

Telling Stories for Common Sense Zero-Shot Action Recognition

Shreyank N Gowda¹ and Laura Sevilla-Lara²

¹ University of Nottingham, UK

² University of Edinburgh, UK

Shreyank.Narayanagowda@nottingham.ac.uk

Abstract. Video understanding has long suffered from reliance on large labeled datasets, motivating research into zero-shot learning. Recent progress in language modeling presents opportunities to advance zero-shot video analysis, but constructing an effective semantic space relating action classes remains challenging. We address this by introducing a novel dataset, *Stories*, which contains rich textual descriptions for diverse action classes extracted from WikiHow articles. For each class, we extract multi-sentence narratives detailing the necessary steps, scenes, objects, and verbs that characterize the action. This contextual data enables modeling of nuanced relationships between actions, paving the way for zero-shot transfer. We also propose an approach that harnesses *Stories* to improve feature generation for training zero-shot classification. Without any target dataset fine-tuning, our method achieves new state-of-the-art on multiple benchmarks, improving top-1 accuracy by up to 6.1%. We believe *Stories* provides a valuable resource that can catalyze progress in zero-shot action recognition. The textual narratives forge connections between seen and unseen classes, overcoming the bottleneck of labeled data that has long impeded advancements in this exciting domain. The data can be found here: <https://github.com/kini5gowda/Stories>.

Keywords: Zero-shot Action Recognition · Action Descriptions · Feature Generation

1 Introduction

Action recognition technology has improved remarkably over the years, with methods becoming more accurate. However, one of the main challenges that remain today is the dependency of these methods on annotated data for novel categories. The availability of large labeled datasets like ImageNet [10] has played a pivotal role in propelling the field of supervised learning seeing better than human level performance on image classification [11, 26, 28] tasks. In practice, obtaining large amounts of annotated examples for each new class we aim to recognize is not realistic and not practical [21]. This is particularly true as we grow the number of classes and wish to incorporate more flexible, natural language, for example, for retrieval. This general problem has led to research in the zero-shot (ZS) domain.




Fig. 1: Comparison of accuracy across state-of-the-art ZS approaches using different semantic embeddings: the proposed *Stories*, word2vec (W2V) and elaborative definitions (ER), on UCF101. Using the proposed *Stories* to create the semantic space of class labels improves the performance by a large margin across all methods, showing that it is model-agnostic.

In the typical ZS setting, there are seen classes that contain visual examples and their class label, and there are unseen classes where only the class labels are available at training time. Given a visual sample of the unseen set at test time, the task is to output the corresponding class label. Approaches typically learn a mapping from the visual space to the class labels using the seen classes and leverage that mapping in different ways to approximate the mapping in the unseen classes. Some examples of solutions include learning to map visual information and class labels to the same space or learning to generate visual features using a generative adversarial network (GAN) [52].

One of the underlying assumptions of these approaches is that the distances between the data points are meaningful both in the visual space and the semantic space. In other words, data points that are close together should be similar in content across both seen and unseen classes. In visual space, this tends to happen naturally as related classes will share objects, scenes, etc. For example, if we compare classes such as “penalty shot” and “playing soccer”, they will share the ball, the soccer field, etc. However, in the space of action labels, which we also refer to as semantic space, this property is not straightforward to achieve. While some similar classes will contain overlapping terms (such as “horse-back riding” and “horse racing”), some others might be similar but not contain overlapping words (such as the previous example of “penalty shot” and “playing soccer”). This makes the step of transferring knowledge between seen and unseen classes harder. Previous efforts to improve the semantic space of class labels have included the use of manually annotated attributes or embedding functions trained on language corpora, such as word2vec [39], sentence2vec [43] and using definitions of actions [9].

In this work, we address the problem of building a meaningful space of action labels by leveraging the story around each action. In particular, we use the

descriptions of the steps needed to achieve each action, which are the *Stories* around this action and encode them using a language embedding. These steps typically contain the objects, tools, scenes, verbs, adjectives, etc., associated with the action label. One could think of all these additional pieces of information around an action as the “common sense” of associations humans would typically consider. For example, in the case of the penalty shot, the steps would describe to first place the ball, run the hand through the grass, fluff the ball, take some steps back, kick, etc. When playing soccer, the steps include kicking the ball with the inside of your shoe for short passes across the grass, tapping the ball from foot to foot, etc. When we compare the stories of steps around these two classes, the overlap of terms becomes much more obvious. It is more likely that these related classes are closer in semantic space, facilitating the transfer of knowledge between seen and unseen classes. We show that this relatively simple approach to creating a semantic space is extremely effective across datasets and methods, improving performance by up to 20% compared to the standard word2vec [39]. Figure 1 shows that all state-of-the-art methods improve significantly and therefore the proposed semantic embedding is general and model-agnostic.

Finally, we leverage these Stories to go beyond simply learning semantic representations, to actually generating additional features that improve the semantic space further. We follow a feature-generating approach [37, 52] using a GAN [19] to synthesize visual data points from these semantic embeddings. These synthetic visual data points are then used to learn a mapping from visual to semantic space in the unseen classes. We show that this method improves state-of-the-art by an additional 6.1%. We observe a strong trend of feature generating networks benefiting particularly from using Stories. The largest jumps using Stories are in feature generating methods (SDR-I3D, SDR-CLIP and OD), as shown in Figure 1.

2 Related Work

Fully Supervised Action Recognition. In this setting, there is a large amount of training samples with their associated labels and the label spaces are the same at train and test time. Much of the advances in this area of research is often used on ZS. Early work in deep learning for action recognition used many tools to represent the spatio-temporal nature of videos, including 3D CNNs [7], 2D CNNs with temporal modules [29, 33], 2D CNNs with relational modules [23] and two-stream networks [49]. More recently, the transformer architecture [11] is particularly well suited to represent sequential information and has been successfully adapted to the video domain [2, 4, 13, 35]. While these are extremely powerful tools, they are difficult to train with limited data. Instead, we use the standard I3D [7] as our backbone feature generator to compare directly with recent state-of-the-art papers [24, 37].

