









# C2C: Component-to-Composition Learning for Zero-Shot Compositional Action Recognition

Rongchang Li<sup>1,2</sup>, Zhenhua Feng<sup>1,2,3,4</sup>, Tianyang Xu<sup>1</sup>, Linze Li<sup>1</sup>,  
Xiao-Jun Wu<sup>1\*</sup>, Muhammad Awais<sup>2,4</sup>, Sara Atito<sup>2,4</sup>, and Josef Kittler<sup>2,3</sup>

<sup>1</sup> School of Artificial Intelligence and Computer Science, Jiangnan University, China

<sup>2</sup> Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, UK

<sup>3</sup> School of Computer Science and Electronic Engineering, University of Surrey, UK

<sup>4</sup> Surrey Institute of People-centred AI (SI-PAI), University of Surrey, UK

{li\_rongchang, linze.li}@stu.jiangnan.edu.cn;

feng-zhenhua@outlook.com; {wu\_xiaojun, tianyang\_xu}@jiangnan.edu.cn;

{muhammad.awais, sara.atito, j.kittler}@surrey.ac.uk;

**Abstract.** Compositional actions consist of dynamic (verbs) and static (objects) concepts. Humans can easily recognize unseen compositions using the learned concepts. For machines, solving such a problem requires a model to recognize unseen actions composed of previously observed verbs and objects, thus requiring so-called compositional generalization ability. To facilitate this research, we propose a novel Zero-Shot Compositional Action Recognition (ZS-CAR) task. For evaluating the task, we construct a new benchmark, Something-composition (Sth-com), based on the widely used Something-Something V2 dataset. We also propose a novel Component-to-Composition (C2C) learning method to solve the new ZS-CAR task. C2C includes an independent component learning module and a composition inference module. Last, we devise an enhanced training strategy to address the challenges of component variations between seen and unseen compositions and to handle the subtle balance between learning seen and unseen actions. The experimental results demonstrate that the proposed framework significantly surpasses the existing compositional generalization methods and sets a new state-of-the-art. The new Sth-com benchmark and code are available at [https://github.com/RongchangLi/ZSCAR\\_C2C](https://github.com/RongchangLi/ZSCAR_C2C).


**Keywords:** Zero-shot compositional action recognition · Action recognition · Compositional generalization

## 1 Introduction

Humans can understand basic dynamic and static concepts (*i.e.*, verbs and objects) in seen compositional actions and generalise the concepts to unseen action categories. For example, by seeing *close a box* and *open a bag*, we can envision *close a bag* even if we have never seen this before. This compositional generalization capability is a fundamental function of human intelligence, by which we can

---

\* Corresponding Author.



**Fig. 1:** Zero-Shot Compositional Action Recognition (ZS-CAR) requires models to recognize unseen actions composed of verbs and objects observed in seen actions.

understand an ‘infinite’ number of complex concepts through assimilating the learned ‘finite’ concepts [8, 14]. Compositional generalization has recently gained attention in the computer vision community, giving rise to the corresponding image analysis task Compositional Zero-Shot Learning (CZSL) [43]. In video analysis, [42] proposed the compositional action recognition task. However, this task only requires recognizing the dynamic component. It makes the task essentially a regular verb classification task characterized by explicit sample selection bias [20, 68] regarding verb-object compositions, rather than verifying compositional generalization ability. Exploring the compositional generalization problem in the video modality is valuable as it presents a challenging video understanding task that requires comprehending both dynamic and static components while also generalising to novel compositions. Also, it provides a reliable way to develop open-scenario-related applications in autonomous vehicles, robots, etc.

To facilitate the related research, we propose a novel task, namely *Zero-Shot Compositional Action Recognition (ZS-CAR)*. We define a compositional *action* as the composition of a dynamic component *verb* (e.g., open, pick up, ...) and a static component *object* (e.g., door, wallet, ...). As shown in Fig. 1, ZS-CAR aims to enable machines to identify unseen action categories consisting of seen verbs and objects. The new proposed task poses a significant challenge. First, a model has to understand the intricate spatial and temporal information to identify the verb and object components, which is more challenging than the existing action recognition task. Second, the task naturally contains *component domain variation* and *component compatibility variation*. The first one is similar to the challenges in domain generalization, caused by the contextuality difference (e.g., transferring knowledge of *book* in *close a book* to recognize *bend a book*). The second one stems from the action category difference in the training and test domains. Overly accepting or rejecting verb-object relationships observed during training can detrimentally impact the generalization capability to recognize unseen actions. Given these challenges, the existing visual compositional generalization solutions cannot handle the ZS-CAR task well, as shown in Tab. 2.

To solve ZS-CAR, we present an innovative Component-to-Composition (C2C) method to serve as the baseline. The pipeline of C2C is shown in Fig. 2. C2C first independently learns the dynamic and static components (i.e., verbs and objects). Then it employs the component consensus to recognize compositions (i.e., actions) from both dynamic and static perspectives, ensuring that both these two components are well learned. The procedure is motivated by the human process

of assembling compositions. It involves an independent observation of a component, assessing its compatibility with another part, and gradually assembling the units into the entire structure. Moreover, C2C enables explicitly addressing the *component domain variation* and *component compatibility variation* problems. Accordingly, we propose an enhanced training strategy that involves minimizing the dependence on spurious information within visual features and balancing the emphasis between fitting seen and imagining unseen actions.

Last, to enable the training of C2C and to evaluate its performance for the ZS-CAR task, we create a new dataset, Something-composition (Sth-com), based on the popular Something-Something V2 (Sth-V2) dataset [16]. The experimental results demonstrate that the vanilla or enhanced version of C2C achieves state-of-the-art results. To summarize, the main contributions of the paper include:

- A novel zero-shot compositional action recognition (ZS-CAR) task that investigates compositional generalization in the video modality.
- A benchmarking dataset for evaluating ZS-CAR. We elaborate on the challenges in ZS-CAR and demonstrate that the existing visual compositional generalization methods cannot handle such problems well.
- A novel Component-to-Composition (C2C) method as the vanilla solution and an enhanced training strategy to address the identified challenges in ZS-CAR. The experimental results demonstrate that C2C solves these challenges effectively and attains SOTA performances.

## 2 Related work

**Video action recognition** aims to identify actions in videos. The main challenge is extracting discriminative action features from videos. Previous CNN-based methods address this challenge in two primary ways: using 2D CNNs with separate temporal operations [27, 28, 30–32, 34, 35, 58, 59, 62], or jointly using 3D CNNs [5, 12, 13, 19, 23, 51, 57, 61] to dig out spatiotemporal features. Recently, Transformer-based models transform a video into 3D tokens and model the spatiotemporal correlations through attention layers [1, 2, 11, 29, 37, 46, 49], or adapt the visual image encoder in large vision-language models to the video modality [47, 48, 66].

