

Siamese Contrastive Embedding Network for Compositional Zero-Shot Learning

Xiangyu Li, Xu Yang*, Kun Wei, Cheng Deng*, Muli Yang

School of Electronic Engineering, Xidian University, Xi'an 710071, China

xdu_xyLi@stu.xidian.edu.cn, {xuyang.xd, weikunsk, chdeng.xd, muliyang.xd}@gmail.com

Abstract

Compositional Zero-Shot Learning (CZSL) aims to recognize unseen compositions formed from seen state and object during training. Since the same state may be various in the visual appearance while entangled with different objects, CZSL is still a challenging task. Some methods recognize state and object with two trained classifiers, ignoring the impact of the interaction between object and state; the other methods try to learn the joint representation of the state-object compositions, leading to the domain gap between seen and unseen composition sets. In this paper, we propose a novel Siamese Contrastive Embedding Network (SCEN)¹ for unseen composition recognition. Considering the entanglement between state and object, we embed the visual feature into a Siamese Contrastive Space to capture prototypes of them separately, alleviating the interaction between state and object. In addition, we design a State Transition Module (STM) to increase the diversity of training compositions, improving the robustness of the recognition model. Extensive experiments indicate that our method significantly outperforms the state-of-the-art approaches on three challenging benchmark datasets, including the recent proposed C-QGA dataset.

1. Introduction

Humans possess the ability to compose their knowledge of known entities to generalize to novel concepts inherently. Given words, such as *green horse*, people can combine the known state *green* with the known object *horse* immediately, although they have never seen the inexistent stuff. To equip an AI system the similar ability, Compositional Zero-Shot Learning (CZSL) [20] is proposed, which aims to recognize unseen compositions composed of a set of seen states and objects. In CZSL setting, each composition comprises two components, namely, state and object, where the compositions of train and test sets are disjoint.

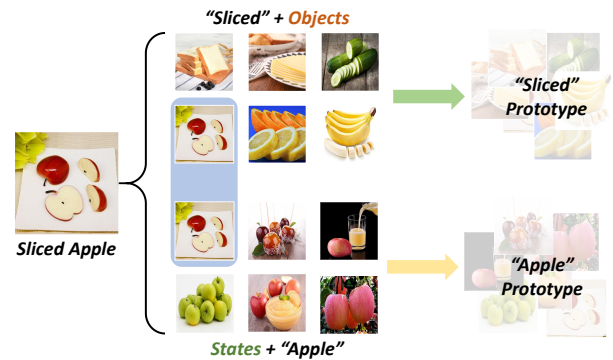


Figure 1. The overall concept of our method. We aim to separately extract discriminative prototypes of state and object based on establishing state and object specific databases, which can generalize to represent corresponding properties.

In order to infer unknown concepts such as *green horse*, CZSL aims to understand the meaning of state *green* and object *horse* after trained on other compositional concepts that separately contain *green* or *horse*, e.g., *green grasses* and *young horse*. The challenge of the task lies in the interaction degree between state and object that we cannot quantify, which gives rise to varying contextuality within different state-object combinations. For instance, we can not equate the state *old* in *old car* and that in *old tiger*, since they are fundamentally distinct in visual presentations, which greatly hinders the recognition of novel compositions.

Existing mainstream methods [17, 20, 26] focus on converting such problem into a general supervised recognition task by training two classifiers for state and object, respectively. They aim to directly predict state and object from the original visual features, ignoring their entanglement. Based on this problem, classifiers cannot capture discriminative state and object features, which potentially limits the recognition accuracy. In addition, other methods [23, 24] aim to learn a common embedding space where the compositions as well as visual features can be projected to narrow the distance between them, such as Euclidean distance. However, these methods, only regarding compositions as enti-

*Corresponding author.

¹Code: <https://github.com/XDUxyLi/SCEN-master>

ties, neglect the domain gap between training and testing compositions, which can be simply confused by similar images from unseen compositions (e.g., *young cat* and *young tiger*). Therefore, it is vital to excavate the discriminative prototypes of state and object to separate the interaction between them and consider the domain transfer between training and testing samples.

To address the problem mentioned above, we propose a Siamese Contrastive Embedding Network (SCEN) for recognizing novel compositions in this paper, aiming to excavate discriminative prototypes of state and object, respectively, as shown in Fig. 1. To be specific, we first project the visual features into state/object-based contrastive spaces to gain the prototypes of state and object. Then, to excavate the discriminative prototypes by contrastive constraints, we set up specific databases named *State-constant* and *Object-constant* databases as positive samples. Besides, a shared irrelevant database is built up as a negative sample set, which is embedded into two contrastive spaces. Benefiting from this learning paradigm, our proposed model can successfully excavate discriminative prototypes to represent the corresponding component. In addition, considering that the distribution between seen and unseen compositions is discrepant, we present a State Transition Module (STM), which generates the virtual but reasonable samples to augment the diversity of training data. In this way, the domain gap between seen and unseen composition sets can be mitigated effectively.

To sum up, our main contributions are as follows:

- We propose a novel Siamese Contrastive Embedding Network (SCEN) to excavate prototypes of state and object for successfully recognizing both seen and unseen compositions.
- We present a State Transition Module (STM) to produce virtual samples and augment the diversity of training compositions, guiding the model to generalize to those compositions not existing in the training process, and alleviating the issue of model migration performance.
- Comprehensive experimental results on three benchmark datasets demonstrate the effectiveness of our proposed approach, which outperforms the state-of-the-art CZSL methods.