Zero-Shot Action Recognition. Early work [48] in this setting used script data in cooking activities to transfer to unseen classes. Considering each action class as a

Nearest Neighbors using Elaborate Descriptions [7]		Nearest Neighbors using Stories	
Hammer Throw a tool with a heavy metal head mounted at right angles at the end of a handle, used for jobs such as breaking things and driving in nails. propel (something) with force through the air by a movement of the arm and hand.	Hammering hit or beat (something) repeatedly with a hammer or similar object.	Hammer Throw Hammer Throw is a popular field and track event in which the athleteby spinning it in air with the help of a wire or string attached to it. At first, you have to stand at the rear point of the circle ...	Shot Put The shot put is a track and field event involving "putting" (throwing) a heavy spherical ball—the shot—as far as possible.... The athlete ... When shot putting, you must stay within the circle during the entire throw...

Fig. 2: Comparing nearest neighbors using Stories. We see an example where ER fails and Stories provides more context and helps in obtaining better neighbors. This is one example of where ER fails, there are multiple such examples. Dataset is UCF101.

domain, Gan et al. [17] address the identification of semantic representations as a multisource domain generalization problem. To obtain semantic embeddings of class names, label embeddings such as word2vec [39] has proven popular as only class names are needed. Some approaches use a common embedding space between video features and class labels [53, 54], error-correcting codes [46], pairwise relationships between classes [15], interclass relationships [16], out-of-distribution detectors [37], and graph neural networks [18]. Recently, it was seen that clustering of joint visual-semantic features helped obtain better representations [24] for ZS action recognition. Similar to CLUSTER, ReST [32] jointly encodes video data and text labels for ZS action recognition. In ResT, transformers are used to perform modality-specific attention. JigSawNet [45] also models visual and textual features jointly but decomposes videos into atomic actions in an unsupervised manner and bridges group-to-group relationships between visual and semantic representations instead of the one-to-one relationships that CLUSTER and ReST do. Unlike JigSawNet, that works on the visual features, we create enriched textual descriptions by decomposing actions into a series of steps and hence obtain richer semantic features. Other solutions to deal with limited labeled data include augmentation [22], relational modules [44].

Semantic Embeddings. To obtain semantic embeddings of action class labels, earlier works use word2vec [39] directly on the class labels. However, word2vec averages the embedding for class labels with multiple words, giving equal weight to each word. This causes class names to lose context. For example, the class “pommel horse” is a gymnastic class which does not involve the animal horse. Unfortunately, using word2vec makes the embedding of that class close to “horse riding” or “horse racing” in the word2vec space, even though they are not gymnastics sports. A recent solution is to use elaborate descriptions [9] based on the principle of Elaborative Rehearsals (ER), which replace each class name with a class definition. An object description was also used to describe the objects in the particular action. This resulted in a significant boost in performance. Still, ER uses descriptions of each word in a class label independently to create the semantic space, which potentially leads to errors.




Fig. 3: Visualization of the features generated from the *embedding* vs ER, using t-SNE. We observe that the samples of each class instance, depicted in a single color, are better clustered together, pointing to a more semantically meaningful space.

Feature Generating Networks for ZS. Bucher et al. [6] proposed to bridge the gap between seen and unseen classes by synthesizing features for the unseen classes using four different generator models. Xian et al. [52] trained their generators by using a conditional generative adversarial network (GAN) [40]. In contrast, Verma et al. [51] trained a variational autoencoder. Mishra et al. [41] introduced the generative approach to the domain of ZS action recognition. Their approach models each action class as a probability distribution in the visual space. Mandal et al. [37] modified the work by Xian et al. [52] to work directly on video features. They also introduced an out-of-distribution detector to help with the generalized zero-shot learning setting. Here we propose a variant of the work on out-of-distribution detectors, which suffers less from long converging times yet it improves its accuracy. More recently Gowda et al. [20], proposed to have a selector that selects generated features based on their importance on training the classifier rather than on realness of the features.

3 The Stories Dataset

Research in semantic representation of action labels has shown over the years that more sophisticated representations help to build a meaningful semantic space for ZS action recognition. Here we go beyond previous work by representing not only the class label but the story around it. This is, we build a representation that captures all the steps needed to perform the action, which include objects, verbs, etc, typically associated with that action. We call these representations *Stories*, and we now describe how we build them.

3.1 Building the Stories Dataset

We leverage textual descriptions of actions from *WikiHow*³, a website that gives instructions on how to perform actions. These instructions consist of long paragraphs that describe each step in completing the action. For example, for the action classes in the HMDB51 dataset, the *WikiHow* articles contain an average

³ <https://www.wikihow.com/>

of 9.8 steps, ranging from 4 to 20 steps. The most closely related work to us, ER [9] uses a single or at max two sentences per class. Instead, for the proposed Stories, the average number of lines is 14.4 for the classes in the Olympics dataset, 9.6 for HMDB51, 13.2 for UCF101, and 13.5 Kinetics. These rich descriptions inherently contain information of the objects needed to perform an action. For example, the class ‘biking’ has a paragraph that explains where to bike, what equipment you need, and the steps to do it. In comparison, ED [9] has only a single line definition of the class.

