



THE UNIVERSITY *of* EDINBURGH

## Edinburgh Research Explorer

### A New Split for Evaluating True Zero-Shot Action Recognition

**Citation for published version:**

N Gowda, S, Sevilla-Lara, L, Kim, K, Keller, F & Rohrbach, M 2022, A New Split for Evaluating True Zero-Shot Action Recognition. in C Bauckhage, J Gall & A Schwing (eds), *Proceedings of the 43rd DAGM German Conference on Pattern Recognition (GCPR 2021)*. 1 edn, Springer Nature Switzerland AG, Cham, pp. 191-205, 43rd DAGM German Conference on Pattern Recognition, 28/09/21.  
[https://doi.org/10.1007/978-3-030-92659-5\\_12](https://doi.org/10.1007/978-3-030-92659-5_12)

**Digital Object Identifier (DOI):**

[10.1007/978-3-030-92659-5\\_12](https://doi.org/10.1007/978-3-030-92659-5_12)

**Link:**

[Link to publication record in Edinburgh Research Explorer](#)

**Document Version:**

Peer reviewed version

**Published In:**

Proceedings of the 43rd DAGM German Conference on Pattern Recognition (GCPR 2021)

**General rights**

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [openaccess@ed.ac.uk](mailto:openaccess@ed.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# A New Split for Evaluating True Zero-Shot Action Recognition

Shreyank N Gowda<sup>1</sup>, Laura Sevilla-Lara<sup>1</sup>, Kiyoon Kim<sup>1</sup>, Frank Keller<sup>1</sup>, and  
Marcus Rohrbach<sup>2</sup>

<sup>1</sup> University of Edinburgh, UK

<sup>2</sup> Facebook AI Research, US

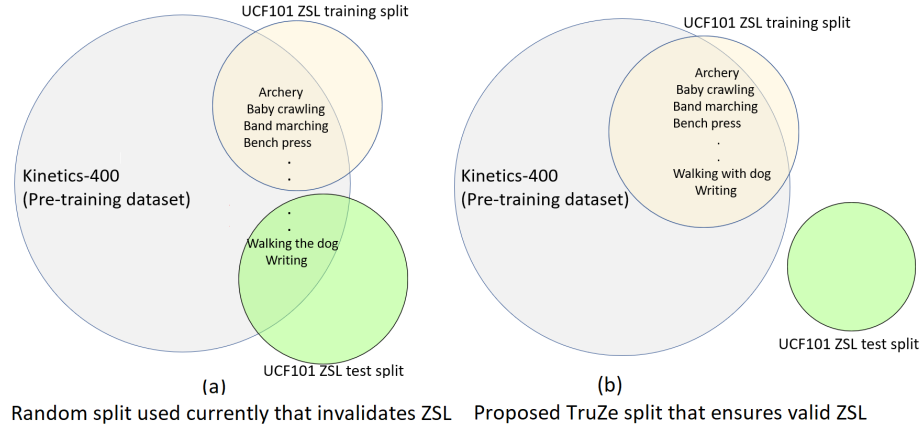
**Abstract.** Zero-shot action recognition is the task of classifying action categories that are not available in the training set. In this setting, the standard evaluation protocol is to use existing action recognition datasets (e.g. UCF101) and *randomly* split the classes into seen and unseen. However, most recent work builds on representations pre-trained on the Kinetics dataset, where classes largely overlap with classes in the zero-shot evaluation datasets. As a result, classes which are supposed to be unseen, are present during supervised pre-training, invalidating the condition of the zero-shot setting. A similar concern was previously noted several years ago for image based zero-shot recognition, but has not been considered by the zero-shot *action* recognition community. In this paper, we propose a new split for *true* zero-shot action recognition with no overlap between unseen test classes and training or pre-training classes. We benchmark several recent approaches on the proposed True Zero-Shot (**TruZe**) Split for UCF101 and HMDB51, with zero-shot and generalized zero-shot evaluation. In our extensive analysis we find that our TruZe splits are significantly harder than comparable random splits as nothing is leaking from pre-training, i.e. unseen performance is consistently lower, up to 8.9% for zero-shot action recognition. In an additional evaluation we also find that similar issues exist in the splits used in few-shot action recognition, here we see differences of up to 17.1%. We publish our splits <sup>1</sup> and hope that our benchmark analysis will change how the field is evaluating zero- and few-shot action recognition moving forward.

## 1 Introduction

Much of the recent progress in action recognition is due to the availability of large annotated datasets. Given how impractical it is to obtain thousands of videos in order to recognize a single class label, researchers have turned to the problem of zero-shot learning (ZSL). Each class label has semantic embeddings that are either manually annotated or inferred through semantic knowledge using word embeddings. These embeddings help obtain relationships between training classes (that have many samples) and test classes (that have zero samples).