**Zero-shot action recognition (ZSAR)** aims to recognize action categories not appearing in the training set. A widely used idea is to measure video features and action embeddings in a common semantic space, hoping this space can generalize to unseen actions. The semantic embeddings are obtained by side semantic information, *e.g.*, word embeddings [4, 41, 63, 69], multiple attributes [15, 36], elaborative rehearsals [7] and action descriptions [9, 52]. Different from ZSAR, the proposed ZS-CAR aims to explore the visual compositionality of actions and generalize the learned composition knowledge to unseen actions.

**Compositional action recognition:** The compositional nature of actions has recently received increasing attention. Action genome [22] takes action dynamics as the relationship changes and decomposes actions into spatiotemporal scene graphs. Something-else [42] proposes compositional action recognition that

only requires recognising verbs in unseen verb-object compositions. Subsequent studies [39, 56, 64, 65] follow this setting and focus on extracting robust dynamic information. However, [42] essentially is a supervised task with explicit sample selection bias, rather than a compositional generalization problem.

**Image-based compositional generalization:** In the image modality, [43] first introduces the Compositional Zero-Shot Learning (CZSL) task that requires models trained on seen attribute-object pairs to recognize unseen ones. A common idea is to find the nearest holistic composition embeddings for input images in a common space [40, 43, 44, 50, 54]. Some studies disentangle visual features to align with component embeddings [18, 24, 26, 33, 54, 55]. Recent works propose to adapt the vision-language model [53] to CZSL by text prompting methods [38, 45] and enhancing the interactions between language and visual features [38].

### 3 Problem Definition and Benchmark

#### 3.1 The Proposed ZS-CAR Task

In **traditional action recognition**, a model predicts the action label defined as either a verb or a verb-object composition. For example, the Sth-V2 dataset [16] defines verb-only actions, such as *open sth* and *pick up sth*. Kinetics-400 [25] defines actions as both verb-object compositions and verb-only actions, such as *play basketball*, *draw*, etc. The recently proposed **compositional action recognition** task considers the interaction between verbs and objects. It requires recognising verbs of which the paired objects in the test set differ from those in the training set. Though introducing action compositionally by this constraint, this task only takes the verb-object relation as a sample selection bias controller, making it not a thorough compositional generalization problem.

In contrast to previous tasks, we introduce the compositional generalization perspective to understand actions, by proposing a novel Zero-Shot Compositional Action Recognition (ZS-CAR) task. We explicitly define a compositional action as a concept of *verb* and *object*. To verify the compositional generalization ability, ZS-CAR requires models to recognize unseen action categories composed of verbs and objects observed in seen actions. Formally, we denote the training set as  $D_{\text{train}} = \{(\mathcal{X}_i, a_i)\}_{i=1}^N$  consisting of video  $\mathcal{X}$  labelled as action  $a \in A_{\text{train}}$ , with  $A_{\text{train}} \subset V \times O = \{(v, o) | v \in V, o \in O\}$ , where  $V$  and  $O$  are the verb and object sets. The validation and test sets are similarly denoted as  $D_{\text{val}}$  and  $D_{\text{test}}$ , with label sets as  $A_{\text{val}}$  and  $A_{\text{test}}$ . We follow the generalised zero-shot learning setting [6, 50], *i.e.*, both seen and unseen action categories appear in the validation/test set:  $A_{\text{train}} \cap A_{\text{val/test}} \neq \emptyset$  and  $\exists a(a \in A_{\text{val/test}} \oplus a \in A_{\text{train}})$ .

#### 3.2 The Something-composition Benchmark

To assess the proposed task, we build a benchmark Something-composition (Sth-com), based on the widely used Something-Something V2 (Sth-V2) dataset [16].

**Table 1:** Statistics of the proposed Sth-com benchmark.

	# Objects	# Verbs	# Verb-Object Actions	# Sample
Train	248	161	3451	38034
Val	225	157	733 Seen + 717 Unseen	18774
Test	233	161	976 Seen + 956 Unseen	22657
All	248	161	5124	79465

Sth-V2 has 220,847 videos of humans performing pre-defined compositional actions with everyday objects. It is widely used by the action recognition community, as its verbs cannot be identified from a single frame.


**Annotation:** We use the training and test split of Sth-else [42], a subset of Sth-V2, as the starting point. Then, we iteratively clean the subsets and ensure that: (1) Each composition (*i.e.*, action) has more than five samples. (2) The components (*i.e.*, verb & object) present in the test set are seen in the training set. This yields 248 objects, 161 verbs, 5124 compositions and 79465 video samples. Next, we select 1/3 of compositions in the training and test sets respectively and interchange 50% of the selected compositions, thereby introducing seen action categories to the test set. Subsequently, following the ratio (3:4) in MIT-State [21], we divide the test set into the validation and test sets, respectively. The statistics of Sth-com are summarised in Tab. 1.

## 4 Component-to-Composition Learning

To address the proposed ZS-CAR task, we propose a novel Component-to-Composition (C2C) learning framework, as shown in Fig. 2, to serve as a vanilla solution. We simulate the process of human assembly to infer compositions. We commence with a specific component and scrutinise the compatibility of various categories of another component, leading to a reasonable composition proposal. These are correspondingly implemented in the proposed *independent component learning* module and *composition to component learning* module. Importantly, we analyse the challenges in ZS-CAR and devise an *enhanced training strategy* to facilitate generalising to unseen actions better.

### 4.1 Independent Component Learning

In this module, our goal is to identify the verb/object score of an input video and acquire their semantic prototypes. Specifically, we first encode the raw video sequence  $\mathcal{X}$  with shape  $T \times 224 \times 224 \times 3$  to a general representation  $\mathbf{F}_x \in \mathbb{R}^{T \times D}$ , where  $T$  and  $D$  denote the frame and channel numbers, respectively. Then, we respectively adopt dynamic and static encoders to extract dynamic and static features, denoted as  $\mathbf{f}_v$  and  $\mathbf{f}_o \in \mathbb{R}^C$ , where  $C$  is the channel number. The static encoder consists of a temporal average pooling layer and two MLP layers, while the dynamics encoder contains two temporal convolution layers and a temporal pooling layer. We encode the component labels as the general representations of verbs and objects, denoted as the component prototypes  $\mathbf{E}_v = \{e_{v,i}\}_{i=1}^{N_v}$  and



**Fig. 2: The proposed Component-to-Composition (C2C) learning framework.** C2C first aligns verb/object prototypes with corresponding visual features to obtain component scores in the **Independent Component Learning** module. Then, actions are inferred through two paths (dynamics and static) in the **Component to Composition** module. In the dynamics path, verb prototypes and visual features are used to compute conditional object scores. Then the independent verb scores and conditional object scores are multiplied to gain the action scores. The static path follows a similar procedure. The final output is a consensus of the results from both paths.