Not all classes have articles on *WikiHow*. For example, if we search for “draw sword” (a class in UCF101), we will get instructions on how to paint or sketch swords instead of the steps needed to essentially remove a sword from its sheath. Hence, collecting clean, meaningful articles requires a more complex process than a simple search. After scraping the articles from *WikiHow* corresponding to all classes, we use sentence-BERT encoders [47] to represent the sentences in the article and use cosine similarity to find the 25 that are the most similar to the class definition [9].

Next, we manually check the sentences for each class. If we find a mismatch between the article and the action class, as was the case with the “draw sword” example, we do a manual search to pick the most relevant article from other sources such as *Wikipedia*. However, these alternative articles do not tend to contain the sequence of steps and hence need more manual intervention to order the sentences into a sequence of steps.

We finally clean each story by re-arranging the sentences in sequential order and removing irrelevant sentences. In total, we had 6 people who manually cleaned the descriptions after the initial stage of noisy collection and a further 10 who verified the descriptions. This was done using the Prolific⁴ platform. The time taken for the cleaning Stories for UCF101, HMDB51 and Kinetics was 7.2 hours, 3.3 hours and 25.3 hours on average respectively. We followed this process to create a dataset of these textual representations for classes in UCF101, HMDB51, Olympics, and Kinetics-400, as they are the most commonly used datasets for ZS action recognition.

3.2 Learning from Stories

In order to test the impact of using richer semantic representations, we used Stories as input to the state-of-the-art methods in ZS. To provide fair points of comparison, we also supply the standard word2vec [39] embeddings as well as the more recently introduced ER [9] embeddings to these same models. By visually depicting the performance gains achieved by the models when using our proposed embeddings instead of the word2vec or ER embeddings in Fig. 1, we are able to clearly demonstrate the substantial improvements obtained through our approach. It increases up to 21% compared to the widely used word2vec embeddings and 11.8% over the ER embeddings. This comprehensive experimental analysis highlights how our semantically richer embeddings can notably boost

⁴ <https://www.prolific.co/>

performance across a diverse range of state-of-the-art models and is hence model agnostic.

We visualize the effect of using ER and Stories to generate features, using the t-SNE [36] for 10 classes in UCF101 (Fig. 3), all related to gymnastics and therefore easier to confuse. We see that using Stories helps keep a more meaningful neighborhood for visual instances, and keep classes apart. The visualization also shows why ER fails as there is not enough information to clearly distinguish classes.

3.3 Why Are Stories Necessary?

The proposed Stories dataset produces semantic embeddings for the class labels that are much more meaningful than previous work, and that reflects in the experimental improvement shown in Fig. 1. Here we delve deeper into what properties make Stories superior.

Capturing meaning of words jointly. Previous work [9, 39] represents a class label by computing a representation for each word in the class label and then computing the average.⁵ While this in general is a sensible choice, it leads often to errors caused by words that have multiple meanings depending on the context. Fig. 2 illustrates this issue with an example. We use the class “Hammer Throw” from the UCF101 [50] dataset, which is a sporting event in which the athlete throws a spherical object. If we retrieve the nearest neighbor with ER [9], we obtain the class “Hammering”, which is not actually related in meaning but both contain the description of the tool hammer. However, if we use Stories, the nearest neighbor is “Shot Put”, which is also a sporting event where the athlete throws a spherical object. Similar problems happen for example with classes such as “Sword Exercise”, or “Swing Baseball”. Overall, this shows the need for the more sophisticated joint description that Stories provides, instead of the per-word definition of ER.

Size of dataset. We also compare Stories to ER from a statistical perspective. We have briefly mentioned before that Stories contains much more detailed descriptions of the classes. Here we look at the numerical difference between ER and Stories, shown in Table 1. The number of sentences is one order of magnitude larger, going from over 1 sentence on average per class, to over 10 sentences on average, depending on the dataset. This ratio is also consistent in the number of nouns verbs, etc.

Diversity of dataset. Another aspect that increases the specificity of the class descriptions in addition to the size of the dataset is the diversity of the vocabulary. We look at the number of unique words in Table 1 and observe that Stories contains more unique words than ER in all datasets and that the difference is particularly remarkable in smaller datasets. We argue that this diversity

⁵ In the case of ER, this is done for some of the classes.

Method	D	N	V	A	Ad	UW	S
Stories	K	68.5	37.9	7.1	0.04	35.1	13.5
ER	K	9.2	3.8	1.0	0.01	30.2	1.8
Stories	U	69.6	37.2	7.7	0.1	32.1	13.2
ER	U	9.7	4.3	1.4	0.06	29.5	1.8
Stories	H	59.2	33.4	6.5	0.3	34.5	9.6
ER	H	5.8	2.1	0.9	0.05	17.7	1.2
Stories	O	68.9	38.6	11.8	2.0	57.2	14.4
ER	O	9.1	4.2	1.2	0.3	31.4	1.7

Table 1: Statistical comparison of Stories to ER. We observe a larger number of Nouns (N), Verbs (V), Adverbs (A), Adjectives (Ad), Unique Words (UW), and Sentences (S) are averages across all the classes in a particular dataset (D). 'K' refers to Kinetics-600, 'U' to UCF101, 'H' to HMDB51, and 'O' to Olympics.

contributes to representing each class label in a more unique way, leading to a more sparse yet meaningful space.

Cleaning data manually. Generally, training with more data tends to produce better results. There often is a tension between using a smaller amount of clean data or a larger amount of noisy data. Here we have explored the effect of cleaning the data of Stories manually. We see that using the noisy version of the dataset (see Supplementary for more details) improves the performance over ER across methods but is still consistently worse than the cleaned version, even though it is roughly twice as large. This shows that the effect of cleaning up the data manually is not trivial. In the ER work, however, there exist errors that can be solved through manual revision. For example, "Table Tennis Shot" has the definition "put (food) into the mouth and chew and swallow it" which clearly corresponds to the wrong action class.

