9.99	5.16	6.55	5.77		
w/o \mathcal{M}	11.97	20.13	15.02	7.23	16.15	9.99		
w/o BO	1.25	13.70	2.28	0.85	13.88	1.60		
w/o CM	10.82	20.81	14.23	6.52	19.61	9.78		
Full	11.85	20.70	15.05	8.01	17.48	10.99		

better compositional learner, this metric should be lower. We illustrate the results in Figure 4, and it can be observed that methods with better UA, like TMN and CompCos, also have a relatively high $\frac{U \rightarrow S}{U \rightarrow U}$ especially when K_s^c increases. This indicates that they are poor at compositional learning, so increasing the sample size makes them tend to overfit the seen compositions. And MetaCGL keeps a low $\frac{U \rightarrow S}{U \rightarrow U}$, especially on RL-CZSL-ACT. Taken in conjunction with the results from Table 2, this suggests that our MetaCGL is a more positive compositional learner.

5.4 Ablation Study

In Table 3, we examine the effectiveness of each component in MetaCGL. The first observation is that the removal of any component from MetaCGL generally results in a worse performance on UA and HM, which upholds the efficacy of our framework design. Among all the components, the compositional graph plays an important role in recognizing samples of seen compositions, and the bi-level optimization significantly contributes to recognizing unseen compositions. Another phenomenon is that our Compositional Mixup method improves UA while sacrificing SA, and it is understandable as the data augmentation inhibits the further use of samples from seen compositions to evaluate the quality of the fast

	$K_s^c = 1$						$K_s^c = 5$					
Method	RL-CZSL-ATTR			RL-CZSL-ACT		RL-CZSL-ATTR			RL-CZSL-ACT			
	UA	SA	HM	UA	SA	HM	UA	SA	HM	UA	SA	HM
VisProd [21]	1.24	20.99	2.34	0.88	21.97	1.68	0.55	15.83	1.07	0.18	16.19	0.35
+MAML	3.53	20.76	6.03	1.72	21.50	3.18	3.76	29.71	6.67	2.23	28.60	4.13
LE [21]	1.01	14.98	1.89	1.27	14.02	2.32	0.72	13.79	1.37	1.23	12.67	2.23
+MAML	4.49	6.24	5.21	6.14	7.93	6.91	6.06	14.15	8.48	5.68	8.06	6.64
SymNet [18]	1.94	17.34	3.48	2.28	17.90	4.01	1.96	18.47	3.54	2.96	17.12	5.04
+MAML	3.62	4.48	4.00	4.91	4.47	4.65	3.40	4.14	3.69	3.61	4.74	4.10
CompCos [19]	2.57	25.14	4.66	3.02	28.19	5.45	1.05	31.62	2.03	1.16	34.32	2.25
+MAML	3.17	5.93	4.07	2.81	6.44	3.90	2.98	7.67	4.28	3.54	8.76	5.02
CGE [22]	4.65	15.40	7.13	4.05	15.51	6.41	4.10	17.03	6.61	2.73	19.12	4.78
+MAML	9.44	18.62	12.54	6.06	17.86	9.05	11.09	19.76	14.21	6.73	18.17	9.82
MetaCGL (Ours)	10.44	19.01	13.47	7.76	15.95	10.44	11.85	20.70	15.05	8.01	17.48	10.99

Table 4: Comparison with CZSL baselines equipped with MAML. Our MetaCGL still achieves the best UA and HM.



Figure 5: Ablation results of different backbone networks. Deeper backbone does not necessarily lead to better performance. Best viewed in color.

adaptation, so as to combat memorization of seen compositional labels and alleviate the overfitting to seen compositions.

5.5 Effect of Backbone Network

Comparing to our default backbone Conv-4 that is often used in FSL, existing CZSL works prefer to use a deeper ResNet [7] backbone network, e.g., ResNet-18 in [22]. Moreover, [3] proposes that increasing the depth of the backbone network improves the FSL methods by reducing intra-class variation. Therefore, we conduct an ablation study by increasing the backbone network from Conv-4 to ResNet-10 and ResNet-18 when $K_s^c = 5$, exploring whether this can also improve the performance in RL-CZSL. Specifically, ResNet-18 is the same as described by [7] with an input size of 84×84 , while ResNet-10 is a simplified version where only one residual building block is used in each layer. We illustrate the results in Figure 5, and it can be observed that while the tendency of the same method is quite unstable on different datasets, different methods also show no consistent pattern on the same dataset when the backbone deepens. In other words, a deeper backbone network will not necessarily result in a boost in RL-CZSL if not paired with the suitable approach.