$\mathbf{E}_o = \{e_{o,i}\}_{i=1}^{N_o}$ , where  $N_v$  and  $N_o$  are the verb and object numbers, respectively. The prototypes are learned by being aligned with the visual features  $\mathbf{f}_v$  and  $\mathbf{f}_o$  using cosine similarity as the gauge. Denoting the action label of the input video as  $a_i = (v_l, o_k)$ , the optimization objective is defined as:

$$\mathcal{L}_{verb} = -\log \frac{\exp(\cos(\mathbf{f}_v, \mathbf{e}_{v,l})/\tau)}{\sum_{v_j \in V} \exp(\cos(\mathbf{f}_v, \mathbf{e}_{v,j})/\tau)}, \quad (1)$$

$$\mathcal{L}_{obj} = -\log \frac{\exp(\cos(\mathbf{f}_o, \mathbf{e}_{o,k})/\tau)}{\sum_{o_j \in O} \exp(\cos(\mathbf{f}_o, \mathbf{e}_{o,j})/\tau)}, \quad (2)$$

where  $\mathbf{e}_{v,l}$  and  $\mathbf{e}_{o,k}$  are the corresponding prototypes.  $\cos(\cdot)$  is the cosine similarity.  $\tau$  is the temperature coefficient for better optimization. We denote  $\mathcal{L}_{comp} = \mathcal{L}_{verb} + \mathcal{L}_{obj}$  as the component loss. By aligning with diverse component features in the training set, the prototypes inherently acquire semantic meanings. Accordingly, the verb and object scores,  $s_v$  and  $s_o$ , are the cosine similarities between the visual features and corresponding component prototypes.

## 4.2 From component to composition

In this module, we infer the action (composition) score  $\mathbf{s}_a \in \mathbb{R}^{N_a}$  using the prototypes and the verbs and objects (components) scores. We define the conditional

scores  $\mathbf{S}_{o|v} \in \mathbb{R}^{N_v \times N_o}$  and  $\mathbf{S}_{v|o} \in \mathbb{R}^{N_o \times N_v}$ , that represents the confidence of one component given the visual input and another component as the conditions. Accordingly, for one action  $a_i = (v_l, o_k)$ , its confidence score can be calculated by either  $s_{a,i} = s_{v,l} \cdot s_{o=o_k|v=v_l}$  or  $s_{a,i} = s_{o,k} \cdot s_{v=v_l|o=o_k}$ . The first one relies on the verbs to infer actions, called the dynamics path. Similarly, the other one is called the static path. As shown in Fig. 2, in the dynamics path, we first encode  $\mathbf{F}_x$  using a static encoder to provide the visual reference for the object component. Then we concatenate the output with the verb prototype  $\mathbf{e}_{v,l}$  and fuse them to the compound prototype-visual feature  $\mathbf{f}_{v-x,l}$ , using an MLP layer. Then we linearly map the cosine similarity  $\cos(\mathbf{f}_{v-x,l}, \mathbf{e}_{o,k})$  to the range of  $[0, 1]$ , obtaining the conditional object score  $s_{o=o_k|v=v_l}$ . The static path follows a similar process except for encoding  $\mathbf{F}_x$  by the dynamics encoder. Last, we average the inference results from the two paths to obtain the final inference output. Taking  $s_{a,i}$  as the score of the groundtruth action  $a_i$ , the entire model is optimized using:


$$\mathcal{L}_{com} = -\log \frac{\exp(s_{a,i}/\tau)}{\sum_{a_j \in A_{train}} \exp(s_{a,j}/\tau)}, \quad (3)$$

### 4.3 Enhanced training strategy

**Risk in generalizing components.** Naturally, there is a high probability of exhibiting different component appearances between training and test domains, namely *component domain variation*. As shown in Fig. 3, the component ‘book’ exhibits a notable difference in appearance between ‘close a book’ and ‘bend a book’. Besides, given the combination of ‘close’ and ‘book’ in the training set, there is a tendency that the model might align features associated with ‘close’ to ‘book’, or vice versa. These dependencies may confuse the model when encountering unseen action categories.

To cope with it, in the independent component learning module, we posit that the ideal dynamic-focused feature  $\mathbf{f}_v$  and static-focused feature  $\mathbf{f}_o$  should encompass two types of information: common and specific. The former denotes features pertinent to both components and the latter denotes the invariant features, that are crucial for identifying the respective components. Notably, the invariant features of verbs are spurious for identifying objects, and vice versa. As shown in Fig. 3, our solution is to minimize the spurious information in  $\mathbf{f}_v$  and  $\mathbf{f}_o$ , and explicitly reduce the mutual associations between their component-specific parts. To achieve this, we propose the following auxiliary losses:

$$\mathcal{L}_{sup,verb} = h(\mathbf{f}_x, \mathbf{f}_v) - h(\mathbf{f}_v, \mathbf{y}_v), \quad (4)$$



**Fig. 3: Component domain variations.** For unseen actions, an object may exhibit different appearances (left). To deal with this (right), we reduce the spurious information and enlarge the independence between component-specific features.

$$\mathcal{L}_{sup,obj} = h(\mathbf{f}_x, \mathbf{f}_o) - h(\mathbf{f}_o, \mathbf{y}_o), \quad (5)$$

$$\mathcal{L}_{ind} = \mathcal{L}_{sup,verb} + \mathcal{L}_{sup,obj} + h(\mathbf{f}'_v, \mathbf{f}'_o), \quad (6)$$

where  $h(\cdot)$  is the Hilbert-Schmidt Independence Criterion (HSIC) [17] for independence measure.  $h(\mathbf{f}, \mathbf{h}) = 0$  if and only if  $\mathbf{f} \perp \mathbf{h}$ .  $\mathbf{f}_x$  is the temporal pooling result of  $\mathbf{F}_x$ .  $\mathbf{y}_v$  and  $\mathbf{y}_o$  are the one-hot label vectors.  $\mathcal{L}_{sup,verb}$  and  $\mathcal{L}_{sup,obj}$  force the model to compress information, but retain sufficient information to identify the components, thus reducing the general spurious information.  $\mathbf{f}'_v$  or  $\mathbf{f}'_o$  is part of  $\mathbf{f}_v$  or  $\mathbf{f}_o$ , defined as  $\mathbf{f}'_v = \mathbf{f}_v[:\rho C]$ ,  $\mathbf{f}'_o = \mathbf{f}_o[:\rho C]$  ( $0 \leq \rho \leq 1$ ). They are seen as component-specific features. So the term  $h(\mathbf{f}'_v, \mathbf{f}'_o)$  ensures the independence of the invariant part of the verb and object features.