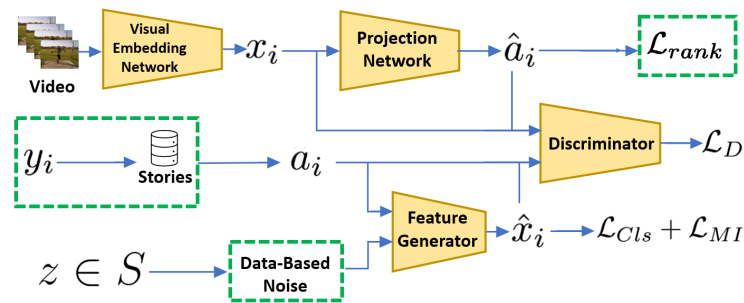


Fig. 4: Using Stories for feature generation. The elements depicted in yellow are the standard vanilla approach to feature generation for ZS. Depicted in green are the elements that we introduce.




Fig. 5: Training the generator using data-driven noise converges much faster than using the standard Gaussian noise.




Fig. 6: The generator loss using data-driven noise is much more stable, leading to faster convergence and better accuracy.

4 Experimental Details

4.1 The Zero-Shot and Generalized Zero-Shot Settings

Let S be the training set of seen classes. S is composed of tuples $(x, y, a(y))$, where x represents the spatiotemporal features of a video, y represents the class label in the set of Y_S seen class labels, and $a(y)$ denotes the category-specific semantic representation of class y , which is either manually annotated or computed automatically, for example using word2vec [39] or the proposed Stories .

Let U be the set of pairs $(u, a(u))$, where u is a class in the set of unseen classes Y_U and $a(u)$ are the corresponding semantic representations. The seen Y_S and the unseen classes Y_U do not overlap.

In the ZS setting, given an input video, the task is to predict a class label in the unseen classes, such as $f_{ZSL} : X \rightarrow Y_U$. In the generalized zero-shot (GSZ) setting, given an input video, the task is to predict a class label in the union of the seen and unseen classes, as $f_{GZSL} : X \rightarrow Y_S \cup Y_U$.

4.2 Vanilla Feature Generation

In the standard feature generation pipeline [27, 34, 37, 41, 52] the high-level idea is to learn to generate visual features for unseen classes using a GAN [19], and then, given these synthetic features, train a classifier that takes in visual features and predicts unseen class labels. Figure 4 illustrates the overall method.

The GAN comprises a generator (G), discriminator (D), and projection network (P). Generator G creates synthetic visual features (\hat{x}_i) from class label semantic embedding (a_i) and noise (z). P maps visual features (x_i) to a semantic embedding approximation (\hat{a}_i). D separates real from synthesized features. They’re jointly trained via a mini-max game using an optimization function (see Eq 1).

$$\begin{aligned}
\mathcal{L}_D = & \mathbb{E}_{(x,a) \sim p_{(x_S \times a_S)}} [D(x, P(x))] \\
& - \mathbb{E}_{z \sim p_z} \mathbb{E}_{a \sim p_a} [D(G(a, z), a)] \\
& - \alpha \mathbb{E}_{z \sim p_z} \mathbb{E}_{a \sim p_a} \left[(\|\nabla_{\hat{x}} D(G(a, z))\|_2 - 1)^2 \right],
\end{aligned} \tag{1}$$

4.3 Technical Details

We now describe a few technical choices that we have made to improve the standard pipeline. We present experimental results for some of our choices (including number of nearest neighbors, alternative feature generating models and the need for manual cleaning) and discuss implementation details in the supplementary material.

First, the typical generator takes as input attributes or semantic embeddings and normally distributed noise to generate visual features. The underlying assumption is that the normal distribution can represent all classes. However, this is not necessarily true [34].

Instead, we use the distribution of the seen classes' features to create the "noise" for the generator [34], such that the synthetic unseen classes will follow the same distribution. To this end, we use a variational auto-encoder (VAE), which takes as input visual features from the seen classes and it is trained to reconstruct them. Once trained, we use the low-dimensional representation of the encoder as the "noise". This simple change in the noise distribution, which we call data-driven noise, benefits in two ways: it improves the overall accuracy and reduces training time by 65% compared to the standard Gaussian noise. See Fig. 5 and Fig. 6.

Second, one of the risks of learning to generate semantic embeddings (through the "Projection Network" in Fig. 4) is that synthetic semantic embeddings can be too similar to each other. To avoid this, we introduce a ranking loss [14] that pushes apart the generated semantic representation (\hat{a}_i) from those of their neighboring classes. Details of this can be found in the supplementary. We refer to this version of the feature generating approach as SDR (for Stories, Data-based Noise and Ranking) and we observe that it achieves state-of-the-art results across all datasets and settings.

$p_{(x_S \times a_S)}$ is the joint distribution of visual and semantic descriptors for seen classes, p_a is the empirical distribution of their semantic embeddings, p_z is noise, and α is a penalty coefficient. Additional losses to enhance generated features are the classification regularized loss (\mathcal{L}_{CLS}) and the mutual information loss [3] (\mathcal{L}_{MI}). These losses form the objective function minimized to train the vanilla pipeline (see Eq 2).

$$\min_G \min_P \max_D \mathcal{L}_D + \lambda_1 \mathcal{L}_{CLS}(G) + \lambda_2 \mathcal{L}_{MI}(G). \tag{2}$$

Once these networks are trained on the seen classes, the generator is used to synthesize visual features for the unseen classes. The final step is to train a simple classifier using these synthetic visual features as input and the class

labels. The loss is the standard cross-entropy loss, and the classifier is a simple multilayer perceptron (MLP).