2. Related Work

Compositional Zero-Shot Learning. The goal of Compositional Zero-Shot Learning [7, 17, 19, 20, 23, 26, 30] is to learn the compositionality of objects and their states from the training data and is tasked with the generalization to an unseen combination of these primitives. Compared

with typical Zero-Shot Learning [15, 16, 29] that utilizes inherent semantic descriptions or attributed vectors to recognize unseen instances, CZSL exploits transferable knowledge by two compositional parts as image labels: objects and states. There are two mainstream methods in this direction. The first mainstream approach is inspired by [2, 11], which learns a single classifier for recognition and a transformation module [20]. In addition, [23] models each state as a linear transformation of objects. [32] aims to learn disentangled and compositional primitives hierarchically. [17] models objects to be symmetric under attribute transformations. Other methods try to learn the joint representation of the state-object compositions [1, 26, 28]. They aim to learn a modular network to rewire the new compositions conditioned on each composition [26, 28]. Recently, GCN [22] is proposed to utilize a causal graph to establish the relationship between state and object reasonably. However, these methods ignore the interaction between state and objects that brings a negative influence for compositions recognition.

As for [1], the author argues to tackle the CZSL problems through a causal graph where the latent features of primitives are independent from each other. However, it also neglects the discriminant analysis of state and object, which cannot excavate discriminative primitives for classification. In addition, there still exists a domain gap between seen and unseen compositions, although they are made up of the same states and objects, which potentially limits the performance of the model.

Contrastive Learning. Inspired by noise contrastive estimation [8, 21, 34], contrastive learning has attracted much attention which leads to major advances in self-supervised representation learning. An efficient way to get better contrastive learning is to employ large numbers of negative examples and design more semantically meaningful augmentations to create different view of images. SimCLR [4] implements two data augmentation paths and a learnable non-linear transformation to train an encoder with a large batch by pulling the features embedding from the same images. Momentum Contrast (MoCo) [9] is presented which enables building a large and consistent dictionary on-the-fly and transfers well to downstream tasks. Aiming at improving generalization in real domains, a contrastive synthetic-to-real generalization model [5] is proposed to prevent overfitting to the synthetic domain by leveraging the pre-trained ImageNet knowledge. More recently, supervised contrastive learning [13] is proposed to extend the self-supervised batch contrastive approach to the fully supervised setting which can effectively leverage label information.

Based on the effectiveness of Contrastive Learning, the differences between this study and existing works are given below. First, we propose a Siamese Contrastive Embedding

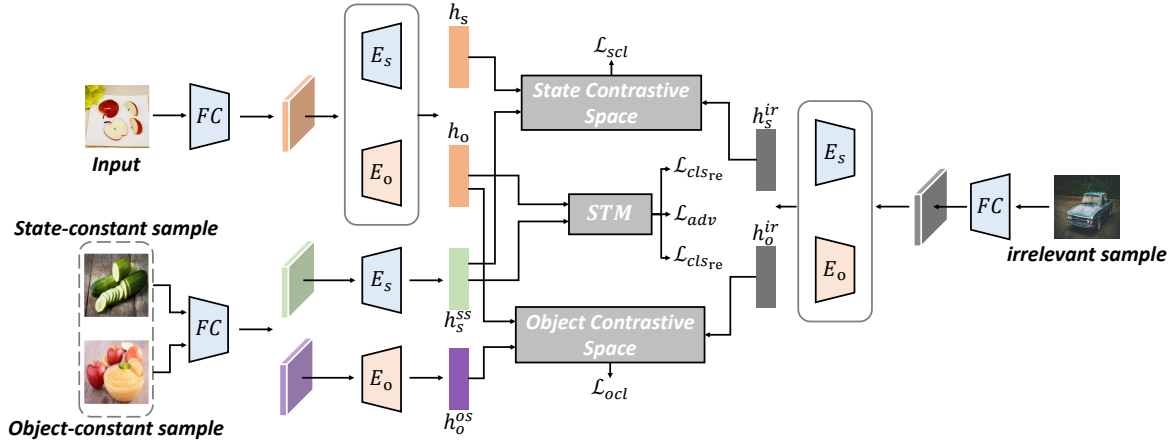


Figure 2. The framework of our SCEN method. The proposed method consists of State-Specific Encoder E_s , Object-Specific Encoder E_o , and a State Transition Module (STM). The prototypes of state h_s and object h_o are separately encoded from Specific Encoders. Prototypes are trained conditioned on two different contrastive embedding spaces, where the anchor and negative samples are shared. In addition, STM aims to generate virtual samples with the help of the adversarial loss function, which diversifies and rationalizes the states of objects effectively. Along with the increasing of the realistic samples, Specific Encoders gradually gain the ability to extract discriminative prototypes that can be generalized to novel compositions.

Network (SCEN) to excavate the discriminative prototypes of state and object, respectively. Besides, we present a State Transition Module (STM) to produce a virtual composition in training to improve the generalization of the proposed model. Second, we construct two contrastive spaces and utilize contrastive constraints to enforce the prototypes of state and object to be discriminative and generalized.