5.6 Effect of Equipping MAML

From the ablation study results of our MetaCGL, it can be observed that the bi-level optimization significantly contributes to recognizing unseen compositions. Therefore, to study whether such bi-level optimization can also improve the baselines on the RL-CZSL task, we equip these methods with the popular bi-level optimization method named model-agnostic meta-learning (MAML) [6]. Note that TMN [25] is not considered as the bi-level optimization will double its large computation overhead, making the computing time unaffordable. We compare the results with our MetaCGL in Table 4. It can be observed that, on the one hand, VisProd and CGE consistently gain from equipping MAML. On the other hand, MAML significantly hurts the seen accuracy of other methods (LE, SymNet, and CompCos) without bringing a stable gain in unseen accuracy, and the performance of SymNet even degenerates to close to that of random prediction. On top of that, our MetaCGL still achieves the best UA and HM in all cases. The results of this experiment show that the bi-level optimization is not a silver bullet for the RL-CZSL problem, and more effective solutions remain to be explored.

6 CONCLUSION

In this paper, we introduce reference-limited compositional zeroshot learning (RL-CZSL), a novel and non-trivial task that mimics the naturalistic unseen environment for compositional learners. Aiming to recognize unseen compositional concepts in the scarcity of referential compositions and samples, we propose a Meta Compositional Graph Learner (MetaCGL) that can efficiently learn the compositionality in the new environment. Moreover, we also build two large-scale benchmark datasets to drive research on the task. We show that in the challenging RL-CZSL setting, our MetaCGL significantly outperforms the state-of-the-art CZSL methods in recognizing unseen compositions, while substantial research space still remains. We hope our work can facilitate and calibrate the development of compositional learning systems.

Acknowledgement This work was supported by STI 2030—Major Projects (2022ZD0208800), and NSFC General Program (Grant No. 62176215).

Reference-Limited Compositional Zero-Shot Learning

REFERENCES

- Dzmitry Bahdanau, Shikhar Murty, Michael Noukhovitch, Thien Huu Nguyen, Harm de Vries, and Aaron C. Courville. 2019. Systematic Generalization: What Is Required and Can It Be Learned?. In Proceedings of the International Conference on Learning Representations.
- [2] Yu-Wei Chao, Zhan Wang, Yugeng He, Jiaxuan Wang, and Jia Deng. 2015. HICO: A Benchmark for Recognizing Human-Object Interactions in Images. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 1017–1025.
- [3] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. 2019. A Closer Look at Few-shot Classification. In Proceedings of the International Conference on Learning Representations.
- [4] Noam Chomsky. 1957. Logical structures in language. American Documentation 8, 4 (1957), 284.
- [5] Guneet Singh Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. 2020. A Baseline for Few-Shot Image Classification. In Proceedings of the International Conference on Learning Representations.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In Proceedings of the International Conference on Machine Learning. 1126–1135.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 770–778.
- [8] Zhi Hou, Xiaojiang Peng, Yu Qiao, and Dacheng Tao. 2020. Visual Compositional Learning for Human-Object Interaction Detection. In Proceedings of the European Conference on Computer Vision. 584–600.
- [9] Siteng Huang, Min Zhang, Yachen Kang, and Donglin Wang. 2021. Attributes-Guided and Pure-Visual Attention Alignment for Few-Shot Recognition. In Proceedings of the AAAI Conference on Artificial Intelligence. 7840–7847.
- [10] Wilhelm von Humboldt. 1988. The diversity of human language-structure and its influence on the mental development of mankind.
- [11] Phillip Isola, Joseph J. Lim, and Edward H. Adelson. 2015. Discovering states and transformations in image collections. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 1383–1391.
- [12] Keizo Kato, Yin Li, and Abhinav Gupta. 2018. Compositional Learning for Human Object Interaction. In Proceedings of the European Conference on Computer Vision. 247–264.
- [13] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In Proceedings of the International Conference on Learning Representations.
- [14] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In Proceedings of the International Conference on Learning Representations.
- [15] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *International Journal of Computer Vision* 123, 1 (2017), 32–73.
- [16] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. 2018. Learning to Generalize: Meta-Learning for Domain Generalization. In Proceedings of the AAAI Conference on Artificial Intelligence. 3490–3497.
- [17] Xiangyu Li, Xu Yang, Kun Wei, Cheng Deng, and Muli Yang. 2022. Siamese Contrastive Embedding Network for Compositional Zero-Shot Learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 9316–9325.
- [18] Yonglu Li, Yue Xu, Xiaohan Mao, and Cewu Lu. 2020. Symmetry and Group in Attribute-Object Compositions. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 11313–11322.
- [19] Massimiliano Mancini, Muhammad Ferjad Naeem, Yongqin Xian, and Zeynep Akata. 2021. Open World Compositional Zero-Shot Learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 5222–5230.
- [20] Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In Proceedings of the Advances in Neural Information Processing Systems. 3111–3119.
- [21] Ishan Misra, Abhinav Gupta, and Martial Hebert. 2017. From Red Wine to Red Tomato: Composition with Context. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 1160–1169.
- [22] Muhammad Ferjad Naeem, Yongqin Xian, Federico Tombari, and Zeynep Akata. 2021. Learning Graph Embeddings for Compositional Zero-shot Learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 953–962.
- [23] Tushar Nagarajan and Kristen Grauman. 2018. Attributes as Operators: Factorizing Unseen Attribute-Object Compositions. In Proceedings of the European Conference on Computer Vision. 172–190.
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith

Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Proceedings of the Advances in Neural Information Processing Systems. 8024–8035.

- [25] Senthil Purushwalkam, Maximilian Nickel, Abhinav Gupta, and Marc'Aurelio Ranzato. 2019. Task-Driven Modular Networks for Zero-Shot Compositional Learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 3592–3601.
- [26] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. 2020. Rapid Learning or Feature Reuse? Towards Understanding the Effectiveness of MAML. In Proceedings of the International Conference on Learning Representations.
- [27] Nirat Saini, Khoi Pham, and Abhinav Shrivastava. 2022. Disentangling Visual Embeddings for Attributes and Objects. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 13648–13657.
- [28] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical Networks for Few-shot Learning. In Proceedings of the Advances in Neural Information Processing Systems. 4077–4087.
- [29] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H. S. Torr, and Timothy M. Hospedales. 2018. Learning to Compare: Relation Network for Few-Shot Learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 1199–1208.
- [30] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B. Tenenbaum, and Phillip Isola. 2020. Rethinking Few-Shot Image Classification: A Good Embedding is All You Need?. In Proceedings of the European Conference on Computer Vision. 266–282.
- [31] Pavel Tokmakov, Yu-Xiong Wang, and Martial Hebert. 2019. Learning compositional representations for few-shot recognition. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 6372–6381.
- [32] Ramakrishna Vedantam, Arthur Szlam, Maximilian Nickel, Ari Morcos, and Brenden M. Lake. 2021. CURI: A Benchmark for Productive Concept Learning Under Uncertainty. In Proceedings of the International Conference on Machine Learning. 10519–10529.
- [33] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching Networks for One Shot Learning. In Proceedings of the Advances in Neural Information Processing Systems. 3630–3638.
- [34] Yaqing Wang, Quanming Yao, James T. Kwok, and Lionel M. Ni. 2020. Generalizing from a Few Examples: A Survey on Few-shot Learning. *Comput. Surveys* 53, 3 (2020), 63:1–63:34.
- [35] Mark Yatskar, Luke S. Zettlemoyer, and Ali Farhadi. 2016. Situation Recognition: Visual Semantic Role Labeling for Image Understanding. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 5534–5542.
- [36] Aron Yu and Kristen Grauman. 2014. Fine-Grained Visual Comparisons with Local Learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 192–199.
- [37] Matthew D. Zeiler and Rob Fergus. 2014. Visualizing and Understanding Convolutional Networks. In Proceedings of the European Conference on Computer Vision. 818–833.
- [38] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. 2018. mixup: Beyond Empirical Risk Minimization. In Proceedings of the International Conference on Learning Representations.
- [39] Tian Zhang, Kongming Liang, Ruoyi Du, Xian Sun, Zhanyu Ma, and Jun Guo. 2022. Learning Invariant Visual Representations for Compositional Zero-Shot Learning. In Proceedings of the European Conference on Computer Vision. 339–355.
- [40] Yixiong Zou, Shanghang Zhang, Ke Chen, Yonghong Tian, Yaowei Wang, and José M. F. Moura. 2020. Compositional Few-Shot Recognition with Primitive Discovery and Enhancing. In Proceedings of the ACM International Conference on Multimedia. 156–164.