

Symmetry and Group in Attribute-Object Compositions

Yong-Lu Li, Yue Xu, Xiaohan Mao, Cewu Lu*

Shanghai Jiao Tong University

{yonglu.li, silicxuyue, mxh1999, lucewu}@sjtu.edu.cn

Abstract

Attributes and objects can compose diverse compositions. To model the compositional nature of these general concepts, it is a good choice to learn them through transformations, such as coupling and decoupling. However, complex transformations need to satisfy specific principles to guarantee the rationality. In this paper, we first propose a previously ignored principle of attribute-object transformation: **Symmetry**. For example, coupling peeled-apple with attribute peeled should result in peeled-apple, and decoupling peeled from apple should still output apple. Incorporating the symmetry principle, a transformation framework inspired by group theory is built, i.e. SymNet. SymNet consists of two modules, Coupling Network and Decoupling Network. With the group axioms and symmetry property as objectives, we adopt Deep Neural Networks to implement SymNet and train it in an end-to-end paradigm. Moreover, we propose a Relative Moving Distance (RMD) based recognition method to utilize the **attribute change** instead of the attribute pattern itself to classify attributes. Our symmetry learning can be utilized for the Compositional Zero-Shot Learning task and outperforms the state-of-the-art on widely-used benchmarks. Code is available at <https://github.com/DirtyHarryLYL/SymNet>.

1. Introduction

Attributes describe the properties of generic objects, e.g. material, color, weight, etc. Understanding the attributes would directly facilitate many tasks that require deep semantics, such as scene graph generation [15], object perception [33, 9, 40, 35, 7, 20], human-object interaction detection [3, 16, 17]. As side information, attributes can also be employed in zero-shot learning [8, 39, 38, 37, 12, 41].

Going along with the road of conventional classification setting, some works [14, 28, 39, 29] address attribute

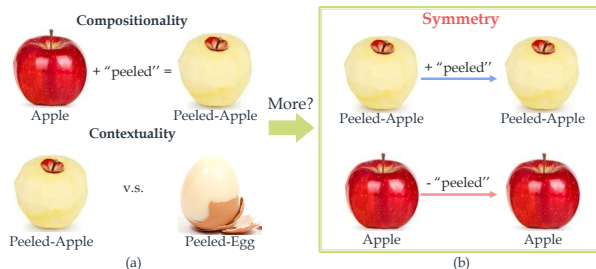


Figure 1. Except for the compositionality and contextuality, attribute-object compositions also have the *symmetry* property. For instance, a peeled-apple should not change after “adding” the peeled attribute. Similarly, an apple should keep the same after “removing” the peeled attribute because it does not have it.

recognition with the typical discriminative models for objects and achieve poor performance. This is because attributes cannot be well expressed independently of the context [24, 25] (Fig. 1(a)). Subsequently, researchers rethink the nature of attributes and treat them as linear operations [25] to operate these two general concepts, e.g. “adding” attribute to object (coupling) or “removing” attribute from objects (decoupling). Though such new insight has promoted this field, the current “add-remove” system is not complete and lacks an axiomatics foundation to satisfy the specific principles of nature. In this paper, we rethink the *physical* and *linguistic* properties of attribute-object, and propose a previously ignored but important principle of attribute-object transformations: **symmetry**, which would promote attribute-object learning. Symmetry depicts the invariance under transformations, e.g. a circle has rotational symmetry under the rotation without changing its appearance. The transformation that “adding” or “removing” attributes should also satisfy the symmetry: An object should remain unchanged if we “add” an attribute which it already has, or “remove” an attribute which it does not have. For instance, a peeled-apple keeps invariant if we “add” attribute peeled upon it. Similarly, “removing” peeled from apple would still result in apple.

As shown in Fig. 1(b), except the compositionality and contextuality, the symmetry property should also be satisfied to guarantee the rationality. In view of this, we first

*Cewu Lu is the corresponding author, member of Qing Yuan Research Institute and MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, China.

introduce the symmetry and propose **SymNet** to depict it. In this work, we aim to bridge attribute-object learning and group theory. Because the elegant properties of group theory would largely help in a more principled way, given its great theoretical potential. Thus, to cover the principles existing in transformations theoretically, the principles from group theory are borrowed to model the symmetry. In detail, we define three transformations {“keep”, “add”, “remove”} and an operation to perform three transformations upon objects, to construct a “group”. To implement these, SymNet adopts Coupling Network (CoN) and Decoupling Network (DecoN) to perform coupling/adding and decoupling/removing. On the other hand, to meet the fundamental requirements of group theory, *symmetry* and the group axioms *closure*, *associativity*, *identity element*, *invertibility element* are all implemented as the learning objectives of SymNet. Naturally, SymNet considers the compositionality and contextuality during coupling and decoupling of various attributes and objects. All above principles will be learned under a unified model in an end-to-end paradigm.

With symmetry learning, we can apply SymNet to address the Compositional Zero-Shot Learning (CZSL), whose target is to classify the unseen compositions composed of seen attributes and objects. We adopt a novel recognition paradigm, **Relative Moving Distance (RMD)** (Fig. 2). That is, given a specific attribute, an object would be manipulated by the “add” and “remove” transformations parallelly in latent space. When those transformations meet the symmetry principle: if the input object already has the attribute, the output after addition should be close to the original input object, and the object after removal should be far from the input. Contrarily, if the object does not have the given attribute, the object after removal should be closer to the input than the object after addition. Thus, attribute classification can be accomplished concurrently by comparing the relative *moving* distances between the input and two outputs. With RMD recognition, we can utilize the robust *attribute change* to classify the attributes, instead of only relying on the dramatically unstable *visual attribute patterns*. Extensive experiments show that our method achieves significant improvements on CZSL benchmarks [12, 41].

The main contributions of this paper are: 1) We propose a novel property of attribute-object composition transformation: symmetry, and design a framework inspired by group theory to learn it under the supervision of group axioms. 2) Based on symmetry learning, we propose a novel method to infer the attributes based on Relative Moving Distance. 3) We achieve substantial improvements in attribute-object composition zero-shot learning tasks.

2. Related Work

Visual Attribute. Visual attribute was introduced into computer vision to reduce the gap between visual patterns

and object concepts, such as reducing the difficulty in object recognition [8] or acting as an intermediate representation for zero-shot learning [14]. After that, attribute has been widely applied in recognition of face [19], people [1], pedestrian [6] or action [34], person Re-ID [18, 30], zero-shot learning [37, 39], caption generation [13, 27] and so on. Therefore, attribute recognition is a fundamental problem to promote the visual concept understanding.