**Balanced learning of seen & unseen actions.** In ZS-CAR, the verb-object connections in the training and test sets differ significantly, as shown in Fig. 4. While the seen connections help optimize the modules, an excessive emphasis on them may hinder generalizing to unseen actions. Our solution is simultaneously enhancing the learning of seen and unseen connections, enabling them to strike an optimal balance.


We first assist the model in *learning seen actions explicitly*. We calculate the observed conditional probabilities in the training set as the observed conditional scores  $\hat{\mathbf{S}}_{v|o}$  and  $\hat{\mathbf{S}}_{o|v}$ . Then, we force the conditional scores in a training batch closer to  $\hat{\mathbf{S}}_{v|o}$  and  $\hat{\mathbf{S}}_{o|v}$ , with the objectives:

$$\mathcal{L}_{con,v} = -\frac{1}{N_v} \sum_{v_i \in V} \hat{\mathbf{s}}_{o|v=v_i} \log \mathbf{s}_{o|v=v_i}, \quad (7)$$

$$\mathcal{L}_{con,o} = -\frac{1}{N_o} \sum_{o_i \in O} \hat{\mathbf{s}}_{v|o=o_i} \log \mathbf{s}_{v|o=o_i}, \quad (8)$$

where  $N_v$  &  $N_o$  denote the verb & object numbers, and the condition loss  $\mathcal{L}_{con}$  is defined as the sum of  $\mathcal{L}_{con,v}$  and  $\mathcal{L}_{con,o}$ .  $\hat{\mathbf{s}}_{o|v=v_i}$  or  $\mathbf{s}_{o|v=v_i}$  denotes the conditional object score vector when verb is  $v_i$ .  $\hat{\mathbf{s}}_{v|o=o_i}$  and  $\mathbf{s}_{v|o=o_i}$  are defined similarly.

Then, inspired by the human ability to create new things by *imagining compositions of known components*, we use CutMix [67] to assemble samples in a training batch as the imagined unseen actions, thereby reducing the risk of overemphasizing seen compositions. Specifically, for an input  $\mathcal{X}_i$  with label  $a_i = (v_l, o_k)$ , we randomly select another sample  $\mathcal{X}_j$  with label  $a_j = (v_m, o_n)$ ,



**Fig. 4: Component compatibility variations.** The component relations are different between the training and test sets (left). To solve this (right), we use the observed conditional score to fit seen relations and encourage the model to imagine unseen actions to avoid being limited to seen relations.



**Table 2:** A comparison with the SOTA methods on Sth-com, using the vision model with word embedding model paradigm. We report both vanilla (only  $\mathcal{L}_{com}$  and  $\mathcal{L}_{comp}$ ) and enhanced (using the advocated enhanced training strategy) versions of C2C. Since ADE is tailored for transformer architectures, we only adapt it to VideoSwin.

Methods	TSM-18						VideoSwin-T					
	verb	obj.	seen	unseen	HM	AUC	verb	obj.	seen	unseen	HM	AUC
TMN [50]	41.4	35.9	31.6	28.1	22.0	7.3	33.4	45.6	31.0	30.2	23.0	7.9
Compos [40]	45.8	40.0	34.8	35.1	27.5	10.6	36.7	50.5	34.2	37.8	27.2	10.9
CGE [44]	35.5	36.4	27.6	25.5	20.0	5.9	31.8	48.1	31.2	30.3	23.1	7.9
OADis [55]	39.1	40.7	33.5	32.3	25.5	9.3	41.6	48.9	38.1	38.3	30.2	12.6
ADE [18]	-	-	-	-	-	-	35.3	49.1	33.8	36.2	26.8	10.3
CAN [60]	40.0	39.2	32.4	31.6	24.5	8.7	32.6	49.5	33.0	33.3	25.0	9.2
<b>C2C (Vanilla)</b>	47.7	42.7	38.1	39.1	30.8	13.2	47.9	50.9	42.1	43.6	34.9	16.4
<b>C2C (Enhanced)</b>	<b>50.1</b>	<b>44.3</b>	<b>39.4</b>	<b>41.0</b>	<b>33.0</b>	<b>14.5</b>	<b>48.8</b>	<b>52.8</b>	<b>43.4</b>	<b>44.1</b>	<b>36.7</b>	<b>17.4</b>

and crop it at a random position of random size. The cropped part replaces the same position of  $\mathcal{X}_i$ , generating the augmented input  $\tilde{\mathcal{X}}_i$ . Denoting the cropping ratio as  $\lambda$ , we modify the previous losses as:  $\tilde{\mathcal{L}}_{com} = (1 - \lambda)\mathcal{L}_{com}^{a_i} + \lambda\mathcal{L}_{com}^{a_j}$  and  $\tilde{\mathcal{L}}_{comp} = (1 - \lambda)\mathcal{L}_{comp}^{v_l, o_k} + \lambda\mathcal{L}_{comp}^{v_m, o_n}$ .  $\mathcal{L}_{com}^{a_i}$  calculates  $\mathcal{L}_{com}$  using Eq. (3) with  $a_i$  as groundtruth.  $\mathcal{L}_{comp}^{v_l, o_k}$  calculates  $\mathcal{L}_{comp}$  with  $v_l$  and  $o_k$  as groundtruth.  $\mathcal{L}_{com}^{a_j}$  and  $\mathcal{L}_{comp}^{v_m, o_n}$  are defined similarly. Also,  $\mathcal{L}_{ind}$  is modified to  $\tilde{\mathcal{L}}_{ind}$  by replacing  $\mathbf{y}_i$  to  $(1 - \lambda)\mathbf{y}_i + \lambda\mathbf{y}_j$  in Eq. (4) and Eq. (5). Particularly, we introduce new compositions  $\bar{a}_1 = (v_l, o_n)$  and  $\bar{a}_2 = (v_m, o_k)$ .  $\bar{a}_1$  and  $\bar{a}_2$  are not constrained in  $A_{test}$  or  $A_{train}$ . It gives a novel loss  $\mathcal{L}_{new} = \mathcal{L}_{com}^{\bar{a}_1} + \mathcal{L}_{com}^{\bar{a}_2}$ , which forces the model to consider all the possible component connections.