4.4 Ablation Study

Stories	DBN	\mathcal{L}_{rank}	ZSL_{HMDB}	ZSL_{UCF}	$GZSL_{HMDB}$	$GZSL_{UCF}$
×	×	×	29.1 ± 3.8	37.5 ± 3.1	32.7 ± 3.4	44.4 ± 3.0
×	×	✓	29.7 ± 3.5	38.0 ± 3.1	35.3 ± 3.1	47.1 ± 3.2
×	✓	×	30.6 ± 2.2	38.6 ± 3.4	33.3 ± 3.0	44.9 ± 2.9
✓	×	×	44.6 ± 2.9	60.4 ± 3.8	44.9 ± 3.6	51.0 ± 2.9
×	✓	✓	31.9 ± 3.2	40.9 ± 2.9	35.7 ± 2.9	47.9 ± 4.1
✓	×	✓	45.0 ± 2.5	60.9 ± 3.5	49.0 ± 3.2	54.4 ± 3.7
✓	✓	×	45.9 ± 2.7	61.4 ± 2.8	47.5 ± 2.6	53.7 ± 3.5
✓	✓	✓	46.5 ± 5.3	61.9 ± 2.5	49.7 ± 2.9	54.9 ± 4.4

Table 2: Ablation study to explore the impact of each proposed component.

We propose modifications to the vanilla pipeline of feature generation and based on this, we show the importance of each component here. We see that every proposed contribution benefits over the baseline. However, crucially, a combination of all three gives us our best results. We note that the improvement from the ranking loss is much more prominent in the GZS setting.

4.5 Datasets and Evaluation Protocol

We use the HMDB-51 [31], UCF-101 [50] and Kinetics [7] datasets, as they are the standard choice in ZS action recognition, so that we can compare them with recent state-of-the-art models [5, 17, 24, 37, 42, 45, 46]. These datasets contain 783, 6766 and 13320 videos, and have 16, 51 and 101 classes, respectively.

We follow the commonly used 50/50 splits of Xu et al. [53], where 50% are the seen classes and 50 are the unseen classes. Similar to previous approaches [17, 30, 38, 46, 55], we report average accuracy and standard deviation over 10 independent runs.

We also report on the recently introduced TruZe [25]. This split accounts for the fact that some classes present on the dataset used for pre-training (Kinetics [7]) overlap with some of the unseen classes in the datasets used in the zero-shot setting. We also report on the Kinetics-220 [7] split as proposed in ER [9]. Here, the 220 classes from Kinetics-600 [8] are treated as unseen classes to a model trained on the Kinetics-400 dataset.

5 Results

We now look at the comparison of the proposed method with recent SOTA models in the zero-shot as well as the generalized zero-shot settings. We also

report the use of *Stories* as a general purpose semantic embedding that improves multiple model’s performance in a simple plug and play fashion highlighting the benefits of a more extensive semantic embedding in zero-shot action recognition.

5.1 Zero-Shot Learning Results

D	E	Bi-Dir GAN	WGAN	OD	E2E	CLASTER	JigSawNet	ActionCLIP	ResT	X-CLIP	SDR+I3D	SDR+CLIP
O	W	40.2 ± 10.6	47.1 ± 6.4	50.5 ± 6.9	61.4 ± 5.5	63.8 ± 5.7	66.4 ± 6.8	-	-	-	55.1 ± 4.8	62.3 ± 4.5
	ER	54.1 ± 6.8	65.5 ± 7.2	67.6 ± 6.2	66.5 ± 4.5	68.4 ± 4.1	71.5 ± 6.1	-	-	-	69.8 ± 2.8	73.5 ± 2.9
	Sto	55.5 ± 6.5	66.2 ± 7.1	69.1 ± 5.6	69.9 ± 5.8	73.1 ± 6.6	74.9 ± 5.2	-	-	-	72.5 ± 2.1	80.1 ± 2.3
	SM	-	-	-	-	-	-	-	-	-	74.8 ± 2.3	82.2 ± 1.6
H	W	21.3 ± 3.2	29.1 ± 3.8	30.2 ± 2.7	33.1 ± 3.4	36.6 ± 4.6	35.4 ± 3.2	40.8 ± 5.4	39.3 ± 3.5	43.7 ± 6.5	35.8 ± 4.7	38.9 ± 3.5
	ER	25.9 ± 2.9	31.6 ± 3.1	36.1 ± 2.9	36.2 ± 1.9	43.2 ± 1.9	39.3 ± 3.9	44.2 ± 4.4	43.6 ± 2.9	46.6 ± 6.1	41.2 ± 4.3	46.2 ± 3.1
	Sto	27.2 ± 2.7	35.5 ± 2.8	39.2 ± 2.8	38.1 ± 3.6	45.5 ± 2.6	42.5 ± 3.2	48.8 ± 3.2	47.1 ± 3.5	50.1 ± 6.1	46.8 ± 5.0	52.7 ± 3.4
	SM	-	-	-	-	-	-	-	-	-	48.9 ± 4.4	54.4 ± 4.1
U	W	21.8 ± 3.6	25.8 ± 3.2	26.9 ± 2.8	46.2 ± 3.8	46.7 ± 5.4	53.3 ± 3.1	58.3 ± 3.4	58.7 ± 3.3	70.1 ± 3.4	41.5 ± 2.5	44.8 ± 4.2
	ER	28.0 ± 3.4	37.9 ± 2.5	42.4 ± 3.4	52.4 ± 3.3	53.9 ± 2.5	56.8 ± 2.8	64.3 ± 3.8	62.6 ± 4.1	72.2 ± 2.3	50.3 ± 1.1	54.2 ± 3.5
	Sto	29.5 ± 3.2	40.1 ± 3.7	50.3 ± 3.0	55.1 ± 3.3	59.6 ± 2.8	61.6 ± 3.5	71.8 ± 2.7	65.3 ± 2.5	73.8 ± 2.9	62.9 ± 1.6	73.4 ± 2.7
	SM	-	-	-	-	-	-	-	-	-	64.9 ± 2.1	75.5 ± 3.2

Table 3: Results on ZSL. E: semantic embedding, W: word2vec embedding, ER: Elaborate Rehearsals, Sto: *Stories*. SM corresponds to the single model training. We use the datasets Olympics (O), HMDB51 (H) and UCF101 (U).