3. Approach

The goal of CZSL is to recognize the novel compositional samples whose labels are composed of a state (e.g., *old*) and an object (e.g., *tiger*). This is particularly challenging since various states significantly change the visual appearance of an object, which hinders the performance of the classifiers.

We propose a novel formulation to tackle the problem, namely Siamese Contrastive Embedding Network (SCEN), which constructs two independent embedding spaces and utilizes contrastive losses to guide corresponding feature extractors to excavate discriminative prototypes of state and object separately. The overview of our approach is shown in Fig. 3.

3.1. Problem Definition

In CZSL setting, each image consists of two primitive concepts. i.e., a state and an object. Given \mathcal{A} and \mathcal{O} as two sets of states and objects, we can compose a set of state-object pairs, i.e., $\mathcal{C} = \mathcal{A} \times \mathcal{O} = \{(a, o) \mid a \in \mathcal{A}, o \in \mathcal{O}\}$. Besides, we denote training set as $\mathcal{D}_{tr} = \{(i, c) \mid i \in \mathcal{I}^s, c \in \mathcal{C}^s\}$, where \mathcal{I}^s is the image set known in training, and \mathcal{C}^s is a subset of \mathcal{C} containing the corresponding

labels. In the traditional Zero-Shot Learning setting, training and testing label are disjoint, i.e., $\mathcal{C}^s \cap \mathcal{C}^u = \emptyset$, where $\mathcal{C}^s, \mathcal{C}^u$ are two subsets of \mathcal{C} seen/unseen in training. In this case, the model only needs to predict the compositions drawn from \mathcal{C}^u in testing [20]. In this paper, we follow the setting of Generalized ZSL [31] where testing samples can be drawn from either seen or unseen compositions, i.e., $\mathcal{C}^s \cup \mathcal{C}^u$, which is more challenging on account of the larger prediction space and the dominant bias to seen compositions [26]. To sum up, CZSL aims to learn a mapping function $\mathcal{I} \rightarrow \mathcal{C}^s \cup \mathcal{C}^u$ that is trained on $\{\mathcal{I}^s, \mathcal{C}^s\}$, in which \mathcal{C} is composed of two primitive concepts drawn from \mathcal{A} and \mathcal{O} .

3.2. Siamese Contrastive Embedding Network

Due to the entanglement between state and object into an image that influences the final classification, we design a Siamese Contrastive Embedding Network (SCEN) to better materialize the discriminative prototypes of state and object, which can effectively improve the accuracy of recognition models. The overall architecture is illustrated in Fig. 2. The SCEN is composed of a feature extractor FC , a State-Specific Encoder E_s , a Object-Specific Encoder E_o , and a State Transition Module (STM).

Specific database. Let us consider a training sample, such as *sliced apple* in Fig. 1. As we all know, from our training set, that object *apple* comes in various states such as *caramelized* and state *sliced* also does not just modify a single object, e.g., *sliced banana*. Therefore, these sample points with overlapping information might have potential relationships. Based on this idea, we set up three specific databases D_s, D_o , and D_{ir} to excavate discriminative state

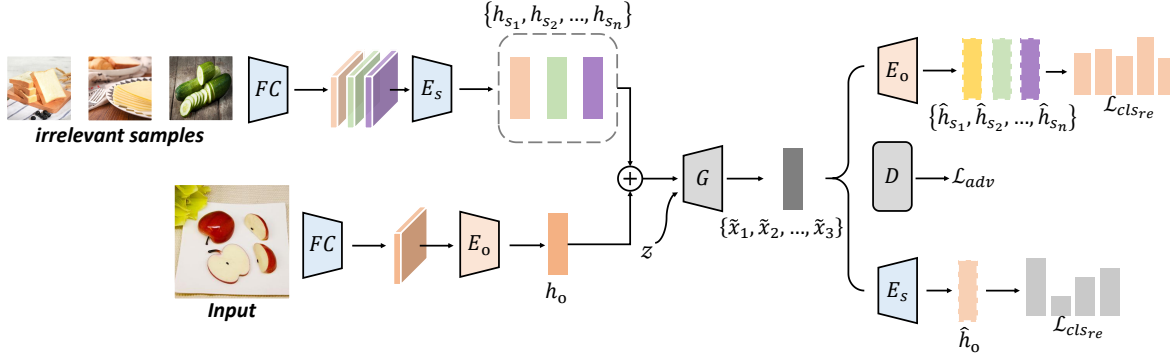


Figure 3. The framework of our State Transition Module (STM). The proposed module aims to improve the generalization of State-Specific Encoder E_s and Object-Specific Encoder E_o . To be specific, the Generator G takes the state h_s and object h_o vectors as input and generates virtual samples, which is the input of Discriminator D . Although we do not gain the label for generated samples, we can utilize E_s and E_o to re-encode them and set classification losses as constraints.

and object factors, respectively. D_s is the set of compositions consisting of constant state and various objects, named State-constant database, while the Object-constant database D_o is defined as the set of compositions made up of various states and a constant object. In addition, D_{ir} is the set of compositions formed from various states and objects, which are both different from the state and object of input instances. For instance, given an image as input which consists of a state \hat{a} and an object \hat{o} , i.e., $x = (\hat{a}, \hat{o}) \in \mathcal{I}^s$, the D_o is denoted as:

$$\mathcal{D}_o = \{(a, o) \mid o = \hat{o}, (a, o) \in \mathcal{C}^s\}. \quad (1)$$