A concern is that the new compositions  $\bar{a}_1$  or  $\bar{a}_2$  may be counterfactual. However, it only introduces noise to the component-to-composition module, as the component labels still match the visual input. In this way, we still reserve these counterfactual labels as they can provide regularization effects that avoid overfitting to seen verb-object relations. Besides, it is also consistent with human imagination which is not only constrained within reasonable compositions.

#### 4.4 More details

**Network details.** Following previous practice, we implement the video encoder and component label encoding with two paradigms:

(1) Vision model with word embedding model. In this paradigm, we adopt an action recognition model (the CNN-based TSM-18 [35] or the transformer-based VideoSwin-T [37]) to extract the general video feature and a word embedding model (fasttext [3]) to encode the labels.

(2) Vision-language models. Recently, large pre-trained vision-language models, such as CLIP [53], have shown superior power. In this paradigm, we adopt the language branch of recent prompt-based compositional generalization methods (Coop [70], CSP [45] and SPM [38]) to encode the labels. As CLIP are trained using image-text pairs, we apply image-to-video adapters introduced in AIM [66] to adapt the CLIP visual encoder (ViT-B/32 [10]) to video modality.

**Table 3:** A comparison with the SOTA CLIP-based methods on Sth-com.

methods	verb	obj.	seen	unseen	HM	AUC
Coop [70]	48.4	55.1	43.9	46.4	36.6	18.1
+C2C(vanilla)	56.2	56.9	48.3	54.6	42.7	24.0
+C2C(enhance)	<b>56.3</b>	<b>58.9</b>	<b>50.7</b>	<b>56.6</b>	<b>44.5</b>	<b>26.0</b>
CSP [45]	49.0	54.8	43.6	47.4	36.0	18.0
+C2C(vanilla)	56.3	56.6	49.0	53.4	42.5	23.8
+C2C(enhance)	<b>57.1</b>	<b>58.6</b>	<b>49.9</b>	<b>56.2</b>	<b>44.3</b>	<b>25.5</b>
SPM [38]	48.9	54.5	44.0	46.3	35.5	17.7
+DFM [38]	47.4	55.4	43.1	47.3	35.8	17.9
+C2C(vanilla)	55.7	57.3	49.3	53.4	42.2	23.8
+C2C(enhance)	<b>56.5</b>	<b>58.4</b>	<b>49.9</b>	<b>55.7</b>	<b>43.8</b>	<b>25.2</b>

**Training Objectives.** When applying the proposed enhanced training strategy, we perform CutMix to generate new compositions with a probability  $p$ . If apply CutMix, the final loss is calculated as:

$$\bar{\mathcal{L}}_{total} = \bar{\mathcal{L}}_{com} + \alpha \bar{\mathcal{L}}_{comp} + \beta \bar{\mathcal{L}}_{ind} + \gamma \mathcal{L}_{new}. \quad (9)$$

If not, the final loss is:

$$\mathcal{L}_{total} = \mathcal{L}_{com} + \alpha \mathcal{L}_{comp} + \beta \mathcal{L}_{ind} + \gamma \mathcal{L}_{con}. \quad (10)$$

As  $\mathcal{L}_{new}$  and  $\mathcal{L}_{con}$  respectively guide the model focus on unseen and seen actions, we use the same coefficient  $\gamma$ .

## 5 Experiments

**Evaluation Metrics.** The metrics include **Verb**, **Object**, **Seen**, **Unseen**, **HM** and **AUC**. **Verb** and **Object** respectively denote the accuracy of verbs and objects. To gain other metrics, we first add a constant to the seen actions scores. The constant ranges from negative infinity to positive infinity, making the model shift from completely biased to unseen actions to completely biased to seen actions and yielding the seen-unseen curve shown in Fig. 5. **Seen** and **Unseen** are the best accuracy of seen/unseen actions, *i.e.* the intersection points with x/y-axis in Fig. 5. **HM** is the best harmonic mean with formulation:  $HM = 2(Seen * Unseen)/(Seen + Unseen)$ . **AUC** is the area under the curve.

**Implementation details.**  $\alpha$ ,  $\beta$   $\gamma$  in Eq. (9) and Eq. (10) are 0.2, 0.1 and 0.1, respectively. The probability  $p$  of conducting CutMix is 0.7. More coefficient ablations and implementation details are reported in supplementary materials.

### 5.1 Comparisons with State-of-the-Art

Note that the existing compositional action recognition methods cannot apply to the ZS-CAR task. To comprehensively demonstrate the proposed C2C, we first adapt the existing CZSL methods to the proposed Sth-com dataset. We replace their image encoders with video encoders and adopt the same word

**Table 4:** A comparison with the SOTA methods on C-GQA.

	multi	seen	unseen	HM	AUC
CGE [44]	×	38.0	17.1	18.5	5.4
OADis [55]	×	38.3	19.8	20.1	7.0
CoT [26]	✓	39.2	22.7	22.1	7.4
<b>C2C</b>	×	39.1	23.0	23.0	7.7
<b>C2C</b>	✓	<b>39.4</b>	<b>23.6</b>	<b>23.3</b>	<b>8.2</b>

**Table 5:** Results of different hypotheses inferring compositions from components.

	verb	obj.	seen	unseen	HM	AUC
Compcos [40]	45.8	40.0	34.8	35.1	27.5	10.6
Independent	47.9	42.2	37.4	37.4	29.8	12.2
Knowledge-agnostic	<b>48.4</b>	<b>43.1</b>	<b>38.7</b>	36.8	30.2	12.6
Dynamic centered	46.7	42.5	37.4	38.0	29.9	12.5
Static centered	46.6	41.1	37.0	36.7	28.9	11.7
Full	47.7	42.7	38.1	<b>39.1</b>	<b>30.8</b>	<b>13.2</b>

embedding model, making sure only the compositional generalization parts are different between the adapted methods and our proposed methods. Besides, we also modify C2C and compare it with the existing SOTA methods on the image-based CZSL task using the challenging C-GQA [44] dataset.

**Comparisons on Sth-com.** We first compare our method with the adapted CZSL methods of the vision model with the word embedding paradigm, using both CNN-based and transformer-based video encoders. The results are reported in Tab. 2. The vanilla version of C2C uses two basic losses. The enhanced version uses the proposed enhanced training strategy.