We look at the use of *Stories* as semantic embedding to a wide variety of models here, from older models all the way to the most recent ones in the ZS literature. We use I3D and CLIP-based features [42] to compare their effect. We list these results as SDR+I3D and SDR+CLIP respectively. We observe an improvement across all of them and across all datasets, demonstrating that *Stories* is clearly model agnostic. Results can be seen in Table 3.

We also observe that the proposed changes to the vanilla feature generation method which we call SDR consistently outperforms all approaches across all datasets, achieving a new state-of-the-art. We experiment with using a single model for all datasets, by training on Kinetics and not doing any fine-tuning for the smaller datasets. This is the last row of the table, which we call “SDR (Ours) + SM”. It is remarkable and quite promising that, without the need to fine-tune, this single model achieves even better performance.

We also evaluate on the Kinetics-220 dataset as proposed in ER [9]. There are fewer methods who report on this split, but it is interesting as it is much larger. Results can be seen in Table 4. We observe that the proposed SDR outperforms all previous work. We see significant gains of up to 5%.

Finally, we evaluate on the stricter TruZe [25] split that ensures no overlap between pre-trained model and test classes. Results are shown in Table 5. We report the mean class accuracy in ZS setting and harmonic mean of seen and unseen class accuracies in GZS. The split refers to the train/test split used.

5.2 Generalized Zero-Shot Learning Results

Generalized ZS action recognition is less explored in comparison to the ZS setting. Nonetheless, we choose models such as Bi-Dir GAN, GGM, WGAN, OD

Method	Top-1 Acc	Top-5 Acc
DEWISE [14]	23.8 \pm 0.3	51.0 \pm 0.6
SJE [1]	22.3 \pm 0.6	48.2 \pm 0.4
ER [9]	42.1 \pm 1.4	73.1 \pm 0.3
JigSawNet [45]	45.9 \pm 1.6	78.8 \pm 1.0
SDR+I3D	50.8 \pm 1.9	82.9 \pm 1.3
SDR+CLIP	55.1 \pm 2.2	86.1 \pm 3.1

Table 4: Results of ZS in Kinetics-220.

Method	UCF101			HMDB51		
	Split	ZSL	GZSL	Split	ZSL	GZSL
WGAN	67/34	22.5	36.3	29/22	21.1	31.8
OD	67/34	22.9	42.4	29/22	21.7	35.5
CLASTER	67/34	45.8	47.3	29/22	33.2	44.5
SDR+I3D	67/34	49.7	51.3	29/22	34.9	45.5
SDR+CLIP	67/34	53.9	56.2	29/22	38.7	49.5
VCAP [12]	0/34	49.1	-	0/22	20.4	-
SDR+I3D SM	0/34	51.5	52.2	0/22	36.1	46.6
SDR+CLIP SM	0/34	55.1	57.7	0/22	40.8	51.1

Table 5: Results of ZS and GZS on TruZe.

and CLASTER as recent state-of-the-art methods evaluating on this setting. We train an out-of-distribution detector (OOD) following OD [37] and two separate classifiers for the seen and unseen classes along with the OOD network.

Table 6 shows the results, with the harmonic mean of the seen and unseen class accuracies. Similar to the zero-shot case, we use both I3D and CLIP-based [42] backbones and list these results as SDR+I3D and SDR+CLIP respectively. We follow the earlier approach of using different semantic embeddings to show the performance gain that using SDR gives us. We provide a detailed seen and unseen class accuracy comparison in the supplementary material.

6 Discussion

6.1 Why Does Using a Single Model Work?

One curious question to ponder would be why the single model trained on a large dataset like Kinetics-400 [7] results in better performance than the models fine-tuned on the smaller datasets. Our hypothesis is that the feature generator trained on a larger dataset has a better distribution of data to learn from as the data-driven noise that we use is more representative of the real visual world. This in turn generates more realistically distributed features, which in turn results in the improved performance.

Dataset	E	Bi-Dir GAN	GGM	WGAN	OD	CLASTER	SDR+I3D	SDR+CLIP
Olympics	W	44.2 ± 11.2	52.4 ± 12.2	59.9 ± 5.3	66.2 ± 6.3	69.1 ± 5.4	67.1 ± 4.2	69.7 ± 4.1
	ER	53.6 ± 6.2	59.1 ± 12.1	63.7 ± 6.6	69.7 ± 6.5	72.5 ± 3.5	70.4 ± 2.3	75.5 ± 3.3
	Sto	55.9 ± 4.2	59.9 ± 11.6	67.1 ± 5.1	72.4 ± 4.9	74.9 ± 6.1	74.5 ± 3.9	79.5 ± 3.5
	SM	-	-	-	-	-	76.6 ± 3.6	81.1 ± 3.0
HMDB51	W	17.5 ± 2.4	20.1 ± 2.1	32.7 ± 3.4	36.1 ± 2.2	48.0 ± 2.4	36.3 ± 5.1	39.5 ± 3.1
	ER	26.1 ± 2.3	28.2 ± 3.3	35.2 ± 3.5	38.1 ± 2.4	50.8 ± 2.8	42.1 ± 4.5	45.4 ± 3.6
	Sto	27.7 ± 2.1	29.1 ± 2.9	38.5 ± 3.3	40.9 ± 3.8	53.5 ± 2.4	49.7 ± 2.9	53.5 ± 3.3
	SM	-	-	-	-	-	50.9 ± 2.6	56.1 ± 3.2
UCF101	W	22.7 ± 2.5	23.7 ± 1.2	44.4 ± 3.0	49.4 ± 2.4	51.3 ± 3.5	41.9 ± 2.7	44.3 ± 4.6
	ER	26.2 ± 4.2	26.5 ± 2.5	46.1 ± 3.5	53.2 ± 3.1	52.8 ± 2.1	45.5 ± 1.1	49.1 ± 2.9
	Sto	29.1 ± 3.4	27.7 ± 3.1	48.3 ± 3.2	55.5 ± 3.3	54.1 ± 3.3	54.9 ± 4.4	57.8 ± 4.1
	SM	-	-	-	-	-	57.2 ± 3.5	59.7 ± 3.1