---

<sup>1</sup> Splits can be found at <https://github.com/kini5gowda/TruZe>



**Fig. 1.** An illustration of the overlap in *classes* of the pretraining dataset (grey), training split (yellow) and zero-shot test split (green). Current evaluation protocol (a) picks classes at random with 51 training classes and 50 test classes. There is always some overlap and also chances of an extremely high overlap. We propose a stricter evaluation protocol (b) where there is no overlap at test time, maintaining the ZSL premise.

Typically, the model predicts the semantic embedding of the input video and matches it to a test class using the nearest neighbor’s search.

However, work in video ZSL [1,10,7] often uses a pre-trained model to represent videos. While pre-trained models help obtaining good visual representations, overlap with test classes can invalidate the premise of zero-shot learning, making it difficult to compare approaches fairly.

In the image domain [24,19,21,26], this problem has also been observed. Typically image models are pre-trained on ImageNet [5]. Xian et al. [26] showed that, in image ZSL, if the pre-training dataset has overlapping classes with the test set, the accuracy is inflated at test time. Hence, the authors propose a new split that avoids that problem, and it is now widely used. Similarly, most video models are pre-trained on Kinetics-400 [3], which has a large overlap with the typical ZSL action recognition benchmarks (UCF101, HMDB51 and Olympics). This pre-training gives leads to inflated accuracies, creating the need for a new split. Figure 1 shows an illustration of these overlap issues.

**Contributions:** First, we show the significant difference in performance caused by pre-training on classes that are included in the test set, across all networks and all datasets. Second, we measure the extent of the overlap between Kinetics-400 and the datasets typically used for ZSL testing: UCF101, HMDB51 and Olympics datasets. We do this by computing both visual and semantic similarity between classes. Finally, we propose a fair split of the classes that takes this class overlap into account, and does not break the premise of ZSL. We show that current models do indeed perform more poorly in this split, which is further proof of the significance of the problem. We hope that this split

will be useful to the community, will avoid the need of random splits, and help an actually fair comparison among methods.

## 2 Related Work

Previous work [26] has studied the effect of pre-training on classes that overlap with test classes in the image domain. The authors compute the extent of overlap between testing datasets and Imagenet [5], where models are typically pre-trained. The overlapping classes correspond to the training classes, while the non-overlapping classes correspond to the test classes. Figure 1 shows an illustration of the proposed evaluation protocol. Unlike the traditional evaluation protocol that chooses classes at random from the list of classes, without typically looking at the list of overlapping classes from the pre-trained dataset, we strictly remove all classes that have a high threshold of visual or semantic similarity (see Sec. 5).

Roitberg et al. [15] proposed to look at the overlapping classes in videos by using a corrective method that would automatically remove categories that are similar. This was done by utilizing a pairwise similarity of labels within the same dataset. While they showed that using pre-trained models resulted in improved accuracy due to class overlap, the evaluation included only one dataset, and only looked at the semantic similarity of labels. Adding visual similarity, helps discovering overlapping classes like “typing” in UCF101 and “using computer” in Kinetics. Therefore, in our proposed split we use both semantic and visual similarity across classes.

Busto et al. [2] provide a mapping of shared classes between UCF101 and Kinetics as part of a domain adaptation problem. They manually find semantic matches based on class names. However, their mapping was not based on visual and semantic similarity as they have classes such as typing and writing on board as part of UCF101 classes not similar to any Kinetics class. Also, they use “floor gymnastics” as an unknown class in UCF101, however, Kinetics has “gymnastics tumbling” which is the same action. We see that samples from the “typing” class consist of a large proportion of people using their computers and this maps directly to the “using computers” class in Kinetics. Based on our visual and semantic similarity approach, we obtain a slightly different set of classes to those proposed by Busto et al.

Recently, end-to-end training [1] has been proposed for ZSL in video classification. As part of the evaluation protocol, to uphold the ZSL premise, the authors propose to train a model on Kinetics by removing the set of overlapping classes (using semantic matching) and using this as a pre-trained model. While this is a promising way to ensure the following of the premise of ZSL, it is very computationally expensive. We also show that having a better backbone (see Sec 6.4) results in better accuracy, and as such, training end-to-end is expensive. As a result, using a proposed split instead whilst having the opportunity to use any backbone seems an easier approach.