Tab. 2 demonstrate that our method outperforms the existing methods by significant margins across all the metrics on ZS-CAR. Furthermore, the proposed enhanced training strategy provides consistent performance gains across different video encoders. Tab. 2 shows that the recently proposed methods based on Compcos are inferior to Compcos. Second, except for OADis, using the recently proposed VideoSwin-T backbone does not yield significant improvements. The first observation may stem from failing to generalise the embedding relations, not disentangling spatial and temporal information well, and not applying the conditional component relations in actions. Besides, jointly learning the compositions may lead to unbalanced learning of the components. These deficiencies degrade the effectiveness of video encoders, leading directly to the second observation. Tab. 2 also compares two video encoders. Though VideoSwin is an advanced model, its verb accuracy is consistently inferior or similar to TSM across all methods. Conversely, it outperforms TSM in object accuracy. This implies that our ZS-CAR task offers a new avenue for assessing biases regarding dynamic and static patterns in action recognition models.

Then we apply our C2C method to recently proposed CLIP-based compositional zero-shot learning methods. The results are reported in Tab. 3. Coop, CSP and SPM are three text prompting methods. We can see that the proposed

C2C method can improve their performance significantly on the ZS-CAR task. DFM [38] was proposed along with SPM. It enables cross-modal interactions between language and visual features for better composition prediction. However, the improvements are slight on the ZS-CAR task. Compared to it, our C2C method provides more significant improvements.


**Comparison on C-GQA.** To further evaluate the proposed C2C method, we modify it for the image-based compositional zero-shot learning task by deleting all the temporal-related operations. We report the results on the recently proposed challenging benchmark C-GQA in Tab. 4, with fine-tuned ViT-B [10] as the backbone. In the table, the term ‘multi’ denotes the use of multi-level backbone outputs. When using multi-level outputs in C2C, we merely average them, much simpler than CoT.

Though designed for video analysis, the proposed C2C exhibits superior performance over existing SOTA methods. Like CoT [26], we also try to utilize multi-level backbone outputs. But compared to the specially designed module in CoT, we merely average the multi-level outputs. This demonstrates that our C2C method can be a general solution to the compositional generalization problem.

## 5.2 Ablation Study

**Different composition generation methods.** We first compare different methods that can be used to generate compositions from verb and object components, including: (1) **Independent** supposes components are completely independent, so the inference is straightforward:  $s_{c,i} = s_{v,i} \cdot s_{o,i}$ . (2) **Knowledge-agnostic** assumes components are not completely independent, but the component relations are agnostic:  $s_{c,i} = s_{v,i} \cdot 1 + s_{o,i} \cdot 1$ . (3) **One path inference** considers the compatibility between compositions but only relies on either inference path, *i.e.*, the dynamics or the static path result (Fig. 2). (4) **Full** is the final scheme, integrating both paths.

For a fair comparison, the experiments only use the base losses  $\mathcal{L}_{com}$  and  $\mathcal{L}_{comp}$ . The results are reported in Tab. 5. We use Compcos as the baseline result. The simplest inference method, *i.e.*, *Independent*, has surpassed Compcos on all the metrics. When using the knowledge-agnostic hypothesis, the seen accuracy and holistic metrics are better. But its unseen accuracy is lower. By considering the component relations, utilizing both inference paths substantially improves the unseen accuracy, yielding the best results for both HM and AUC. The results suggest that our proposed inference approach achieves a more balanced performance between seen and unseen actions. Fig. 5 reports the seen-unseen curves, implying that the proposed method is more robust.



**Fig. 5: Seen-unseen curves**, drawn by gradually forcing the model biased from unseen to seen actions.

**Table 6:** Effectiveness of the elements in the enhanced training strategy.

$\mathcal{L}_{con}$	$\mathcal{L}_{comp}$	$\mathcal{L}_{ind}$	CutMix	$\mathcal{L}_{con}$	$\mathcal{L}_{new}$	HM	AUC
✓						29.8	12.5
✓	✓					30.8	13.2
✓	✓	✓				31.4	13.6
✓	✓	✓	✓			31.5	13.9
✓	✓	✓	✓	✓		32.1	14.0
✓	✓	✓	✓	✓	✓	<b>33.0</b>	<b>14.5</b>

**Table 7:** The accuracy (%) on actions containing verbs that cause deformations.

methods	verb		object		composition	
	seen	unseen	seen	unseen	seen	unseen
Vanilla	70.1	60.2	<b>68.7</b>	34.8	56.9	27.5
Vanilla + $\mathcal{L}_{ind}$	<b>71.8</b>	<b>65.2</b>	68.2	<b>45.1</b>	<b>58.0</b>	<b>36.7</b>


**Table 8:** Unbiased accuracy (%) of seen and unseen actions and the corresponding harmonic mean (HM).

methods	seen	unseen	HM
baseline	36.8	17.2	23.5
+ $\mathcal{L}_{con}$	36.5(-0.3)	16.4(-0.8)	22.6(-0.9)
+cutmix	35.7(-1.1)	17.7(+0.5)	23.6(+0.1)
+cutmix+ $\mathcal{L}_{con}$	36.1(+0.4)	17.6(-0.1)	23.7(+0.1)
+cutmix+ $\mathcal{L}_{new}$	37.0(+1.3)	18.2(+0.5)	24.4(+0.8)
+cutmix+ $\mathcal{L}_{con}$ + $\mathcal{L}_{new}$	<b>37.5(+0.5/+1.4)</b>	<b>18.7(+0.5/+1.1)</b>	<b>24.9(+0.5/+1.2)</b>

**The enhanced training strategy.** In Tab. 6, we demonstrate the effectiveness of our enhanced training strategy. It shows that the results exhibit consistent improvements as gradually applying the proposed strategy.

To further validate the proposed method, we verify the merits of the independent loss  $\mathcal{L}_{ind}$  on actions containing verbs that can cause object deformations (such as *squeeze* and *twist*). The results are shown in Tab. 7. They show that the unseen accuracy improves significantly after adding  $\mathcal{L}_{ind}$ . However, the object accuracy of seen actions slightly decreases. This might be because spurious features are useful for recognizing seen components.