Table 6: Results on GZSL. E: semantic embedding, W: word2vec embedding, ER: Elaborate Rehearsals, Sto: Stories. SM corresponds to the single model training.

6.2 Possible Limitations

One possible limitation of our approach is that the Stories may focus on one specific way of performing each action, while other valid methods may exist to do the same action. For example, the story for "shuffling cards" details the riffle shuffling technique, but other shuffling techniques such as the overhand shuffling could occur in videos of this class. Similarly, some parts of the stories may describe non-visual aspects like memorizing lines for the "acting in play" class that are not depicted in the videos. Our current approach does not explicitly address potential mismatches between the textual stories and visual contents.

Still, we believe non-visual cues actually help making the semantic embeddings more distinct, as these cues are unique to each class's story. However, this remains a limitation worth noting. One area of future work is exploring how to make the stories more comprehensive by incorporating multiple variations of actions. We also plan to investigate techniques for identifying and excluding non-visual sentences that do not translate to visual features. Overall, handling the diversity of real videos compared to procedural descriptions remains an open challenge that we aim to address.

7 Conclusion

The introduction of the novel Stories dataset provides rich textual narratives that establish connections between diverse action classes. Leveraging this contextual data enables modeling of nuanced relationships between actions, overcoming previous limitations in ZS action recognition. Stories enables significant improvement for multiple SOTA models. Building on Stories, our proposed feature generating approach harnessing Stories achieves new state-of-the-art performance on multiple benchmarks without any target dataset fine-tuning. This demonstrates the value of Stories as a resource to enable ZS transfer and significant progress in video understanding without reliance on large labeled datasets.

References

1. Akata, Z., Reed, S., Walter, D., Lee, H., Schiele, B.: Evaluation of output embeddings for fine-grained image classification. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2927–2936 (2015) [13](#)
2. Arnab, A., Dehghani, 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) [3](#)
3. Belghazi, M.I., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., Hjelm, D.: Mutual information neural estimation. In: International conference on machine learning. pp. 531–540. PMLR (2018) [10](#)
4. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: ICML. vol. 2, p. 4 (2021) [3](#)
5. 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) [11](#)
6. Bucher, M., Herbin, S., Jurie, F.: Generating visual representations for zero-shot classification. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 2666–2673 (2017) [5](#)
7. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: IEEE Conf. Comput. Vis. Pattern Recog. (2017) [3](#), [11](#), [13](#)
8. Carreira, J., Noland, E., Banki-Horvath, A., Hillier, C., Zisserman, A.: A short note about kinetics-600. arXiv preprint arXiv:1808.01340 (2018) [11](#)
9. 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) [2](#), [4](#), [6](#), [7](#), [11](#), [12](#), [13](#)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009) [1](#)
11. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020) [1](#), [3](#)
12. Estevam, V., Laroca, R., Menotti, D., Pedrini, H.: Tell me what you see: A zero-shot action recognition method based on natural language descriptions. arXiv preprint arXiv:2112.09976 (2021) [13](#)
13. 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) [3](#)
14. Frome, A., Corrado, G.S., Shlens, J., Bengio, S., Dean, J., Ranzato, M., Mikolov, T.: Devise: A deep visual-semantic embedding model. Advances in neural information processing systems **26** (2013) [10](#), [13](#)
15. Gan, C., Lin, M., Yang, Y., De Melo, G., Hauptmann, A.G.: Concepts not alone: Exploring pairwise relationships for zero-shot video activity recognition. In: Thirtieth AAAI conference on artificial intelligence (2016) [4](#)
16. Gan, C., Lin, M., Yang, Y., Zhuang, Y., Hauptmann, A.G.: Exploring semantic inter-class relationships (sir) for zero-shot action recognition. In: Proceedings of the National Conference on Artificial Intelligence (2015) [4](#)