Analogously, the State-specific database D_s is denoted as:

$$\mathcal{D}_s = \{(a, o) \mid a = \hat{a}, (a, o) \in \mathcal{C}^s\}. \quad (2)$$

And the irrelevant database D_{ir} is denoted as:

$$\mathcal{D}_{ir} = \{(a, o) \mid a \neq \hat{a}, o \neq \hat{o}, (a, o) \in \mathcal{C}^s\}. \quad (3)$$

Siamese Contrastive Space. Based on sets of specific databases being set up, the visual feature \mathbf{x} , extracted by the feature extractor FC , are separately projected into two independent contrastive embedding spaces to extract prototypes of state h_s and object h_o :

$$\begin{aligned} h_s &= E_s(\mathbf{x}), \\ h_o &= E_o(\mathbf{x}). \end{aligned} \quad (4)$$

We hope that h_s and h_o contain information that is separately sensitive to classifiers for compositions recognition. Therefore, we aim to utilize contrastive learning as a constraint condition to extract discriminative prototypes of state and object. However, a general contrastive loss simultaneously cannot extract their discriminative representations due to their interaction. Based on this problem, we define state-based contrastive loss \mathcal{L}_{scl} and object-based contrastive loss

\mathcal{L}_{ocl} as constraints to enforce the model to extract discriminative primitives.

To be specific, we set h_s as an anchor in the state-based contrastive space. Meanwhile, h_s^{ss} selected from D_s is set as a positive point while k negative samples selected from the D_{ir} are denoted as $\{h_{s_1}^{ir}, \dots, h_{s_k}^{ir}\}$. We aim to decrease the distance between the anchor h_s and the positive instance h_s^{ss} , increasing the distance between h_s and each negative point $h_{s_i}^{ir}$ to extract discriminative prototype of state. Therefore, the state-specific loss function \mathcal{L}_{scl} in the state-based contrastive space can be calculated as follows:

$$\mathcal{L}_{scl} = -\log \frac{\exp((h_s)^\top h_s^{ss} / \tau_s)}{\exp((h_s)^\top h_s^{ss} / \tau_s) + \sum_{i=1}^K \exp((h_s)^\top h_{s_i}^{ir} / \tau_s)}, \quad (5)$$

where $\tau_s > 0$ is the temperature parameter for the contrastive embedding and K is the number of negative samples. It is obvious that the larger K we set, the longer time the training process cost. The larger number of negative samples encourages State-Specific Encoders E_s to excavate a more representative state prototype, which can be generalized to novel compositions.

As a Siamese Contrastive Embedding space, similar to the state-based contrastive space, we denote h_o embedded from the same input visual features as an anchor in the object-based contrastive space, and $\{h_{o_1}^{ir}, \dots, h_{o_k}^{ir}\}$ as negative points embedded from irrelevant database D_{ir} . Therefore, the object-specific loss function \mathcal{L}_{ocl} can be defined in the object contrastive space as:

$$\mathcal{L}_{ocl} = -\log \frac{\exp((h_o)^\top h_o^{os} / \tau_o)}{\exp((h_o)^\top h_o^{os} / \tau_o) + \sum_{j=1}^K \exp((h_o)^\top h_{o_j}^{ir} / \tau_o)}, \quad (6)$$

where $\tau_o > 0$ is the temperature parameter for the con-

trastive embedding and K is the number of negative samples.

Considering that the prototypes of state and object should be optimized in the same direction, we share irrelevant samples for two contrastive spaces $\{h_{o_1}^{ir}, \dots, h_{o_k}^{ir}\}$ as negative points, which effectively avoids the problem of unbalanced optimization between E_s and E_o .

Finally, we introduce classification losses to guide classifiers to recognize prototypes of state and object, respectively, which is formulated as:

$$\mathcal{L}_{cls} = C_a(h_s, a) + C_o(h_o, o), \quad (7)$$

where C_a and C_o are both fully connected layers with the cross-entropy loss trained to classify state and object respectively. The prototypes of state and object can be further preserved in the composition with the supervision of classification losses.

Thus, the total loss function in Siamese Contrastive Space \mathcal{L}_{cts} can be formulated as:

$$\mathcal{L}_{cts} = \mathcal{L}_{scl} + \mathcal{L}_{ocl} + \mathcal{L}_{cls}. \quad (8)$$

State Transition Module. In order to enforce the SCEN to be generalized to novel samples that do not appear in the training stage, we aim to produce virtual samples to augment the diversity of training compositions, alleviating the domain gap between training and testing data. Therefore, we propose a State Transition Module (STM), which consists of a State-Specific encoder E_s , a Object-Specific encoder E_o , a Generator G and a Discriminator D [6]. The architecture is shown in Fig. 3.

Let us consider two objects, namely *apple* and *banana*. As is known to all, from our training set, that *apple* can be *ripe* while *banana* can be *caramelized* since there appears at least one image in training. However, there exists *caramelized apple* and *ripe banana* compositions in the testing set while the training set does not. Therefore, we can conclude that the object has the possibility of forming a new combination with various states. Based on this discovery, we aim to utilize the Generator G to produce virtual compositions with the input of various states and a given object. However, such generation with the random combination can produce many irrational compositions, which will actually widen the domain gap between seen and unseen data. For instance, *cored banana* and *squished apple* do not appear in the testing set or even exist in reality. Thus, we design a Discriminator D to distinguish whether a composition is composed by the generator G .