### 3 ZSL preliminaries

Consider  $S$  to be the training set of seen classes, composed of tuples  $(x, y, a(y))$ , where  $x$  represents the visual features of each sample in  $S$  (spatio-temporal features in the case of video),  $y$  corresponds to the class label in the set of seen class labels  $Y_s$ , and  $a(y)$  represents the semantic embedding of class  $y$ . These semantic embeddings are either annotated manually or computed using a language-based embedding, e.g. word2vec [11] or sen2vec [13].

Let  $U$  be the set of unseen classes, composed of tuples  $(u, a(u))$ , where  $u$  is a class in the label set  $Y_u$ , and  $a(u)$  are the corresponding semantic representations.  $Y_s$  and  $Y_u$  do not overlap, i.e.

$$Y_s \cap Y_u = \emptyset \quad (1)$$

In ZSL for video classification, given an input video, the task is to predict a class label in the unseen set of classes,  $f_{ZSL} : X \rightarrow Y_u$ . An extension of the problem is the related generalized zero-shot learning (GZSL) setting, where given a video, the task is to predict a class label in the union of the seen and unseen classes, as  $f_{GZSL} : X \rightarrow Y_s \cup Y_u$ .

When relying on a pre-trained model to obtain visual features, we denote the pre-trained classes as the set  $Y_p$ . For the ZSL premise to be maintained, there must be no overlap with the unseen classes:

$$Y_p \cap Y_u = \emptyset. \quad (2)$$

The core problem we address in this paper is that while prior work generally adheres to Eq. 1, recent use of pre-trained models does not adhere to Eq. 2. Instead, we propose the TruZe split in Section 5.2, which adheres to both Eq. (1) and Eq. (2).

#### 3.1 Visual and Semantic Embeddings

Early work computed **visual embeddings** (or representations) using hand-crafted features such as Improved Dense Trajectories (IDT) [20], which include tracked trajectories of detected interest points within a video, and four descriptors. More recent work often uses deep features such as those from 3D convolutional networks (e.g., I3D [3] or C3D [18]). These 3D CNNs are used to learn spatio-temporal representation of the video. In our experiments, we will use both types of visual representations.

To obtain **semantic embeddings**, previous work [9] uses manual attribute annotations for each class. For example, the action of kicking would have motion of the leg and motion of twisting the upper body. However, such attributes are not available for all datasets. An alternative approach is to use word embeddings such as word2vec [11] for each class label. This gets rid of the requirement of manual attributes. More recently, Gowda et al.[7] showed that using sen2vec [13] instead of word2vec yields better results as action labels are typically multi-worded and averaging them using word2vec makes it lose context. Based on this, in our experiments we use sen2vec.

## 4 Evaluated Methods

We consider early approaches that use IDT features such as ZSL by bi-directional latent embedding learning (BiDiLEL), ZSL by single latent embedding (Latem) [23] and synthesized classifiers (SYNC) [4]. Using features that are not learned, allows us to control for the effect of pre-training when using random splits, and when using the proposed split (PS). We then evaluate recent state-of-the-art approaches such as feature generating networks (WGAN) [25], out-of-distribution detection networks (OD) [10] and end-to-end learning for ZSL (E2E) [1] as well. Let us briefly have a look at these methods.

**Latem** [23] uses piece-wise linear compatibility to understand the visual-semantic embedding relationship with the help of latent variables. Here, each latent variable is encoded to account for the various visual properties of the input data. The authors project the visual embedding to the semantic embedding space.

**BiDiLEL** [21] projects both the visual and semantic embeddings into a common latent space (instead of projecting to the semantic space) so that the intrinsic relationship between them is maintained.

**SYNC** [4] uses a weighted bipartite graph in order to learn a projection between the semantic embeddings and the classifier model space. They generate the graph by using a set of "phantom" classes synthesized in order to ensure aligned semantic embedding space and classifier model space and minimize the distortion error.

**WGAN** [25] uses a Wasserstein GAN to synthesize the unseen features of classes, with additional losses in the form of cosine and cycle-consistency losses. These losses help enhancing the feature generation process.

**OD** [10] trains an out-of-distribution detector to distinguish the generated features from those of the seen class features and in turn to help with classification in the generalized zero-shot learning setting.

**E2E** [1] is a recent approach that leverages end-to-end training to alleviate the problem of overlapping classes. This is done by removing all overlapping classes in the pre-training dataset and then using a CNN trained on the remaining classes to generate the visual features for the ZSL videos.

**CLUSTER** [7] uses clustering of visual-semantic embeddings optimised by reinforcement learning.

## 5 Evaluation protocol

### 5.1 Datasets

The three most popular benchmarks for ZSL in videos are UCF101 [17], HMDB51 [8] and Olympics [12]. The typical evaluation protocol in video ZSL is to use a 50-50 split of each dataset, where 50 % of the labels are used as the training set and 50 % as the test set. In order to provide comparisons to prior work [21,23,4,10,1,7] and for the purpose of communicating replicable research results, we study UCF101, HMDB51, and Olympics, as well as the the relationship to the pre-training dataset Kinetics-400.