The typical approach for attribute recognition is to train a multi-label discriminative model same as object classification [14, 28, 39, 29], which ignores the intrinsic properties of attributes, such as compositionality and contextuality. Farhadi *et al.* [8] propose a visual feature selection method to recognize the attributes, under the consideration of cross-category generalization. Later, some works start to consider the properties by exploiting the attribute-attribute or attribute-object correlations [11, 4, 22]. Considering the contextuality of attributes, Nagarajan *et al.* [25] regard attributes as linear transformations operated upon object embeddings, and Misra *et al.* [24] map the attributes into model weight space to attain better representations.

Compositional Zero-Shot Learning. CZSL is a crossing filed of compositional learning and zero-shot learning. In the CZSL setting, test compositions are unseen during training, while each component is seen in both training set and test set. Chen *et al.* [4] construct linear classifiers for unseen compositions with tensor completion of weight vectors. Misra *et al.* [24] consider that the model space is more smooth, thus project the attributes or objects into model space by training binary linear SVMs for the corresponding components. To deal with the CZSL task, it composes the attribute and object embeddings in model space as composition representation. Wang *et al.* [36] address the attribute-object compositional problem via conditional embedding modification which relies on attribute word embedding [23] transformation. Nan *et al.* [26] map the image features and word vectors [31] into embedding space with the reconstruction constraint. Nagarajan *et al.* [25] regard attributes as linear operations for object embedding and map the image features and transformed object embeddings into a shared latent space. However, linear and explicit matrix transformation may be insufficient to represent various attribute concepts of different complexity, *e.g.* representing “red” and “broken” as matrices with the same capacity. Previous methods usually ignored or incompletely considered the natural principles within the coupling and decoupling of attributes and objects. In light of this, we propose a unified framework inspired by group theory to learn these important principles such as symmetry.

3. Approach

Fig. 2 gives an overview of our approach. Our goal is to learn the symmetry within attribute-object compositions.

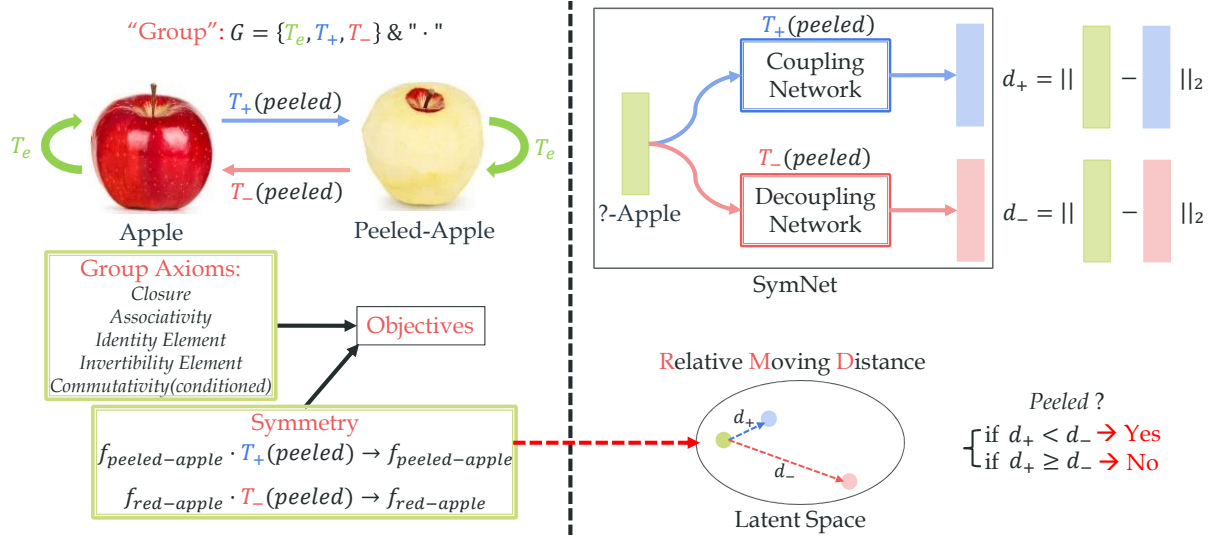


Figure 2. Overview of our proposed method. We construct a “group” to learn the symmetry and operate the composition learning.

Thus we can utilize it to obtain a deeper understanding of attribute-object, *e.g.*, to address the CZSL task [12, 41]. To learn the symmetry in transformations, we need a comprehensive framework to cover all principles. Inspired by the group theory, we define a unified model named SymNet.

We define $G = \{T_e, T_+, T_-\}$ which contains identity (“keep”), coupling (“add”) and decoupling (“remove”) transformations (Sec. 3.1) for each specific attribute and utilize Deep Neural Networks to implement them (Sec. 3.2). To depict symmetry theoretically, it is a natural choice to adopt group theory as the close associations between symmetry and group in physics and mathematics. Since a group should satisfy the group axioms, *i.e.*, *closure*, *associativity*, *identity element*, and *invertibility element*, we construct the learning objectives based on these axioms to train the transformations (Sec. 3.3). In addition, SymNet also satisfies the *commutativity* under conditions. With the above constraints, we can naturally guarantee compositionality and contextuality. *Symmetry* allows us to use a novel method, Relative Moving Distance, to identify whether an object has a certain attribute with the help of T_+ and T_- (Sec. 3.4) for CZSL task (Sec. 3.5).

3.1. Group Definition

To depict the symmetry, we need to first define the transformations. Naturally, we need two reciprocal transformations to “add” and “remove” the attributes. Further, we need an axiomatic system to restrain the transformations and keep the rationality. Thus, we define three transformations $G = \{T_e, T_+, T_-\}$ and an operation “.”. In practice, it is difficult to strictly follow the theory considering the physical and linguistic truth. For example, the operation between attribute transformations “peeled · broken” is

odd. Thus the “operation” here is defined to be operated upon object only.

Definition 1. Identity transformation T_e keep the attributes of object. **Coupling** transformation T_+ couples a specific attribute with an object. **Decoupling** transformation T_- decouples a specific attribute from an object.

Definition 2. Operation “.” performs transformations $\{T_e, T_+, T_-\}$ upon object. Noticeably, operation “.” is not the dot product and we use this notation to maintain the consistence with group theory.