In particular, the Generator G takes as input the prototype of an object h_o and another state \tilde{h}_s to generate virtual compositions that never appear in training. Then, the Discriminator D takes the real samples $x_{a,o}$ as input and determines which are produced by the Generator G . G and D

can be optimized by the following adversarial objective:

$$\max_D \min_{G, E_s, E_o} V(G, D) = \mathbb{E}_{s,o} (\log D(x_{a,o})) + \mathbb{E}_{h_{\tilde{s}}, h_o} (\log(1 - D(G(h_{\tilde{s}}, h_o))))), \quad (9)$$

where $G(h_{\tilde{s}}, h_o) = \hat{x}_{\tilde{s},o}$. G tries to minimize $V(G, D)$ while D tries to maximize it.

The goal of improving the E_s and E_o performance is to be generalized to novel compositions in testing, but the generated samples do not have labels as supervision. Thus, we re-encode the generated samples to extract prototypes of state and object again, and design a re-classification loss to constrain them, which is formulated as follows:

$$\mathcal{L}_{cls_{re}} = C_a(E_s(G(h_{\tilde{s}}, h_o)), \tilde{a}) + C_o(E_o(G(h_{\tilde{s}}, h_o)), o). \quad (10)$$

The total loss function of State Transition Module \mathcal{L}_{stm} is formulated as follows:

$$\mathcal{L}_{stm} = \max_D \min_{G, E_s, E_o} V(G, D) + \mathcal{L}_{cls_{re}}. \quad (11)$$

Eventually, the final loss of our proposed framework is formulated as:

$$\mathcal{L}_{total} = \alpha \mathcal{L}_{cts} + \beta \mathcal{L}_{stm}, \quad (12)$$

where α and β are the weighting coefficients to balance the influence of each loss function, respectively.

3.3. Inference

In the training stage, the model is trained to estimate the likelihood $p(I = i | A = a, O = o)$ for image i conditioned on state s and object o . The inference takes place in both \mathcal{C}^s and \mathcal{C}^o . In the inference, model embeds an image as \mathbf{x} and extract prototypes of state and object, i.e., h_a and h_s with trained E_s and E_o , respectively. Then the state and object of the most similar prototypes are taken as the prediction. The inference rule can be parameterized as:

$$c_t = (a_t, o_t) = \arg \max_{(a,o) \in \mathcal{C}} p(i | E_s(\mathbf{x}), E_o(\mathbf{x})). \quad (13)$$

4. Experiment

In this section, all datasets and evaluation protocols are introduced concretely. Then, we present the implementation details and the comparison of experimental results with other state-of-the-art methods. Eventually, ablation studies prove the effectiveness of the method we proposed.

4.1. Experimental Setup

Datasets. Our proposed method is evaluated on three CZSL benchmark datasets, i.e., MIT-States [12], UT-Zappos [33], and C-GQA [22].

MIT-States contains 53753 images, e.g., *young cat* and *rusty bike*, with 115 states and 245 objects in total. MIT-States has 1962 available compositions where 1262 state-object pairs are seen in the training stage, leaving 700 pairs unseen. UT-Zappos contains 50025 images of shoes, e.g., *Cotton Sandals* and *Suede Slippers*, with 16 states and 12 objects. In UT-Zappos, there are 116 state-object pairs, 83 pairs of which are used for training, while the other 33 pairs are unseen in training. As for C-GQA dataset, it contains over 9500 compositions that make it most extensive dataset for CZSL. The detailed information of each dataset is summarized in Tab. 1.

Evaluation Metrics. We evaluate the performance according to prediction accuracy for recognizing seen and unseen compositions. Following the setting of [26], we compute the accuracy in two situations: 1) *Seen*, testing only on seen compositions; 2) *Unseen*, testing only on unseen compositions. Based on these, we can compute *Harmonic Mean HM* of the two metrics, which balances the performance between seen and unseen accuracies. Eventually, we compute 4) *Area Under the Curve (AUC)* to quantify the overall performance of both seen and unseen accuracy at different operating points with respect to the bias. Following [3, 26], we utilize a calibration bias to trade off between the prediction scores of seen and unseen pairs. As the calibration bias varies, we can draw a seen-unseen accuracy curve where the AUC metric can be computed.

Implementation Details. For each image, we extract a 1024 dimensional visual feature vector using ResNet-18 [10] pre-trained on the ImageNet dataset [27]. We separately extract a 300-dimensional feature vectors for both states and objects with E_s and E_o , which is implemented with two fully-connected layers and ReLU activation. Our model is implemented with PyTorch [25] and optimized by ADAM optimizer [14] on an NVIDIA GTX 1080Ti GPU. In addition, we set the learning rate as 0.00004, batch size as 128, and the number of negative samples K as 10. For the MIT-States dataset, the training time is approximately 3 hours for 800 epochs. For the UT-Zappos dataset, it takes around 1 hour for 500 epochs in training. As for C-GQA, it spends around 4 hours for 1000 epochs in training.

4.2. Comparison with State-of-the-Arts

We compare our experiments with the state-of-the-art in Tab. 2 and show that our Siamese Contrastive Embedding Network (SCEN) outperforms all previous methods in three benchmark datasets, which includes recent proposed C-GQA dataset [22]. Our detailed observations are as follows.