Dataset	Videos	Classes	Random Split* (Seen/Unseen)	Overlapping classes with Kinetics	TruZe Split (Seen/Unseen)
Olympics	783	16	8/8	13	-
HMDB51	6766	51	26/25	29	29/22
UCF101	13320	101	51/50	70	70/31

**Table 1.** Datasets and their splits used for ZSL in action recognition. Traditionally, ‘Random Split’ was followed where the seen and unseen classes were randomly selected. However, we can see the extent of overlap in the ‘overlapping classes’ column. Using the extent of overlap we define our ‘TruZe split’. For the full list of seen and unseen classes, please look at the supplementary. \*Note that for the all experiments in this paper we use a random split which matches the number of classes of our TruZe, e.g. 29/22 for HMDB51.

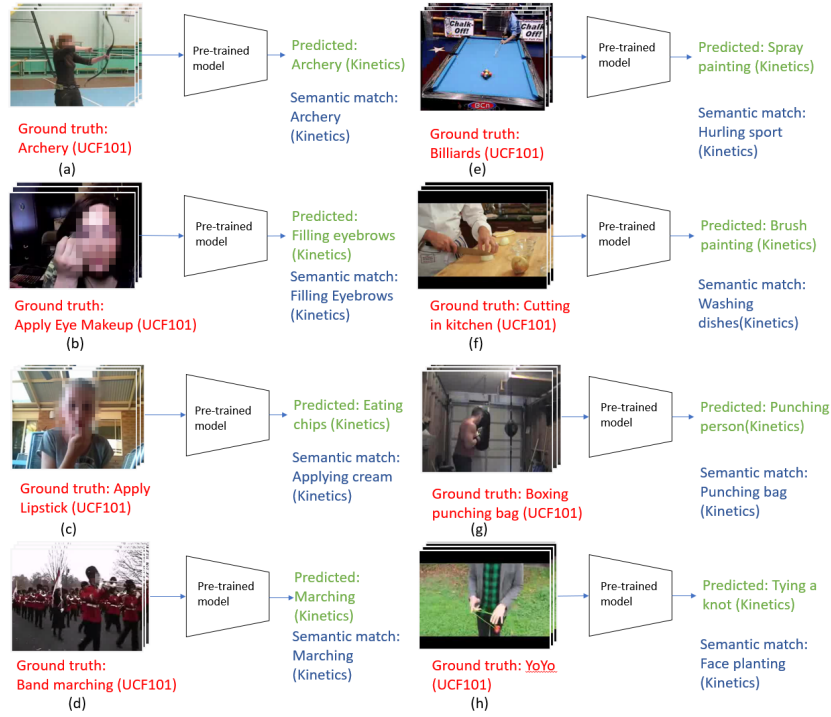
In our experiments (see Section 6), we find overlapping classes between Kinetics-400 and each of the ZSL datasets, and move them to the training split. Thus, instead of using 50-50, we need to use 70-31 (number of labels for train and test) for UCF101 and 29-22 (number of labels for train and test) for HMDB51. We see that the number of overlapping classes in the case of Olympics is 13 out of 16, and hence we choose not to proceed further with it. More details can be found in Table 1. For a fair comparison between the TruZe and random split, we use the same proportions (i.e., 70-31 in UCF101 and so on) in the experiments with random splits. We create ten such random splits and use these same splits for all models.

## 5.2 TruZe Split

We now describe the process of creating the proposed TruZe split, to avoid the coincidental influence of pre-training on ZSL. First, we identify overlapping classes between the pre-training Kinetics-400 dataset and each ZSL dataset. To do this, we compute visual and semantic similarities, and discard those classes that are too similar.

To calculate visual similarity, we use an I3D model pre-trained on Kinetics-400 and evaluate all video samples in UCF101, HMDB51 and Olympics using the Kinetics labels. This helps us to detect similarities that are often not recognized in terms of semantic similarities. Some examples include typing (class in UCF101) that the model detects as using computer (class in Kinetics), applying eye makeup (class in UCF101) that the model detects as filling eyebrows (class in Kinetics).