More formally, for object $o \in \mathcal{O}$ and attribute $a^i, a^j \in \mathcal{A}$, $a^i \neq a^j$, where \mathcal{O} denotes object set and \mathcal{A} denotes attribute set, operation “.” performs transformations in G upon an object/image embedding:

$$\begin{aligned} f_o^i \cdot T_+(a^j) &= f_o^{ij}, \\ f_o^{ij} \cdot T_-(a^j) &= f_o^i, \\ f_o^i \cdot T_e &= f_o^i, \end{aligned} \quad (1)$$

where f_o^i means o has one attribute a^i and f_o^{ij} means o has two attributes a^i, a^j . Here we do not sign a specific object category and use o for simplicity.

Definition 3. G has the **symmetry** property if and only if $\forall a^i, a^j \in \mathcal{A}, a^i \neq a^j$:

$$f_o^i = f_o^i \cdot T_+(a^i), \quad f_o^i = f_o^i \cdot T_-(a^j). \quad (2)$$

3.2. Group Implementation

In practice, when performing T_e upon f_o^i , we directly use f_o^i as the $f_o^i \cdot T_e$ to implement the identity transformation for

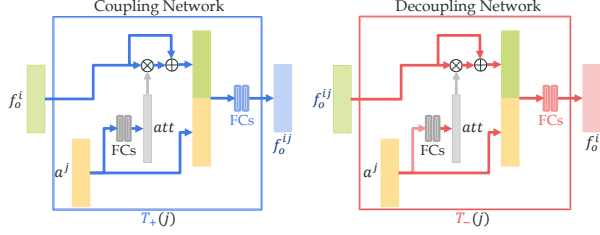


Figure 3. The structure of CoN and DecoN. They take the attribute embedding to assign a specific attribute a^j . f_o^i, f_o^{ij} are the object embeddings extracted from ResNet-18 [10].

simplicity. For other two transformations T_+ , T_- , we propose **SymNet** consists of two modules: Coupling Network (**CoN**) and Decoupling Network (**DecoN**). CoN and DecoN have the same structure but independent weights and are trained with different tasks. As seen in Fig. 3, CoN and DecoN both take the image embedding f_o^i of an object and the embedding of attribute a^j as inputs, and output the transformed object embedding. We use the attribute category word embeddings such as GloVe [31] or onehot vector to represent the attributes. f_o^i is extracted by an ImageNet [5] pre-trained ResNet [10] from image I , i.e. $f_o^i = F_{res}(I)$.

Intuitively, attributes affect objects in different ways, e.g. “red” changes the color, “wooden” changes the texture. In CoN and DecoN, we use an *attribute-as-attention* strategy, i.e. using $att = g(a^j)$ as attention, where $g(\cdot)$ means two fully-connected (FC) and a Softmax layers. We concatenate $f_o^i \circ att + f_o^i$ with original a^j as the input and use two FC layers to perform the transformation.

3.3. Group Axioms as Objectives

According to group theory, SymNet should satisfy four group axioms: *closure*, *associativity*, *identity element*, and *invertibility*. Under certain conditions, attribute-object also satisfy *commutativity*. Besides, SymNet must obey the symmetry property of the attribute transformations.

In practice, we use Deep Neural Networks to implement transformations. Thus, we can construct training objectives to approach the theoretic transformations following the axioms. Considering the *actual characteristics* of attribute-object compositions, we slightly adjust the axioms to construct the objectives. Besides, there are two situations with different forms of axioms: 1) coupling or decoupling an attribute a^i that the object f_o^i already has, or 2) coupling or decoupling an attribute a^j that object f_o^i does not have.

Symmetry. First of all, SymNet should satisfy the symmetry property as depicted in Eq. 2, i.e., $f_o^i = f_o^i \cdot T_+(a^i)$, $f_o^i = f_o^i \cdot T_-(a^j)$. The symmetry is essential to keep the semantic meaning during coupling and decoupling. For example, given a `peeled-egg`, adding the attribute `peeled` again should not change the object state. Similarly, a `cup` without attribute `broken` should remain unchanged after removing

broken. Thus, we construct the **symmetry loss**:

$$\mathcal{L}_{sym} = \|f_o^i - f_o^i \cdot T_+(a^i)\|_2 + \|f_o^i - f_o^i \cdot T_-(a^j)\|_2. \quad (3)$$

where $a^i, a^j \in \mathcal{A}, i \neq j$. We use L_2 norm loss to measure the distance between two embeddings.

Closure. For all elements in set G , their operation results should also be in G . In SymNet, for the attribute a^i that f_o^i has, $f_o^i \cdot T_+(a^i) \cdot T_-(a^i)$ should be equal to $f_o^i \cdot T_-(a^i)$. For the attribute a^j that f_o^i does not have, $f_o^i \cdot T_-(a^j) \cdot T_+(a^j)$ should be equal to $f_o^i \cdot T_+(a^j)$. Thus, we construct:

$$\begin{aligned} \mathcal{L}_{clo} = & \|f_o^i \cdot T_+(a^i) \cdot T_-(a^i) - f_o^i \cdot T_-(a^i)\|_2 + \\ & \|f_o^i \cdot T_-(a^j) \cdot T_+(a^j) - f_o^i \cdot T_+(a^j)\|_2, \end{aligned} \quad (4)$$

Identity Element. The properties of identity element T_e are automatically satisfied since we implement T_e as a skip connection, i.e. $f_o^i \cdot T_*(a^i) \cdot T_e = f_o^i \cdot T_e \cdot T_*(a^i) = f_o^i \cdot T_*(a^i)$ where T_* denotes any element in G .

Invertibility Element. According to the definition, T_+ is the invertibility element of T_- , vice versa. For the attribute a^i that f_o^i has, $f_o^i \cdot T_-(a^i) \cdot T_+(a^i)$ should be equal to $f_o^i \cdot T_e = f_o^i$. For the attribute a^j that f_o^i does not have, $f_o^i \cdot T_+(a^j) \cdot T_-(a^j)$ should be equal to $f_o^i \cdot T_e = f_o^i$. Therefore, we have:

$$\begin{aligned} \mathcal{L}_{inv} = & \|f_o^i \cdot T_+(a^j) \cdot T_-(a^j) - f_o^i \cdot T_e\|_2 + \\ & \|f_o^i \cdot T_-(a^i) \cdot T_+(a^i) - f_o^i \cdot T_e\|_2. \end{aligned} \quad (5)$$

Associativity. In view of the practical physical meaning of attribute-object compositions, we only define the operation “.” that operates a transformation upon an object embedding in Sec. 3.1, but do not define the operation between transformations. Therefore, we relax the constraint here and do not construct an objective according to associativity in practice.