17. Gan, C., Yang, T., Gong, B.: Learning attributes equals multi-source domain generalization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 87–97 (2016) [4](#), [11](#)
18. Gao, J., Zhang, T., Xu, C.: I know the relationships: Zero-shot action recognition via two-stream graph convolutional networks and knowledge graphs. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8303–8311 (2019) [4](#)
19. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. Communications of the ACM **63**(11), 139–144 (2020) [3](#), [9](#)
20. Gowda, S.N.: Synthetic sample selection for generalized zero-shot learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 58–67 (2023) [5](#)
21. Gowda, S.N., Hao, X., Li, G., Sevilla-Lara, L., Gowda, S.N.: Watt for what: Rethinking deep learning’s energy-performance relationship. arXiv preprint arXiv:2310.06522 (2023) [1](#)
22. Gowda, S.N., Rohrbach, M., Keller, F., Sevilla-Lara, L.: Learn2augment: Learning to composite videos for data augmentation in action recognition. In: European Conference on Computer Vision. pp. 242–259. Springer (2022) [4](#)
23. Gowda, S.N., Rohrbach, M., Sevilla-Lara, L.: Smart frame selection for action recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 1451–1459 (2021) [3](#)
24. Gowda, S.N., Sevilla-Lara, L., Keller, F., Rohrbach, M.: Cluster: clustering with reinforcement learning for zero-shot action recognition. In: European Conference on Computer Vision. pp. 187–203. Springer (2022) [3](#), [4](#), [11](#)
25. Gowda, S.N., Sevilla-Lara, L., Kim, K., Keller, F., Rohrbach, M.: A new split for evaluating true zero-shot action recognition. arXiv preprint arXiv:2107.13029 (2021) [11](#), [12](#)
26. Gowda, S.N., Yuan, C.: Colornet: Investigating the importance of color spaces for image classification. In: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018. pp. 581–596. Springer (2019) [1](#)
27. Han, Z., Fu, Z., Li, G., Yang, J.: Inference guided feature generation for generalized zero-shot learning. Neurocomputing **430**, 150–158 (2021) [9](#)
28. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) [1](#)
29. Kim, K., Gowda, S.N., Mac Aodha, O., Sevilla-Lara, L.: Capturing temporal information in a single frame: Channel sampling strategies for action recognition. arXiv preprint arXiv:2201.10394 (2022) [3](#)
30. Kodirov, E., Xiang, T., Fu, Z., Gong, S.: Unsupervised domain adaptation for zero-shot learning. In: Proceedings of the IEEE international conference on computer vision. pp. 2452–2460 (2015) [11](#)
31. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: Hmdb: a large video database for human motion recognition. In: 2011 International Conference on Computer Vision. pp. 2556–2563. IEEE (2011) [11](#)
32. Lin, C.C., Lin, K., Wang, L., Liu, Z., Li, L.: Cross-modal representation learning for zero-shot action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19978–19988 (2022) [4](#)
33. 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) [3](#)

34. Liu, J., Bai, H., Zhang, H., Liu, L.: Beyond normal distribution: More factual feature generation network for generalized zero-shot learning. *IEEE MultiMedia* (2022) [9](#), [10](#)
35. 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) [3](#)
36. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008) [7](#)
37. 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 Conference on Computer Vision and Pattern Recognition*. pp. 9985–9993 (2019) [3](#), [4](#), [5](#), [9](#), [11](#), [13](#)
38. Mettes, P., Snoek, C.G.: Spatial-aware object embeddings for zero-shot localization and classification of actions. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4443–4452 (2017) [11](#)
39. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Advances in neural information processing systems*. pp. 3111–3119 (2013) [2](#), [3](#), [4](#), [6](#), [7](#), [9](#)
40. Mirza, M., Osindero, S.: Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014) [5](#)
41. Mishra, A., Verma, V.K., Reddy, M.S.K., Arulkumar, S., Rai, P., Mittal, A.: A generative approach to zero-shot and few-shot action recognition. In: *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. pp. 372–380. *IEEE* (2018) [5](#), [9](#)
42. Ni, B., Peng, H., Chen, M., Zhang, S., Meng, G., Fu, J., Xiang, S., Ling, H.: Expanding language-image pretrained models for general video recognition. In: *European Conference on Computer Vision*. pp. 1–18. *Springer* (2022) [11](#), [12](#), [13](#)
43. Pagliardini, M., Gupta, P., Jaggi, M.: Unsupervised learning of sentence embeddings using compositional n-gram features. In: *Proceedings of NAACL-HLT*. pp. 528–540 (2018) [2](#)
44. Perrett, T., Masullo, A., Burghardt, T., Mirmehdi, M., Damen, D.: Temporal-relational crosstransformers for few-shot action recognition. *arXiv preprint arXiv:2101.06184* (2021) [4](#)
45. Qian, Y., Yu, L., Liu, W., Hauptmann, A.G.: Rethinking zero-shot action recognition: Learning from latent atomic actions. In: *European Conference on Computer Vision*. pp. 104–120. *Springer* (2022) [4](#), [11](#), [13](#)
46. Qin, J., Liu, L., Shao, L., Shen, F., Ni, B., Chen, J., Wang, Y.: Zero-shot action recognition with error-correcting output codes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2833–2842 (2017) [4](#), [11](#)
47. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics (11 2019), <https://arxiv.org/abs/1908.10084> [6](#)
48. Rohrbach, M., Regneri, M., Andriluka, M., Amin, S., Pinkal, M., Schiele, B.: Script data for attribute-based recognition of composite activities. In: *Eur. Conf. Comput. Vis.* (2012) [3](#)
49. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: *Advances in neural information processing systems*. pp. 568–576 (2014) [3](#)

50. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012) [7](#), [11](#)
51. Verma, V.K., Arora, G., Mishra, A., Rai, P.: Generalized zero-shot learning via synthesized examples. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4281–4289 (2018) [5](#)
52. Xian, Y., Lorenz, T., Schiele, B., Akata, Z.: Feature generating networks for zero-shot learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5542–5551 (2018) [2](#), [3](#), [5](#), [9](#)
53. Xu, X., Hospedales, T., Gong, S.: Transductive zero-shot action recognition by word-vector embedding. *International Journal of Computer Vision* **123**(3), 309–333 (2017) [4](#), [11](#)
54. Xu, X., Hospedales, T.M., Gong, S.: Multi-task zero-shot action recognition with prioritised data augmentation. In: European Conference on Computer Vision. pp. 343–359. Springer (2016) [4](#)
55. Zhu, Y., Long, Y., Guan, Y., Newsam, S., Shao, L.: Towards universal representation for unseen action recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9436–9445 (2018) [11](#)