To calculate semantic similarity, we use a sen2vec model pre-trained on Wikipedia that helps us compare action phrases and outputs a similarity score. We combine the visual and semantic similarity to obtain a list of extremely similar classes to the ones present in Kinetics. This list of classes is present in the supplementary. The classes that even have a slight overlap or are a subset of a class in Kinetics are all chosen as part of the seen set (for example, cricket



**Fig. 2.** A few examples of how the classes are selected. (a) is an example of an exact match between the testing dataset (in this case UCF101) and the pre-trained dataset (Kinetics). (b), (d) and (g) are examples of visual-semantic similar matches where the output and semantically closest classes are the same. (c), (e), (f) and (h) are examples of classes without overlap in terms of both visual and semantic similarity.

bowling and cricket shot in UCF101 are part of the seen set due to the superclass playing cricket in Kinetics). A few examples of the selection of classes is show in Figure 2.

We discard classes from the test set based on the following rules:

- Discard exact matches. For example, archery in UCF101 is also present in Kinetics.
- Discard matches that can be either superset or subset. For example, UCF101 has classes such as cricket shot and cricket bowling while Kinetics has playing cricket (superset). We manually do this based on the output of the closest semantic match.
- Discard matches that predict the same visual and semantic match. For example, *apply eye makeup* (UCF101 label) predicts *filling eyebrows* as the visual match using Kinetics labels and the closest semantic match to classes in Kinetics is also *filling eyebrows*. We also manually confirm this.

Method	UCF101			HMDB51		
	Random	TruZe	Diff	Random	TruZe	Diff
Latem [23]	21.4	15.5	5.9	17.8	9.4	8.4
SYNC [4]	22.1	15.3	6.8	18.1	11.6	6.5
BiDiLEL [21]	21.3	15.7	5.6	18.4	10.5	7.9
OD [10]	28.4	22.9	5.5	30.6	21.7	8.9
E2E [1]	46.6	45.5	1.1	33.2	31.5	1.7
CLUSTER [7]	47.1	45.2	1.9	36.6	33.2	3.4

**Table 2.** Results with different splits for **Zero-Shot Learning (ZSL)**. Column ‘Random’ corresponds to the accuracy using splits in the traditional fashion (random selection of train and test classes, but with the same number of classes in train/test as in TruZe), ‘TruZe’ corresponds to the accuracy using our proposed split and ‘Diff’ corresponds to the difference in accuracy between using random splits and our proposed split. We run 10 independent runs for different random splits and report the average accuracy. We see positive differences in the ‘Diff’ column which we believe is due to the overlapping classes in Kinetics.

We move all the discarded classes to the training set. This leaves a 70-31 split on UCF101 and a 29-22 split on HMDB51. We also see that in the Olympics dataset, there are 13 directly overlapping classes out of 16 classes and hence dropped the dataset from further analysis. One particular interesting scenario is the “pizza tossing” class in UCF101. In Kinetics, there is a class called “making pizza”, however, the action of tossing is not performed in them and hence we use “pizza tossing” as an unseen class.

## 6 Experimental Results

### 6.1 Results on ZSL and Generalized ZSL

We first consider the results on ZSL. Here, as explained before, only samples from the unseen class are passed as input to the model at test time. Since TruZe separates the overlapping classes from the pre-training dataset, we expect a lower accuracy on this split compared to the traditionally used random splits. We compare BiDiLEL [21], Latem [23], SYNC [4], OD [10], E2E [1] and CLUSTER [7] and report the results in Table 2. As expected, we see in the ‘Diff’ column for both UCF101 and HMDB51 a positive difference, indicating that the accuracy is lower for the TruZe split.

Generalized ZSL (GZSL) looks at a more realistic scenario, wherein the samples at test time belong to both seen and unseen classes. The reported accuracy is then the harmonic mean of the seen and unseen class accuracies. Since we separate out the overlapping classes, we expect to see an increase in the seen class accuracy and a decrease in the unseen class accuracy. We report GZSL results on OD, WGAN and CLUSTER in Table 3. The semantic embedding used for all models is sen2vec. We use 70 classes for training chosen at random along

Method	Acc <sub>U</sub>			Acc <sub>S</sub>			Harmonic mean			Dataset
	Rand	TruZe	Diff	Rand	TruZe	Diff	Rand	TruZe	Diff	
WGAN [25]	27.9	21.3	6.6	58.2	63.2	-5.0	37.7	31.8	5.9	HMDB51
WGAN [25]	28.2	23.9	4.3	74.9	75.6	-0.7	41.0	36.3	4.7	UCF101
OD [10]	34.1	24.7	9.4	58.5	62.8	-4.3	43.1	35.5	7.6	HMDB51
OD [10]	32.6	29.1	3.5	76.1	78.4	-2.3	45.6	42.4	3.2	UCF101
CLUSTER [7]	41.8	38.4	3.4	52.3	53.1	-0.8	46.4	44.5	1.9	HMDB51
CLUSTER [7]	37.5	35.6	1.9	68.8	70.6	-1.8	48.5	47.3	1.2	UCF101