Commutativity. Because of the speciality of attribute-object, SymNet satisfies the *commutativity* when coupling and decoupling *multiple* attributes. Thus, $f_o^i \cdot T_+(a^i) \cdot T_-(a^j)$ should be equal to $f_o^i \cdot T_-(a^j) \cdot T_+(a^i)$:

$$\begin{aligned} \mathcal{L}_{com} = & \|f_o^i \cdot T_+(a^i) \cdot T_-(a^j) - \\ & f_o^i \cdot T_-(a^j) \cdot T_+(a^i)\|_2. \end{aligned} \quad (6)$$

Although above definitions do not strictly follow the theory, but the *loosely* conducted axiom objectives still contribute to the robustness and effectiveness a lot (ablation study in Sec. 3.5) and open a door to a more theoretical way.

The last but not the least, CoN and DecoN need to keep the **semantic consistency**, i.e. before and after the transformation, the *object category* should not change. Hence we use a cross-entropy loss \mathcal{L}_{cls}^o for the object recognition of the input and output embeddings of CoN and DecoN. In the same way, before and after coupling and decoupling, the *attribute changes* provide the attribute classification loss \mathcal{L}_{cls}^a . We use typical visual pattern-based classifiers consisting of FC layers for the object and attribute classifications.

3.4. Relative Moving Distance

As shown in Fig. 4, we utilize the Relative Moving Distance (RMD) based on the symmetry property to operate the attribute recognition. Given an image embedding f_o^x of an object with unknown attribute a^x , we first input it to both CoN and DecoN with all kinds of attributes word embeddings $\{a^1, a^2, \dots, a^n\}$ where n is the number of attributes. Two transformers would take attribute embeddings as conditions and operate the coupling and decoupling in parallel, then output $2n$ transformed embeddings $\{f_o^x \cdot T_+(a^1), f_o^x \cdot T_+(a^2), \dots, f_o^x \cdot T_+(a^n)\}$ and $\{f_o^x \cdot T_-(a^1), f_o^x \cdot T_-(a^2), \dots, f_o^x \cdot T_-(a^n)\}$. We compute the distances between f_o^x and the transformed embeddings:

$$\begin{aligned} d_+^i &= \|f_o^x - f_o^x \cdot T_+(a^i)\|_2, \\ d_-^i &= \|f_o^x - f_o^x \cdot T_-(a^i)\|_2. \end{aligned} \quad (7)$$

To compare two distances, we define *Relative Moving Distance* as $d^i = d_-^i - d_+^i$ and perform binary classification for each attribute (Fig. 4): 1) If $d^i \geq 0$, i.e. $f_o^x \cdot T_+(a^i)$ is closer to f_o^x than $f_o^x \cdot T_-(a^i)$, we tend to believe f_o^x has attribute a^i . 2) If $d^i < 0$, i.e. $f_o^x \cdot T_-(a^i)$ is closer, we tend to predict that f_o^x does not have attribute a^i . Previous zero/few-shot learning methods usually classify the instances via measuring the distance between the embedded instances and **fixed** points like prototype/label/centroid embeddings. Differently, Relative **Moving** Distance compares the distance before and after applying the coupling and decoupling operation.

Training. To enhance the RMD-based classification performance, we further use a triplet loss function. Let \mathcal{X} denote the set of attributes that f_o^x has, the loss can be described as:

$$\mathcal{L}_{tri} = \sum_i^{\mathcal{X}} [d_+^i - d_-^i + \alpha]_+ + \sum_j^{\mathcal{A}-\mathcal{X}} [d_-^j - d_+^j + \alpha]_+, \quad (8)$$

where $\alpha=0.5$ is triplet margin. d_+^i should be less than d_-^i for the attributes that f_o^x has and greater than d_-^i for the attributes f_o^x does not have. The total loss of SymNet is

$$\begin{aligned} \mathcal{L}_{total} &= \lambda_1 \mathcal{L}_{sym} + \lambda_2 \mathcal{L}_{axiom} \\ &+ \lambda_3 \mathcal{L}_{cls}^a + \lambda_4 \mathcal{L}_{cls}^o + \lambda_5 \mathcal{L}_{tri}, \end{aligned} \quad (9)$$

where $\mathcal{L}_{axiom} = \mathcal{L}_{clo} + \mathcal{L}_{inv} + \mathcal{L}_{com}$.

Inference. In practice, for n attribute categories, we use RMDs $d = \{d^i\}_{i=1}^n$ as the attribute scores, i.e. $\mathcal{S}_a = \{\mathcal{S}_a^i\}_{i=1}^n = \{d^i\}_{i=1}^n$ and obtain attribute probability with Sigmoid function: $p_a^i = Sigmoid(\mathcal{S}_a^i)$. Notably, we also consider the scale and use a factor γ to adjust the score before Sigmoid. Our method can be operated in parallel, i.e., we simultaneously compute the RMD values of n attributes. We input $[B, n, 300]$ sized tensor where B is the mini-batch size and 300 is the object embedding size. CoN and DecoN would output two $[B, n, 300]$ sized embeddings

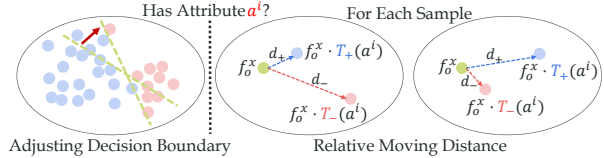


Figure 4. Comparison between typical method and our Relative Moving Distance (RMD) based recognition. Previous methods mainly try to adjust the decision boundary in latent space. Our RMD based approach moves the embedding point with T_+ and T_- and classifies by comparing their moving distances.

after transformation. Then we can compute RMDs $\{d^i\}_{i=1}^n$ at the same time. Our method has approximately the same speed as the typical FC classifier. The inference speed from features to RMD is about 41.0 FPS and the FC classifier speed is about 45.8 FPS. The gap can be further omitted if considering the overhead of the feature extractor.

3.5. Discussion: Composition Zero-Shot Learning

With robust and effective symmetry learning for attribute-object, we can further apply SymNet to CZSL [12, 41]. The goal of CZSL is to infer the unseen attribute-object pairs in test set, i.e. a prediction is true positive if and only if both attribute and object classifications are accurate. The pair candidates are available during testing, thus the predictions of impossible pairs can be masked.