Generalized CZSL performance. For the CZSL task, our SCEN achieves a test AUC of 5.3%, which achieves the best result on MIT-States. In addition, our method significantly boosts the state-of-the-art harmonic mean, i.e.,

17.2% to 18.4%. When it refers to state and object prediction accuracy, we can observe an improvement from 27.9% to 28.2 % for states and 31.8% to 32.2% for objects.

Similar observations are confirmed on UT-Zappos, in which we can achieve a superior improvement on state-of-the-arts with an AUC of 32.0% compared to 28.7% from Compos. In addition, our proposed model performs the best harmonic mean 47.8% and improves around 4.5% compared with the Compos.

Finally, on the recent proposed splits of the C-GQA dataset, which is shown in Tab. 3, we also achieve the best test AUC of 5.5%. Since C-GQA is a large number of compositions (over 9.3k concepts), which is more complex than MIT-States and UT-Zappos for recognition. The state and the object accuracies of our method are 28.1% and 32.8%, which are both higher than state-of-the-arts. In addition, our best seen and unseen accuracies (28.9% and 25.4%) also achieve the best results on this new dataset.

According to the significant improvement on three challenging datasets, we can conclude that our proposed Siamese Contrastive Embedding Network (SCEN) can not only effectively extract discriminative prototypes of state and object, but also improve the robust of the model for unseen compositions recognition.

4.3. Ablation Study

We now make an ablation study to evaluate the effectiveness of the Siamese Contrastive Embedding Network. We take a single classification model as a base model, which trains two classifiers to recognize states and objects separately. Meanwhile, we train three variants by adding Siamese Contrastive space, adding State Transition Module (STM), or adding both of them, which is denoted as *base model*, *base model* + \mathcal{L}_{cts} , and *base model* + \mathcal{L}_{cts} + \mathcal{L}_{stm} , respectively. According to the showing experiment of each setting as shown in Tab. 4, every variant tends to be more superior performance than the *base model*. The combination of two variants we proposed achieves the best improvement, which demonstrates that different components promote each other and work together to improve the performance of SCEN significantly. In addition, the result of the *base model* with \mathcal{L}_{cts} proves that our proposed model successfully excavates discriminative prototypes of states and objects, which is better for compositions recognition. Meanwhile, the significant improvement after adding \mathcal{L}_{stm} indicates the necessity of the proposed State Transition Module, which can effectively alleviate the domain gap between training and testing data. Finally, the result of the *base model* with adding \mathcal{L}_{cts} and \mathcal{L}_{stm} achieves the best, which shows that they improve the performance of the model together, not affected by each other.

Dataset	s	o	Training		Validation			Test		
			c_s	i	c_s	c_u	i	c_s	c_u	i
MIT-States [12]	115	245	1262	30k	300	300	10k	400	400	13k
UT-Zappos [33]	16	12	83	23k	15	15	3k	18	18	3000
C-GQA [22]	453	870	6963	26k	1173	1368	7k	1022	1047	5k

Table 1. Datasets used in our experiments and their statistics. We use three datasets to benchmark our method against the baseline states. states s , objects o , seen compositions c_s , unseen compositions c_u , images i .

Method	MIT-States							UT-Zappos						
	AUC		Best					AUC		Best				
	Val	Test	HM	Seen	Unseen	s	o	Val	Test	HM	Seen	Unseen	s	o
AttOp [23]	2.5	1.6	9.9	14.3	17.4	21.1	23.6	21.5	25.9	40.8	59.8	54.2	38.9	69.6
LE+ [20]	3.0	2.0	10.7	15.0	20.1	23.5	26.3	26.4	25.7	41.0	53.0	61.9	41.2	69.2
TMN [26]	3.5	2.9	13.0	20.2	20.1	23.3	26.5	36.8	29.3	45.0	58.7	60.0	40.8	69.9
SymNet [17]	4.3	3.0	16.1	24.4	25.2	26.3	28.3	25.9	23.9	39.2	53.3	57.9	40.5	71.2
CGE [22]	6.8	5.1	17.2	28.7	25.3	27.9	32.0	38.7	26.4	41.2	56.8	63.6	45.0	73.9
CompCos [18]	5.9	4.5	16.4	25.3	24.6	27.9	31.8	38.6	28.7	43.1	59.8	62.5	44.7	73.5
Ours	7.2	5.3	18.4	29.9	25.2	28.2	32.2	40.2	32.0	47.8	63.5	63.1	47.3	75.6

Table 2. The state-of-the-art comparisons on UT-Zappos and MIT-States. U and S are the accuracies tested on unseen classes and seen classes in CZSL, respectively. H is the harmonic mean value of U and S . The best results are marked in bold.

Method	C-GQA						
	AUC		Best				
	Val	Test	HM	Seen	Unseen	s	o
AttOp [23]	2.5	1.6	9.9	14.3	17.4	21.1	23.6
LE+ [20]	3.0	2.0	10.7	15.0	20.1	23.5	26.3
TMN [26]	3.5	2.9	13.0	20.2	20.1	23.3	26.5
SymNet [17]	4.3	3.0	16.1	24.4	25.2	26.3	28.3
CGE [22]	6.8	5.1	17.2	28.7	25.3	27.9	32.0
CompCos [18]	5.9	4.5	16.4	25.3	24.6	27.9	31.8
Ours	6.9	5.5	17.5	28.9	25.4	28.1	32.8

Table 3. The state-of-the-art comparisons on recent proposed C-GQA dataset.