**Table 3.** Results with different splits for **Generalized Zero-Shot Learning (GZSL)**. ‘Rand’ corresponds to the splits using random classes over 10 independent runs, ‘TruZe’ corresponds to the proposed split. Acc<sub>U</sub> and Acc<sub>S</sub> correspond to unseen class accuracy and seen class accuracy respectively. The semantic embedding used is sen2vec. ‘diff’ corresponds to the difference between ‘Rand’ and ‘TruZe’. We see consistent positive difference in performance on the unseen classes and negative difference in the performance of the seen classes while using the ‘TruZe’.



**Fig. 3.** Graphical representation of the difference in performances of different models on GZSL. We see consistent positive difference in performance on the unseen classes and negative difference in the performance of the seen classes while using the ‘TruZe’. The x-axis corresponds to difference in accuracy (Random splits accuracy - TruZe split accuracy) and the y-axis to different methods.

with 31 test classes (also chosen at random) for UCF101 and 29 training with 22 testing for HMDB51. As expected, the average unseen class accuracy drops in the proposed split and the average seen class accuracy increases. We expect this as the unseen classes are more disjoint in the proposed split than using random splits. For easier understanding, we convert the differences in Table 3 to a graph and this can be seen in Figure 3.

## 6.2 Extension to Few-shot Learning

Few-shot learning (FSL) is another scenario we consider. Since the premise is the same as ZSL, except that we have a few samples instead of zero. Again, usually, the splits used are random, and as such, the pre-trained model has seen hundreds of samples of classes that are supposed to belong to the test

Method	SS					TruZe					Diff				
	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
C3D-PN [16]	57.1	66.4	71.7	75.5	78.2	50.9	61.9	67.5	72.9	75.4	6.2	4.5	4.2	2.6	2.8
ARN [27]	66.3	73.1	77.9	80.4	83.1	61.2	70.7	75.2	78.8	80.2	5.1	2.4	2.7	1.6	2.9
TRX [14]	77.5	88.8	92.8	94.7	96.1	75.2	88.1	91.5	93.1	93.5	2.5	0.7	1.3	1.6	2.6

**Table 4. Few Shot Learning (FSL)** with different splits on UCF101. Accuracies are reported for 5-way, 1, 2, 3, 4, 5-shot classification. 'SS' corresponds to the split used in [27,14] and 'TruZe' corresponds to the proposed split. We can see that using our proposed split results in a drop in performance of up to 6.2 % for UCF101. This shows TruZe is much harder even in the FSL scenario.

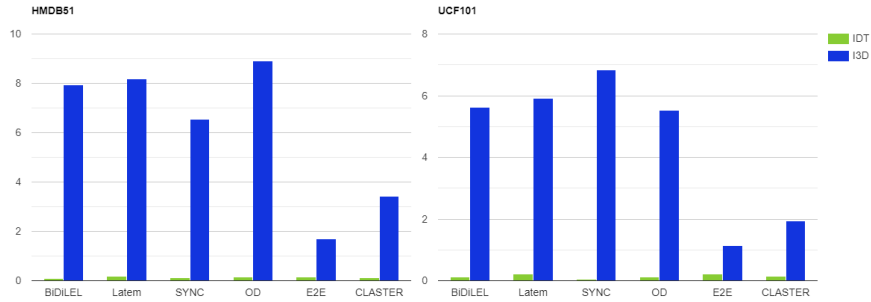
Method	SS					TruZe					Diff				
	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
C3D-PN [16]	38.1	47.5	50.3	55.6	57.4	28.8	38.5	43.4	46.7	49.1	9.3	9.0	6.9	8.9	8.3
ARN [27]	45.5	50.1	54.2	58.7	60.6	31.9	42.3	46.5	49.8	53.2	12.6	7.8	7.7	8.9	7.4
TRX [14]	50.5	62.7	66.9	73.5	75.6	33.5	46.7	49.8	57.9	61.5	17.0	16.0	17.1	15.6	14.1

**Table 5. Few Shot Learning (FSL)** with different splits on HMDB51. Accuracies are reported for 5-way, 1, 2, 3, 4, 5-shot classification. 'SS' corresponds to the split used in [27,14] and 'TruZe' corresponds to the proposed split. We can see that using our proposed split results in a drop in performance of up to 17.1 % for HMDB51. This shows TruZe is much harder even in the FSL scenario.

set. We report results on the 5-way, 1,2,3,4,5-shot case for temporal relational cross-transformers (TRX) [14], action relation network (ARN) [27], and C3D prototypical net (C3D-PN) [16]. Results are reported in Table 4 and Table 5. The standard split (SS) used here is taken from the one proposed in ARN [27]. Similar to the SS, we divide the classes in UCF101 and HMDB51 to (70,10,21) and (31,10,10), respectively, where the order corresponds to the number of training classes, validation classes and test classes. We see that the proposed split is much harder than SS. Consistent drops in performance can be seen on every split and for every model. Performance drops of upto 6.2% on UCF101 and 17.1% on HMDB51 can be seen. Our proposed splits are available in the supplementary material.