We propose a novel method to address this task based on Relative Moving Distance (RMD). With Relative Moving Distance $d^i = d_-^i - d_+^i$, the probabilities of attribute category are computed as $p_a^i = Sigmoid(d^i)$. For object category, we input the object embedding to 2-layer FC with Softmax to obtain the object scores $\mathcal{S}_o = \{\mathcal{S}_o^j\}_{j=1}^m$, where m is the number of object categories. The object category probability $p_j = Softmax(\mathcal{S}_o^j)$ and $p_o = \{p_o^i\}_{i=1}^m$. We then use p_{ao}^{ij} to represent the probability of an attribute-object pair in test set which is composed of the i -th attribute category and j -th object. The pair probabilities are given by $p_{ao}^{ij} = p_a^i \times p_o^j$. The impossible compositions would be masked according to the benchmarks [12, 41].

4. Experiment

4.1. Data and Metrics

Our experiments are conducted on MIT-States [12] and UT-Zappos50K [41]. MIT-States contains 63440 images covering 245 objects and 115 attributes. Each image is attached with one single object-attribute composition label and there are 1262 possible pairs in total. We follow the setting of [24] and use 1262 pairs/34562 images for training and 700 pairs/19191 images as the test set. UT-Zappos50K is a fine-grained dataset with 50025 images of shoes annotated with shoe type-material pairs. We follow the setting and split from [25], using 83 object-attribute pairs/24898

images as train set and 33 pairs/4228 images for testing. The training and testing pairs are non-overlapping for both datasets, *i.e.* the test set contains unseen attribute-object pairs composed of seen attributes and objects. We report the Top-1, 2, 3 accuracies on the unseen test set as evaluation metrics. We also evaluate our model under the generalized CZSL setting of TMN [32], since the "open world" setting from [25] brings biases towards unseen pairs [2].

4.2. Baselines

We compare SymNet with baselines following [24] and [25], as well as previous state-of-the-arts. If not specified, the adopted methods are based on ResNet-18 backbone.

Visual Product trains two simple classifiers for attributes and objects independently and fuses the outputs by multiplying their margin probabilities: $P(a, o) = P(a)P(o)$. The classifiers can be either linear SVMs [24] or single layer softmax regression models [25].

LabelEmbed (LE) is proposed by [24]. It combines the word vectors [31] of attribute and object and uses 3-layer FC to transform the pair embedding into a transform matrix. The classification score is the product of transform matrix and visual feature: $\mathcal{T}(e_a, e_b)^T \phi(I)$. It has three variants:

1. **LabelEmbed Only Regression (LEOR)** [24] changes the target to minimize the Euclidean distance between $\mathcal{T}(e_a, e_b)$ and the weight of pair SVM classifier w_{ab} .
2. **LabelEmbed With Regression (LE+R)** [24] combines the losses of LE and LEOR aforementioned.
3. **LabelEmbed+** [25] embeds the attribute, object vectors, and image features into a semantic space and also optimizes the input representations during training.

AnalogousAttr [4] trains linear classifiers for seen compositions and uses tensor completion to generalize to the unseen pairs. We report the reproduced results from [25].

Red Wine [24] uses SVM weights as the attribute or object embeddings to replace the word vectors in LabelEmbed.

AttrOperator [25] regards attributes as linear transformations and object word vectors [31] after transformation as pair embeddings. It takes the pair with the closest distance to the image feature as the recognition result. Besides the top-1 accuracy directly reported in [25], we evaluate the top-2, 3 accuracies with the open-sourced code.

TAFE-Net [36] uses word vectors [23] of attribute-object pair as task embedding of its meta learner. It generates a binary classifier for each existing composition. We report the results based on VGG-16 which is *better* and more complete than the result based on ResNet-18.

GenModel [26] projects the visual features of images and semantic language embeddings of pairs into a shared latent space. The prediction is given by comparing the distance between visual features and all candidate pair embeddings.

TMN [32] adopts a set of small FC-based modules and configure them via a gating function in a task-driven way. It can

Method	MIT-States			UT-Zappos		
	Top-1	Top-2	Top-3	Top-1	Top-2	Top-3
Visual Product [24]	9.8/13.9*	16.1	20.6	49.9*	/	/
LabelEmbed (LE) [24]	11.2/13.4*	17.6	22.4	25.8*	/	/
- LEOR [24]	4.5	6.2	11.8	/	/	/
- LE + R [24]	9.3	16.3	20.8	/	/	/
- LabelEmbed+ [25]	14.8*	/	/	37.4*	/	/
AnalogousAttr [4]	1.4	/	/	18.3	/	/
Red Wine [24]	13.1	21.2	27.6	40.3	/	/
AttrOperator [25]	14.2	19.6	25.1	46.2	56.6	69.2
TAFE-Net [36]	16.4	26.4	33.0	33.2	/	/
GenModel [26]	17.8	/	/	48.3	/	/
SymNet (Ours)	19.9	28.2	33.8	52.1	67.8	76.0

Table 1. Results of CZSL on MIT-States and UT-Zappos.

be generalized to unseen pairs via re-weighting these primitive modules.

4.3. Implementation Details

For two datasets, we use ImageNet pre-trained ResNet-18 [10] as the backbone to extract image features and do not fine-tune it following previous methods. We use the 300-dimensional pre-trained GloVe [31] vectors as the word embeddings. The 512-dimensional ResNet-18 feature is first transformed to 300-dimensional by a single FC. The main modules of our SymNet, CoN and DecoN, have the same structures but independent weights as depicted in Fig. 3: two FC layers of sizes 768/300 with Sigmoid convert the attribute embedding to 300-dimensional attention and be multiplied to the input image representation. The representation after attention is concatenated to the attribute embedding and then compressed to the original dimension by the other two 300-sized FC layers. Each hidden FC in CoN and DecoN is followed by BatchNorm and ReLU layers.