4.4. Qualitative Results

We show some qualitative results for the novel compositions with top-3 predictions in Fig. 4. The first three columns present some examples where the top prediction matches the label. For MIT-States and UT-Zappos datasets, we notice that the remaining two answers of the model can fundamentally capture at least one factor, which proves that the superior performance of our method. As for more complex C-GQA dataset, our model can give the correct answer in top-3 predictions, which shows the robustness of our proposed framework.

Meanwhile, the model can predict more combinations of unseen compositions, rather than being limited to that of seen compositions, which effectively alleviate the domain gap between seen and unseen samples.

In addition, the last two columns show the wrong prediction. For instance, in column 4 and row 2, the image of

the *Slippers* is misclassified as *Sandals* or *Boost*. This is because there exists a large number of training compositions so that the negative sample dataset may not contain the entire negative samples, such as *Sandals* and *Boost*, thus the model does not pay more attention to these pairs that are not included in D_{ir} . Besides, limited by the compositional class accuracy dependent on the number of groups associated with an object in the label space, the model may only focus on the state of the object in a certain aspect and ignore the state of the object labeled by the tag. For example, in column 4, row 1 the image of the *cat* consists of texture and age both present in the label space of the dataset and the output of the model. However the label for this image only contains its age.

4.5. Hyper-Parameter Analysis.

We perform an experiment to demonstrate the effect of the weighting co-efficients α and β for the loss functions \mathcal{L}_{cts} and \mathcal{L}_{stm} in our proposed model. As is shown in Fig. 5 and Fig. 6. With the different α and β setting, the *Harmonic Mean HM* and *AUC* have a certain degree of change, which indicates that \mathcal{L}_{cts} and \mathcal{L}_{stm} dominate the performance of the entire model. Based on this situation, we set and fix $\alpha = 0.1$, changing the value of β to observe the performance on different datasets. Finally, on MIT-States, UT-Zappos, and C-GQA datasets, we set $\beta = 0.5$, $\beta = 0.1$, and $\beta = 0.1$, which can achieve the best results, respectively.

Method	MIT-States							UT-Zappos						
	AUC		Best					AUC		Best				
	Val	Test	HM	Seen	Unseen	s	o	Val	Test	HM	Seen	Unseen	s	o
Base	2.5	1.6	9.9	14.3	17.4	21.1	23.6	21.5	25.9	40.8	59.8	54.2	38.9	69.6
+ \mathcal{L}_{cts}	5.8	4.3	16.5	25.3	24.7	27.9	31.6	32.6	26.4	44.2	60.0	61.2	42.5	71.3
+ \mathcal{L}_{stm}	4.8	3.9	13.3	22.5	21.9	27.6	30.2	34.6	28.5	43.9	59.8	61.1	46.9	72.1
$\mathcal{L}_{cts}+\mathcal{L}_{stm}$	7.2	5.3	18.4	29.2	25.2	28.2	32.2	39.0	29.7	47.8	63.3	62.5	47.3	74.4

Table 4. Ablation studies for Compositional Zero-Shot Learning on MIT-States and UT-Zappos datasets.

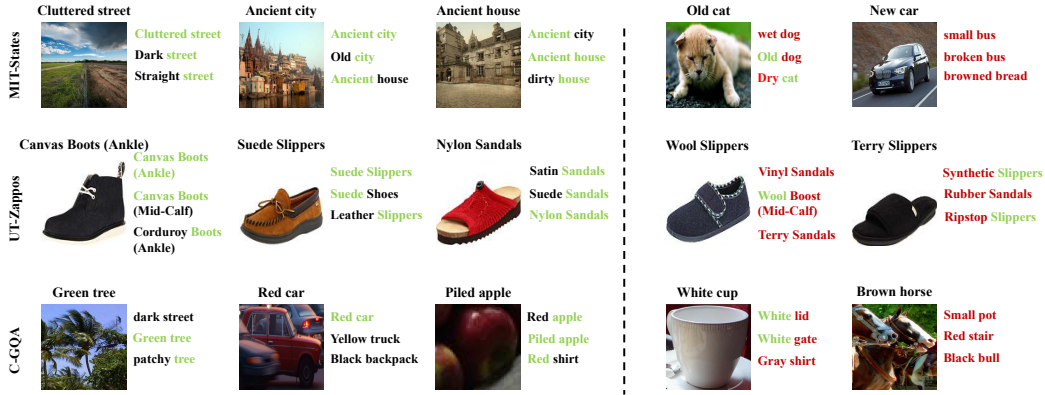


Figure 4. Qualitative results. We show the top-3 predictions of our proposed model for some instances. From the first three columns, we can observe that all predictions of the model indicate it can categorize accurately. However, the model is only incentivized when it matches the label. The task of CZSL is a multi-label one, and future datasets need to account for this. The last two columns show some examples of suboptimal labels and wrong predictions.




Figure 5. The influence of the weighting coefficient β for H value.




Figure 6. The influence of the weighting coefficient β for AUC .