### 6.3 Is overlap the reason for performance difference between Random and our TruZe split?

In order to understand the difference in model performance due to the overlapping classes, we compare the performance of each model for the random split (with five runs) vs the proposed split by using visual features represented by IDT and I3D. We depict the difference in performance in the form of a bar graph for better visual understanding. This is seen in Figure 4. The higher the difference, the bigger the impact of performance. We can see that there is a big



**Fig. 4.** The difference of accuracy for different models using IDT and I3D using manual annotations as the semantic embedding. The larger the bar, the more significant the difference. We can see a clear difference when using I3D and this difference is due to the presence of overlapping classes in the test set. The y-axis corresponds to the difference in performance in percentage and the x-axis corresponds to various models.

difference when using I3D features compared to using IDT features (where there is a minimal difference). Since IDT features are independent of any pre-training model, the difference in performance is negligible. The difference while using I3D features can be attributed to the presence of overlapping classes in the random splits compared to the proposed split.

#### 6.4 Use of different backbone networks

An end-to-end approach was proposed in [1] where a 3D CNN was trained in an end-to-end manner on classes in Kinetics not semantically similar to UCF101 and HMDB51 to overcome the overlapping classes conundrum. While this approach is useful, training more complex models end-to-end is not feasible for everyone due to the high computational cost involved. We show that using more recent state-of-the-art approaches as the backbone, there is a slight improvement in model performance and hence believe that having a proposed split instead of training end-to-end would be more easily affordable for the general public. Table 6 shows the results of using different backbones for extracting visual features on some of the recent state-of-the-art ZSL approaches. We use Non-Local networks [22] that build on I3D by adding long-term spatio-temporal dependencies in the form of non-local connections (referred as NL-I3D in Table 6). We also use slow-fast networks [6] that is a recent state-of-the-art approach that uses two pathways, a slow and a fast, to capture motion and fine temporal information. We can see minor but consistent improvements using stronger backbones, and this suggests that having a proposed split is an economical way of maximising the use of state-of-the-art models as backbone networks. We see gains of up to 0.6% in UCF101 and 0.8% in HMDB51.

Method	Backbone	UCF101 Accuracy	HMDB51 Accuracy
WGAN [25]	I3D	22.5	21.1
WGAN [25]	NL-I3D	22.7	21.3
WGAN [25]	SlowFast	<b>23.1</b>	<b>21.5</b>
OD [10]	I3D	22.9	21.7
OD [10]	NL-I3D	23.2	22.0
OD [10]	SlowFast	<b>23.4</b>	<b>22.5</b>
CLASTER [7]	I3D	45.2	33.2
CLASTER [7]	NL-I3D	45.3	33.6
CLASTER [7]	SlowFast	<b>45.5</b>	<b>33.9</b>

**Table 6.** Results comparison using different backbones to extract visual features for the ZSL models. We evaluate OD, E2E and CLASTER using I3D, NL-I3D and SlowFast networks as backbones. All results are on the proposed split. We see that stronger backbones result in improved performance of the ZSL model.

## 7 Implementation Details

### 7.1 Visual features

We use either IDT [20] or I3D [3] for the visual features. Using the fisher vector obtained from a 256 component Gaussian mixture model, we generate visual feature representations using IDT (contains four different descriptors). To reduce this, PCA is used to obtain a 3000-dimensional vector for each descriptor. Concatenating these (all four descriptors), we obtain a 12000-dimensional vector for each video. In the case of I3D features, we use RGB and flow features taken from the *mixed 5c* layer from a pre-trained I3D (pre-trained on Kinetics-400). The output of the flow network is averaged across the temporal dimension and pooled by four in the spatial dimension, and then flattened to a vector of size 4096. We then concatenate the two.

### 7.2 Semantic embedding

While manual annotations are available for UCF101 in the form of a vector of size 40, there is no such annotation available for HMDB51. Hence, we use sen2vec embeddings of the action classes where the sen2vec model is pre-trained on Wikipedia. While most approaches use word2vec and average embeddings for each word in the label, we use sen2vec which obtains an embedding for the entire label.