For each training image, we randomly sample another image with the same object label but different attribute as the negative sample to compute the losses (Sec. 3.3). We train SymNet with SGD optimizer on single NVIDIA GPU. We use cross-validation to determine the hyper-parameters, *e.g.*, learning rate, weights, epochs. For MIT-States, the model is trained with learning rate $5e-4$ and batch size 512 for 320 epochs. The loss weights are $\lambda_1 = 0.05, \lambda_2 = 0.01, \lambda_3 = 1, \lambda_4 = 0.01, \lambda_5 = 0.03$. For UT-Zappos, the model is trained with learning rate $1e-4$ and batch size 256 for 600 epochs. The loss weights are $\lambda_1 = 0.01, \lambda_2 = 0.03, \lambda_3 = 1, \lambda_4 = 0.5, \lambda_5 = 0.5$. Notably, the weights on two datasets are different. Because MIT-States has diverse attributes and objects, while UT-Zappos contains similar fine-grained shoes. Different range and scale lead to distinct embedding spaces and different parameters for RMD.

4.4. Compositional Zero-Shot Learning

To evaluate the symmetry learning in compositional zero-shot task, we conduct experiments on widely-used benchmarks: MIT-States [12] and UT-Zappos [41].

Composition Learning. The results of CZSL are shown in Tab. 1, where the first five rows are baselines from

Model	Val AUC			Test AUC			Seen	Unseen	HM
	1	2	3	1	2	3			
AttrOperator [25]	2.5	6.2	10.1	1.6	4.7	7.6	14.3	17.4	9.9
Red Wine [24]	2.9	7.3	11.8	2.4	5.7	9.3	20.7	17.9	11.6
LabelEmbed+ [25]	3.0	7.6	12.2	2.0	5.6	9.4	15.0	20.1	10.7
GenModel [26]	3.1	6.9	10.5	2.3	5.7	8.8	24.8	13.4	11.2
TMN [32]	3.5	8.1	12.4	2.9	7.1	11.5	20.2	20.1	13.0
SymNet (Ours)	4.3	9.8	14.8	3.0	7.6	12.3	24.4	25.2	16.1

Table 2. Results of generalized CZSL on MIT-States. All methods (Sec. 4.2) use ResNet-18 [10] as the backbone.

Method	MIT-States		UT-Zappos	
	Attribute	Object	Attribute	Object
AttrOperator [25]	14.6	20.5	29.7	67.5
GenModel [26]	15.1	27.7	18.4	68.1
SymNet	18.9	28.8	38.0	65.4

Table 3. Attribute learning results on two benchmarks.

[24, 25] (the scores with * are reproduced by [25], the others are from [24]). SymNet outperforms all baselines on two benchmarks. Although we use a simple product to compose the attribute and object scores, we still achieve 2.1% and 3.8% improvements over the state-of-the-art [26] on two benchmarks respectively. On UT-Zappos, most previous approaches do not surpass the Visual Product baseline, while ours outperforms it by 2.2%. To further evaluate our SymNet, we additionally conduct the comparison on generalized CZSL setting from recent state-of-the-art TMN [32]. The results are shown in Tab. 2. SymNet also outperforms previous methods significantly, which strongly proves the effectiveness of our method.

Attribute Learning. We also compare the attribute accuracy alone on two benchmarks in Tab. 3. We reproduce the results of AttrOperator [25] with its open-sourced code. For all methods involved, the individual attribute and object accuracy do not consider the relations between attributes and objects. The object recognition module of our method is a simple 3-layer MLP classifier with the visual image features from ResNet-18 backbone. SymNet outperforms previous methods by a large margin, *i.e.* 3.8% on MIT-States and 8.3% on UT-Zappos. Our RMD-based attribute recognition is particularly effective. In addition, our object classification performance is comparable to AttrOperator [25] and GenModel [26]. Accordingly, the main contribution of the CZSL improvement of SymNet comes from attribute learning rather than object recognition.

4.5. Image Retrieval after Attribute Manipulation

To qualitatively evaluate SymNet, we further report the image retrieval results after attribute manipulation. We first train SymNet on MIT-States or UT-Zappos, then use trained CoN and DeCoN to manipulate the image embeddings. For an image with pair label (a, o) , we remove the attribute a with DeCoN and add an attribute b with CoN, then we retrieve the top-5 nearest neighbors of the manipulated embeddings. This task is much more difficult than the normal attribute-object retrieval [25, 24, 32] because of the com-



Top-5 Attribute Manipulation Retrieval of UT-Zappos



Figure 5. Image Retrieval on MIT-States, UT-Zappos. We conduct the retrieval after the *attribute manipulation*.

plex semantic manipulation and recognition. Retrieval results are shown in Fig. 5, where the imaged on the left are original ones and right are the nearest neighbors after manipulation. SymNet is capable of retrieving a certain number of correct samples among top-5 nearest neighbors, especially in a fine-grained dataset like UT-Zappos. This suggests that our model has well exploited the learned symmetry in attribute transformation and learned the contextuality and compositionality of attributes.

4.6. Visualization in Latent Space

To verify the robustness and principles in transformations, we use t-SNE [21] to visualize the image embeddings before or after transformations in latent space in Fig. 6. Specifically, we first visualize the group axioms related transformations: 1) **Closure** is verified by comparing $\{f_o^i \cdot T_+(a^i) \cdot T_-(a^i)\}$ v.s. $\{f_o^i \cdot T_-(a^i)\}$ and $\{f_o^i \cdot T_+(a^j) \cdot T_+(a^j)\}$ v.s. $\{f_o^i \cdot T_+(a^j)\}$. 2) **Invertibility** is verified by comparing $\{f_o^i \cdot T_+(a^j) \cdot T_-(a^j)\}$ v.s. $\{f_o^i \cdot T_e\}$ and $\{f_o^i \cdot T_-(a^i) \cdot T_+(a^i)\}$ v.s. $\{f_o^i \cdot T_e\}$. 3) **Commutativity** is verified by comparing $\{f_o^i \cdot T_+(a^i) \cdot T_-(a^j)\}$ v.s. $\{f_o^i \cdot T_-(a^j) \cdot T_+(a^i)\}$. The results are shown in Fig. 6 (a,b). We observe that SymNet can robustly operate the transformations and the axiom objectives are well satisfied during embedding transformations.

Then, to verify the **symmetry** property, we visualize the sample embeddings in Relative Moving Space in Fig. 6(c,d): 1) for the sample f_o^i which do not have attribute a^j , $f_o^i \cdot T_+(a^j)$ should be far from f_o^i . On the contrary, $f_o^i \cdot T_-(a^j)$ are relatively close to f_o^i because of the symmetry. 2) For the sample f_o^i with attribute a^i , $f_o^i \cdot T_+(a^i)$ should be close to f_o^i and $f_o^i \cdot T_-(a^i)$ should be far from f_o^i . We can also find that the relative moving distance rules are all satisfied, *i.e.* the symmetry is well learned by our SymNet.