Next, we validate the proposed balanced learning of actions. With the vanilla version with  $\mathcal{L}_{ind}$  as the baseline, we report the accuracy of all seen and unseen actions and HM without adding bias in Tab. 8. It shows that adding only  $\mathcal{L}_{con}$  leads to performance degradation, but the degradation in seen accuracy is minor. Similarly, adding just CutMix results in a slight improvement, but with an obvious degradation in seen accuracy. When using CutMix, we add either  $\mathcal{L}_{con}$  or  $\mathcal{L}_{new}$ . The results indicate that  $\mathcal{L}_{new}$  yields a universal improvement, while  $\mathcal{L}_{con}$  contributes positively to seen accuracy. Last, when  $\mathcal{L}_{new}$  and  $\mathcal{L}_{con}$  are integrated, they both bring improvements on all the metrics, and the gains are more pronounced than adding either one alone. This indicates that our training strategy enables the model to achieve a better balance between seen and unseen actions.



**Fig. 6:** The top-3 compatible components **Fig. 7:** Examples of successes and failures with the same verb or object components. of our methods.

### 5.3 Qualitative analysis

In Fig. 6, we show the top three compatible components for different samples, while keeping another component the same. When the verb is *squeeze*, for the first sample, the second and third compatible objects are *cloth* and *paper*, which are close to the visual contents and compatible with *squeeze*. For the second sample, the model identifies different objects. The second object *basket* is not included in the visual content but is compatible with *squeeze*. The third object *bed* cannot be combined with *squeeze* but is visually close to the sample. When the fixed component is an object, the situation is similar. It demonstrates that the proposed model can identify component compatibility according to both visual contents and component attributions.

In Fig. 7, we visualize the top three recognized actions of some successful and failed samples. In the top row, our model accurately predicts the given actions. Besides, the second and third predictions relate to the visual content, such as *fold paper* for the first sample and *tear paper just a little bit* for the second sample. The bottom row shows some failures. Though wrong, the predictions still reflect the video content, such as *move part of plant* and *tear leaf into two pieces* for the second example. However, the failures also demonstrate some limitations of the proposed method. For example, it struggles with the task of distinguishing objects or verbs of similar visual appearances.

## 6 Conclusion

This paper introduces a novel task, Zero-Shot Compositional Action Recognition (ZS-CAR), that investigates the compositional generalization in the video modality. A benchmarking dataset Sth-com and a baseline method C2C are proposed along with the novel task. Besides, an enhanced training strategy is presented to address two inherent challenges in ZS-CAR: component domain variations and component compatibility variations. The approach demonstrates superior performance compared to existing compositional generalization solutions in both image- and video-based zero-shot compositional learning settings.

## Acknowledgements

This work is supported in part by the National Key Research and Development Program of China (2023YFF1105102, 2023YFF1105105), the National Natural Science Foundation of China (62020106012, 62332008, 62106089, U1836218, 62336004), the 111 Project of Ministry of Education of China (B12018), the Post-graduate Research & Practice Innovation Program of Jiangsu Province (KYCX22\_2307), and the UK EPSRC (EP/V002856/1, EP/T022205/1).

## References

1. Arnab, A., Deghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6836–6846 (2021)
2. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: International Conference on Machine Learning. pp. 813–824 (2021)
3. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. *Transactions of the association for computational linguistics* **5**, 135–146 (2017)
4. Brattoli, B., Tighe, J., Zhdanov, F., Perona, P., Chalupka, K.: Rethinking zero-shot video classification: End-to-end training for realistic applications. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4613–4623 (2020)
5. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)
6. Chao, W.L., Changpinyo, S., Gong, B., Sha, F.: An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In: Proceedings of the European Conference on Computer Vision. pp. 52–68 (2016)
7. Chen, S., Huang, D.: Elaborative rehearsal for zero-shot action recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13638–13647 (2021)
8. Chomsky, N.: *Aspects of the Theory of Syntax*. No. 11, MIT press (2014)
9. Doshi, K., Yilmaz, Y.: Zero-shot action recognition with transformer-based video semantic embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4858–4867 (2023)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Deghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
11. Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., Feichtenhofer, C.: Multiscale vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6824–6835 (2021)
12. Feichtenhofer, C.: X3d: Expanding architectures for efficient video recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 203–213 (2020)

13. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6202–6211 (2019)
14. Fodor, J.A., Pylyshyn, Z.W.: Connectionism and cognitive architecture: A critical analysis. *Cognition* **28**(1-2), 3–71 (1988)
15. Fu, Y., Hospedales, T.M., Xiang, T., Gong, S.: Attribute learning for understanding unstructured social activity. In: Proceedings of the European Conference on Computer Vision. pp. 530–543 (2012)
16. Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., et al.: The "something something" video database for learning and evaluating visual common sense. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5842–5850 (2017)
17. Gretton, A., Bousquet, O., Smola, A., Schölkopf, B.: Measuring statistical dependence with hilbert-schmidt norms. In: International Conference on Algorithmic Learning Theory. pp. 63–77 (2005)
18. Hao, S., Han, K., Wong, K.Y.K.: Learning attention as disentangler for compositional zero-shot learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15315–15324 (2023)
19. Hara, K., Kataoka, H., Satoh, Y.: Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6546–6555 (2018)
20. Heckman, J.J.: Sample selection bias as a specification error. *Econometrica: Journal of the econometric society* pp. 153–161 (1979)
21. Isola, P., Lim, J.J., Adelson, E.H.: Discovering states and transformations in image collections. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1383–1391 (2015)
22. Ji, J., Krishna, R., Fei-Fei, L., Niebles, J.C.: Action genome: Actions as compositions of spatio-temporal scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10236–10247 (2020)
23. Ji, S., Xu, W., Yang, M., Yu, K.: 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**(1), 221–231 (2012)
24. Karthik, S., Mancini, M., Akata, Z.: Kg-sp: Knowledge guided simple primitives for open world compositional zero-shot learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9336–9345 (2022)
25. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950* (2017)
26. Kim, H., Lee, J., Park, S., Sohn, K.: Hierarchical visual primitive experts for compositional zero-shot learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5675–5685 (2023)
27. Kwon, H., Kim, M., Kwak, S., Cho, M.: Motionsqueeze: Neural motion feature learning for video understanding. In: Proceedings of the European Conference on Computer Vision. pp. 345–362 (2020)
28. Kwon, H., Kim, M., Kwak, S., Cho, M.: Learning self-similarity in space and time as generalized motion for video action recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13065–13075 (2021)
29. Li, K., Wang, Y., Gao, P., Song, G., Liu, Y., Li, H., Qiao, Y.: Uniformer: Unified transformer for efficient spatiotemporal representation learning. *arXiv preprint arXiv:2201.04676* (2022)