5. Conclusion

In this paper, we propose a novel Siamese Contrastive Embedding Network (SCEN) to excavate discriminative prototypes of state and object for the CZSL task. We firstly project the visual feature into two contrastive spaces, where

we set up *state-constant* and *object-constant* databases. Meanwhile, we design state-specific and object-specific loss functions as constraints, forcing them to contain discriminative corresponding information. In addition, we design a State Transition Module (STM) to produce virtual but rational compositions that never appear in training, which effectively augment the diversity of training data. The proposed module can provide a robust model that can excavate prototypes for seen samples and be generalized to novel compositions, where linear softmax classifiers can be trained to recognize compositions from both seen and unseen instances. The comparison and ablation study experiments demonstrate that our proposed CZSL framework has achieved state-of-the-arts on three challenging datasets.

6. Acknowledgements

Our work was supported in part by the National Key R&D Program of China under Grant 2017YFE0104100, in part by the National Natural Science Foundation of China under Grant 62132016, Grant 62171343, and Grant 62071361, in part by Key Research and Development Program of Shaanxi under Grant 2021ZDLGY01-03, and in part by the Fundamental Research Funds for the Central Universities ZDRC2102.

References

- [1] Yuval Atzmon, Felix Kreuk, Uri Shalit, and Gal Chechik. A causal view of compositional zero-shot recognition. *arXiv preprint arXiv:2006.14610*, 2020. [2](#)
- [2] Irving Biederman. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2):115, 1987. [2](#)
- [3] 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 *European conference on computer vision*, pages 52–68. Springer, 2016. [6](#)
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [2](#)
- [5] Wuyang Chen, Zhiding Yu, Shalini De Mello, Sifei Liu, Jose M Alvarez, Zhangyang Wang, and Anima Anandkumar. Contrastive syn-to-real generalization. *arXiv preprint arXiv:2104.02290*, 2021. [2](#)
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. [5](#)
- [7] Yanan Gu, Cheng Deng, and Kun Wei. Class-incremental instance segmentation via multi-teacher networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1478–1486, 2021. [2](#)
- [8] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. [2](#)
- [9] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. [2](#)
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [6](#)
- [11] Donald D Hoffman and Whitman A Richards. Parts of recognition. *Cognition*, 18(1-3):65–96, 1984. [2](#)
- [12] Phillip Isola, Joseph J Lim, and Edward H Adelson. Discovering states and transformations in image collections. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1383–1391, 2015. [5, 7](#)
- [13] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning, 2021. [2](#)
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [15] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):453–465, 2013. [2](#)
- [16] Xiangyu Li, Zhe Xu, Kun Wei, and Cheng Deng. Generalized zero-shot learning via disentangled representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1966–1974, 2021. [2](#)
- [17] Yong-Lu Li, Yue Xu, Xiaohan Mao, and Cewu Lu. Symmetry and group in attribute-object compositions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11316–11325, 2020. [1, 2, 7](#)
- [18] Massimiliano Mancini, Muhammad Ferjad Naeem, Yongqin Xian, and Zeynep Akata. Open world compositional zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5222–5230, 2021. [7](#)
- [19] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013. [2](#)
- [20] Ishan Misra, Abhinav Gupta, and Martial Hebert. From red wine to red tomato: Composition with context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1792–1801, 2017. [1, 2, 3, 7](#)
- [21] Andriy Mnih and Koray Kavukcuoglu. Learning word embeddings efficiently with noise-contrastive estimation. *Advances in neural information processing systems*, 26:2265–2273, 2013. [2](#)
- [22] Muhammad Ferjad Naeem, Yongqin Xian, Federico Tombari, and Zeynep Akata. Learning graph embeddings for compositional zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 953–962, 2021. [2, 5, 6, 7](#)
- [23] Tushar Nagarajan and Kristen Grauman. Attributes as operators: factorizing unseen attribute-object compositions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 169–185, 2018. [1, 2, 7](#)
- [24] Zhixiong Nan, Yang Liu, Nanning Zheng, and Song-Chun Zhu. Recognizing unseen attribute-object pair with generative model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8811–8818, 2019. [1](#)
- [25] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimeshain, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019. [6](#)
- [26] Senthil Purushwalkam, Maximilian Nickel, Abhinav Gupta, and Marc’Aurelio Ranzato. Task-driven modular networks for zero-shot compositional learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3593–3602, 2019. [1, 2, 3, 6, 7](#)
- [27] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. [6](#)

- [28] Xin Wang, Fisher Yu, Trevor Darrell, and Joseph E Gonzalez. Task-aware feature generation for zero-shot compositional learning. *arXiv preprint arXiv:1906.04854*, 2019. 2
- [29] Kun Wei, Cheng Deng, and Xu Yang. Lifelong zero-shot learning. In *IJCAI*, pages 551–557, 2020. 2
- [30] Kun Wei, Muli Yang, Hao Wang, Cheng Deng, and Xianglong Liu. Adversarial fine-grained composition learning for unseen attribute-object recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3741–3749, 2019. 2
- [31] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2251–2265, 2018. 3
- [32] Muli Yang, Cheng Deng, Junchi Yan, Xianglong Liu, and Dacheng Tao. Learning unseen concepts via hierarchical decomposition and composition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10248–10256, 2020. 2
- [33] Aron Yu and Kristen Grauman. Fine-grained visual comparisons with local learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 192–199, 2014. 5, 7
- [34] Han Zhao, Xu Yang, Zhenru Wang, Erkun Yang, and Cheng Deng. Graph debiased contrastive learning with joint representation clustering. In *Proc. IJCAI*, pages 3434–3440, 2021. 2