4.7. Ablation Study

To evaluate different components of our method, we design ablation studies and report the results in Tab. 4.

Objectives. To evaluate the objectives constructed from

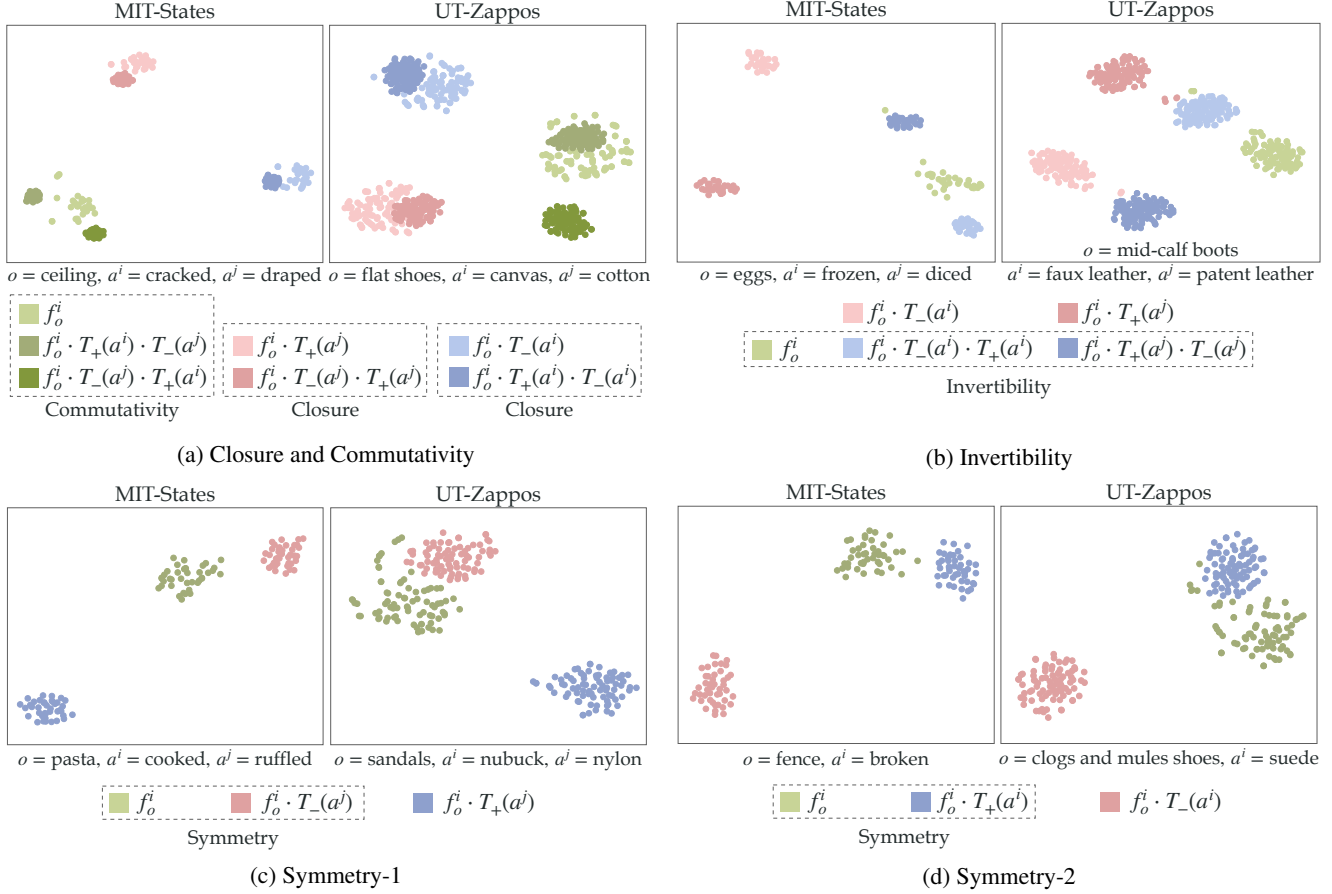


Figure 6. Visualization of symmetry and the group axioms by t-SNE [21]. The points with colors in a same dotted box should be close.

Method	MIT-States			UT-Zappos		
	Top-1	Top-2	Top-3	Top-1	Top-2	Top-3
SymNet	19.9	28.2	33.8	52.1	67.8	76.0
SymNet w/o \mathcal{L}_{sym}	18.3	27.5	33.4	51.1	67.0	76.0
SymNet w/o \mathcal{L}_{axiom}	16.9	25.5	30.9	47.6	65.4	73.6
SymNet w/o \mathcal{L}_{inv}	17.9	26.7	32.5	50.8	67.4	76.1
SymNet w/o \mathcal{L}_{com}	17.8	27.0	32.7	51.2	67.6	75.8
SymNet w/o \mathcal{L}_{clo}	18.0	27.0	32.8	51.1	67.2	76.0
SymNet w/o \mathcal{L}_{cls}	10.3	18.9	25.9	28.7	51.2	65.2
SymNet w/o \mathcal{L}_{tri}	17.8	26.8	32.6	49.2	65.3	74.2
SymNet w/o \mathcal{L}_{sym} & \mathcal{L}_{tri}	17.7	27.0	33.0	50.1	66.1	75.6
SymNet w/o \mathcal{L}_{tri} & \mathcal{L}_{cls}	10.5	19.4	26.7	28.6	51.4	65.6
SymNet w/o \mathcal{L}_{sym} & \mathcal{L}_{cls}	9.3	17.0	22.7	27.4	48.2	64.1
SymNet only \mathcal{L}_{sym}	9.4	16.9	22.5	20.4	38.9	53.5
SymNet w/o attention	18.0	26.9	32.7	48.5	65.0	75.6
SymNet L_1 dist.	7.1	11.2	14.3	37.5	53.3	62.3
SymNet Cos dist.	11.3	20.7	28.5	18.7	41.1	60.0

Table 4. Results of ablation studies.

group axioms and the core principle symmetry, we conduct tests of these objectives by removing them. In Tab. 4, SymNet shows obvious degradations without the constraints of these principles. This is in line with our assumption that a transformation framework that covers the essential principles can largely promote compositional learning.

Attention. Removing the attention module drops 1.9% and 3.6% accuracy on two benchmarks.