30. Li, R.C., Wu, X.J., Wu, C., Xu, T.Y., Kittler, J.: Dynamic information enhancement for video classification. *Image and Vision Computing* **114**, 104244 (2021)
31. Li, R., Wu, X.J., Xu, T.: Video is graph: Structured graph module for video action recognition. arXiv preprint arXiv:2110.05904 (2021)
32. Li, R., Xu, T., Wu, X.J., Shen, Z., Kittler, J.: Perceiving actions via temporal video frame pairs. *ACM Transactions on Intelligent Systems and Technology* **15**(3), 1–20 (2024)
33. Li, X., Yang, X., Wei, K., Deng, C., Yang, M.: Siamese contrastive embedding network for compositional zero-shot learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9326–9335 (2022)
34. Li, Y., Ji, B., Shi, X., Zhang, J., Kang, B., Wang, L.: Tea: Temporal excitation and aggregation for action recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 909–918 (2020)
35. Lin, J., Gan, C., Han, S.: Tsm: Temporal shift module for efficient video understanding. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7083–7093 (2019)
36. Liu, J., Kuipers, B., Savarese, S.: Recognizing human actions by attributes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3337–3344 (2011)
37. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3202–3211 (2022)
38. Lu, X., Guo, S., Liu, Z., Guo, J.: Decomposed soft prompt guided fusion enhancing for compositional zero-shot learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 23560–23569 (2023)
39. Ma, L., Zheng, Y., Zhang, Z., Yao, Y., Fan, X., Ye, Q.: Motion stimulation for compositional action recognition. *IEEE Transactions on Circuits and Systems for Video Technology* (2022)
40. Mancini, M., Naeem, M.F., Xian, Y., Akata, Z.: Open world compositional zero-shot learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5222–5230 (2021)
41. Mandal, D., Narayan, S., Dwivedi, S.K., Gupta, V., Ahmed, S., Khan, F.S., Shao, L.: Out-of-distribution detection for generalized zero-shot action recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9985–9993 (2019)
42. Materzynska, J., Xiao, T., Herzig, R., Xu, H., Wang, X., Darrell, T.: Something-else: Compositional action recognition with spatial-temporal interaction networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1049–1059 (2020)
43. Misra, I., Gupta, A., Hebert, M.: From red wine to red tomato: Composition with context. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1792–1801 (2017)
44. Naeem, M.F., Xian, Y., Tombari, F., Akata, Z.: Learning graph embeddings for compositional zero-shot learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 953–962 (2021)
45. Nayak, N.V., Yu, P., Bach, S.: Learning to compose soft prompts for compositional zero-shot learning. In: *The Eleventh International Conference on Learning Representations* (2022)
46. Neimark, D., Bar, O., Zohar, M., Asselmann, D.: Video transformer network. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3163–3172 (2021)

47. Pan, J., Lin, Z., Zhu, X., Shao, J., Li, H.: St-adapter: Parameter-efficient image-to-video transfer learning. *Advances in Neural Information Processing Systems* **35**, 26462–26477 (2022)
48. Park, J., Lee, J., Sohn, K.: Dual-path adaptation from image to video transformers. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2203–2213 (2023)
49. Piergiovanni, A., Kuo, W., Angelova, A.: Rethinking video vits: Sparse video tubes for joint image and video learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2214–2224 (2023)
50. Purushwalkam, S., Nickel, M., Gupta, A., Ranzato, M.: Task-driven modular networks for zero-shot compositional learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3593–3602 (2019)
51. Qian, R., Lin, W., See, J., Li, D.: Controllable augmentations for video representation learning. *Visual Intelligence* **2**(1), 1 (2024)
52. Qian, Y., Yu, L., Liu, W., Hauptmann, A.G.: Rethinking zero-shot action recognition: Learning from latent atomic actions. In: *Proceedings of the European Conference on Computer Vision*. pp. 104–120 (2022)
53. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PMLR (2021)
54. Ruis, F., Burghouts, G., Bucur, D.: Independent prototype propagation for zero-shot compositionality. *Advances in Neural Information Processing Systems* **34**, 10641–10653 (2021)
55. Saini, N., Pham, K., Shrivastava, A.: Disentangling visual embeddings for attributes and objects. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13658–13667 (2022)
56. Sun, P., Wu, B., Li, X., Li, W., Duan, L., Gan, C.: Counterfactual debiasing inference for compositional action recognition. In: *Proceedings of the ACM International Conference on Multimedia*. pp. 3220–3228 (2021)
57. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4489–4497 (2015)
58. Wang, L., Tong, Z., Ji, B., Wu, G.: Tdn: Temporal difference networks for efficient action recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1895–1904 (2021)
59. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Van Gool, L.: Temporal segment networks: Towards good practices for deep action recognition. In: *Proceedings of the European Conference on Computer Vision*. pp. 20–36 (2016)
60. Wang, Q., Liu, L., Jing, C., Chen, H., Liang, G., Wang, P., Shen, C.: Learning conditional attributes for compositional zero-shot learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11197–11206 (2023)
61. Xie, S., Sun, C., Huang, J., Tu, Z., Murphy, K.: Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In: *Proceedings of the European Conference on Computer Vision*. pp. 305–321 (2018)
62. Xu, T., Zhu, X.F., Wu, X.J.: Learning spatio-temporal discriminative model for affine subspace based visual object tracking. *Visual Intelligence* **1**(1), 4 (2023)
63. Xu, X., Hospedales, T., Gong, S.: Transductive zero-shot action recognition by word-vector embedding. *International Journal of Computer Vision* **123**, 309–333 (2017)

64. Yan, R., Huang, P., Shu, X., Zhang, J., Pan, Y., Tang, J.: Look less think more: Rethinking compositional action recognition. In: Proceedings of the ACM International Conference on Multimedia. pp. 3666–3675 (2022)
65. Yan, R., Xie, L., Shu, X., Zhang, L., Tang, J.: Progressive instance-aware feature learning for compositional action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
66. Yang, T., Zhu, Y., Xie, Y., Zhang, A., Chen, C., Li, M.: Aim: Adapting image models for efficient video action recognition. arXiv preprint arXiv:2302.03024 (2023)
67. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6023–6032 (2019)
68. Zadrozny, B.: Learning and evaluating classifiers under sample selection bias. In: Proceedings of the International Conference on Machine Learning. p. 114 (2004)
69. Zellers, R., Choi, Y.: Zero-shot activity recognition with verb attribute induction. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing. pp. 946–958 (2017)
70. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. *International Journal of Computer Vision* **130**(9), 2337–2348 (2022)