Distance Metrics. SymNet with other distance metrics, *i.e.*, L_1 and cosine distances, perform much worse than L_2 .

5. Conclusion

In this paper, we propose the symmetry property of attribute-object compositions. Symmetry reveals profound principles in composition transformations. To an object, giving it an attribute it already has, or erasing an attribute it does not have, would all result in the same object. To learn the symmetry, we construct a framework inspired by group theory to couple and decouple attribute-object compositions, and use group axioms and symmetry as the learning objectives. When applied to CZSL, our method achieves state-of-the-art performance. In the future, we consider to study the transformation with varying degrees, *e.g.*, not-peeled, half-peeled and totally-peeled and apply SymNet to GAN-related tasks.

Acknowledgement: This work is supported in part by the National Key R&D Program of China, No. 2017YFA0700800, National Natural Science Foundation of China under Grants 61772332 and Shanghai Qi Zhi Institute.

References

- [1] Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Describing people: A poselet-based approach to attribute classification. In *ICCV*, 2011. 2
- [2] Wei-Lun Chao, Soravit Changpinyo, Boqing Gong, and Fei Sha. An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *ECCV*, 2016. 6
- [3] Yu-Wei Chao, Yunfan Liu, Xieyang Liu, Huayi Zeng, and Jia Deng. Learning to detect human-object interactions. In *WACV*, 2018. 1
- [4] Chao-Yeh Chen and Kristen Grauman. Inferring analogous attributes. In *CVPR*, 2014. 2, 6
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 4
- [6] Yubin Deng, Ping Luo, Chen Change Loy, and Xiaoou Tang. Pedestrian attribute recognition at far distance. In *ACMMM*, 2014. 2
- [7] Hao-Shu Fang, Jianhua Sun, Runzhong Wang, Minghao Gou, Yong-Lu Li, and Cewu Lu. Instaboost: Boosting instance segmentation via probability map guided copy-pasting. In *ICCV*, 2019. 1
- [8] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *CVPR*, 2009. 1, 2
- [9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 1
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4, 6, 7
- [11] Sung Ju Hwang, Fei Sha, and Kristen Grauman. Sharing features between objects and their attributes. In *CVPR*, 2011. 2
- [12] Phillip Isola, Joseph J Lim, and Edward H Adelson. Discovering states and transformations in image collections. In *CVPR*, 2015. 1, 2, 3, 5, 6
- [13] Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C Berg, and Tamara L Berg. Babytalk: Understanding and generating simple image descriptions. *TPAMI*, 2013. 2
- [14] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 2009. 1, 2
- [15] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaoang Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, 2017. 1
- [16] Yong-Lu Li, Liang Xu, Xijie Huang, Xinpeng Liu, Ze Ma, Mingyang Chen, Shiyi Wang, Hao-Shu Fang, and Cewu Lu. Hake: Human activity knowledge engine. *arXiv preprint arXiv:1904.06539*, 2019. 1
- [17] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, and Cewu Lu. Transferable interactiveness knowledge for human-object interaction detection. In *CVPR*, 2019. 1
- [18] Yutian Lin, Liang Zheng, Zhedong Zheng, Yu Wu, Zhi-lan Hu, Chenggang Yan, and Yi Yang. Improving person re-identification by attribute and identity learning. *Pattern Recognition*, 2019. 2
- [19] Ziwei Liu, Ping Luo, Xiaoang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 2
- [20] Cewu Lu, Hao Su, Yonglu Li, Yongyi Lu, Li Yi, Chi-Keung Tang, and Leonidas J Guibas. Beyond holistic object recognition: Enriching image understanding with part states. In *CVPR*, 2018. 1
- [21] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 2008. 7, 8
- [22] Dhruv Mahajan, Sundararajan Sellamanickam, and Vinod Nair. A joint learning framework for attribute models and object descriptions. In *ICCV*, 2011. 2
- [23] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013. 2, 6
- [24] Ishan Misra, Abhinav Gupta, and Martial Hebert. From red wine to red tomato: Composition with context. In *CVPR*, 2017. 1, 2, 5, 6, 7
- [25] Tushar Nagarajan and Kristen Grauman. Attributes as operators: factorizing unseen attribute-object compositions. In *ECCV*, 2018. 1, 2, 5, 6, 7
- [26] Zhixiong Nan, Yang Liu, Nanning Zheng, and Song-Chun Zhu. Recognizing unseen attribute-object pair with generative model. In *AAAI*, 2019. 2, 6, 7
- [27] Vicente Ordonez, Girish Kulkarni, and Tamara L Berg. Im2text: Describing images using 1 million captioned photographs. In *NIPS*, 2011. 2
- [28] Devi Parikh and Kristen Grauman. Relative attributes. In *ICCV*, 2011. 1, 2
- [29] Genevieve Patterson and James Hays. Coco attributes: Attributes for people, animals, and objects. In *ECCV*, 2016. 1, 2
- [30] Peixi Peng, Yonghong Tian, Tao Xiang, Yaowei Wang, Massimiliano Pontil, and Tiejun Huang. Joint semantic and latent attribute modelling for cross-class transfer learning. *TPAMI*, 2017. 2
- [31] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 2, 4, 6
- [32] Senthil Purushwalkam, Maximilian Nickel, Abhinav Gupta, and Marc’Aurelio Ranzato. Task-driven modular networks for zero-shot compositional learning. In *ICCV*, 2019. 6, 7
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 1
- [34] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 2
- [35] Wenguan Wang, Xiankai Lu, Jianbing Shen, David J Crandall, and Ling Shao. Zero-shot video object segmentation via attentive graph neural networks. In *ICCV*, 2019. 1
- [36] Xin Wang, Fisher Yu, Ruth Wang, Trevor Darrell, and Joseph E Gonzalez. Tafe-net: Task-aware feature embeddings for low shot learning. In *CVPR*, 2019. 2, 6

- [37] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010. [1](#), [2](#)
- [38] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning-a comprehensive evaluation of the good, the bad and the ugly. *TPAMI*, 2018. [1](#)
- [39] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010. [1](#), [2](#)
- [40] Wenqiang Xu, Yonglu Li, and Cewu Lu. Srda: Generating instance segmentation annotation via scanning, reasoning and domain adaptation. In *ECCV*, 2018. [1](#)
- [41] Aron Yu and Kristen Grauman. Semantic jitter: Dense supervision for visual comparisons via synthetic images. In *ICCV*, 2017. [1](#), [2](#), [3](#), [5](#), [6](#)