### 7.3 Hyperparameters for evaluated methods

We use the optimal parameters reported in BiDiLEL [21]. The values for  $\alpha$  and  $k_G$  values are set to 10 and  $d_y$  is set to 150. SYNC [4] has a parameter  $\sigma$  that models correlation between a real class and a phantom class and this is set to

1, while the balance coefficient is set to  $2^{-10}$ . For Latem [23] the learning rate, number of epochs, and number of embeddings are 0.1, 200 and 10 respectively. For OD [10], WGAN [25], E2E [1] and CLASTER [7] we follow the settings provided by the authors. For few-shot learning, we use the hyperparameters defined in the papers [14,27]. We compare against the standard split proposed in [27]. For the proposed split, we change the classes slightly for fair comparison to the standard split. Now the splits for HMDB51 and UCF101 are (31,10,10) and (70,10,21) where the order corresponds to (train,val,test).

## 8 Discussion and Conclusion

As we see in Figure 4 using IDT features which do not require pre-training on Kinetics resulted in a negligible change in performance comparing the TruZe split vs the random splits. However, using I3D features saw a stark difference due to the overlapping classes in the pre-trained dataset.

We see that the proposed split is harder in all scenarios (ZSL, GZSL, and FSL) whilst maintaining the premise of the problem. The differences are significant in most cases: between 0.7-6.2 % for UCF101 and 7.4-17.1 % for HMDB51 in FSL, an increase of 1.2-4.7 % for UCF101 and 1.9-7.6 % for HMDB51 (with respect to the harmonic mean of seen and unseen classes) in GZSL and an increase of 1.1-6.8 % for UCF101 and 1.7-8.1 % for HMDB51 in ZSL. It is also important to note that different methods are differently affected, suggesting that some method in the past have claimed improvements due to not adhering to the zero-shot premise, which is highly concerning.

We also see that changing the backbone network increases the performance slightly for each model, and as a result, the end-to-end pre-training [1] can prove very expensive. As such, having a proposed split makes things easier as we can directly use pre-trained models off the shelf. We see gains of up to 1.3% in UCF101 and 1.1% in HMDB51.

Details to our TruZe splits can be found in supplemental material and we will release them publicly so the research community can fairly compare zero-shot and few-shot action recognition approaches and compare to the benchmark results provided in this paper.

## References

1. 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)
2. Busto, P.P., Iqbal, A., Gall, J.: Open set domain adaptation for image and action recognition. *IEEE transactions on pattern analysis and machine intelligence* **42**(2), 413–429 (2018)
3. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)

4. Changpinyo, S., Chao, W.L., Gong, B., Sha, F.: Synthesized classifiers for zero-shot learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5327–5336 (2016)
5. 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)
6. 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)
7. Gowda, S.N., Sevilla-Lara, L., Keller, F., Rohrbach, M.: Cluster: Clustering with reinforcement learning for zero-shot action recognition. arXiv preprint arXiv:2101.07042 (2021)
8. 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)
9. Liu, J., Kuipers, B., Savarese, S.: Recognizing human actions by attributes. In: CVPR 2011. pp. 3337–3344. IEEE (2011)
10. 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)
11. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. arXiv preprint arXiv:1310.4546 (2013)
12. Niebles, J.C., Chen, C.W., Fei-Fei, L.: Modeling temporal structure of decomposable motion segments for activity classification. In: European conference on computer vision. pp. 392–405. Springer (2010)
13. Pagliardini, M., Gupta, P., Jaggi, M.: Unsupervised learning of sentence embeddings using compositional n-gram features. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). pp. 528–540 (2018)
14. Perrett, T., Masullo, A., Burghardt, T., Mirmehdi, M., Damen, D.: Temporal-relational crosstransformers for few-shot action recognition. arXiv preprint arXiv:2101.06184 (2021)
15. Roitberg, A., Martinez, M., Haurilet, M., Stiefelwagen, R.: Towards a fair evaluation of zero-shot action recognition using external data. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 0–0 (2018)
16. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. arXiv preprint arXiv:1703.05175 (2017)
17. 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)
18. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 4489–4497 (2015)
19. 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)
20. Wang, H., Schmid, C.: Action recognition with improved trajectories. In: Proceedings of the IEEE international conference on computer vision. pp. 3551–3558 (2013)

21. Wang, Q., Chen, K.: Zero-shot visual recognition via bidirectional latent embedding. *International Journal of Computer Vision* **124**(3), 356–383 (2017)
22. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 7794–7803 (2018)
23. Xian, Y., Akata, Z., Sharma, G., Nguyen, Q., Hein, M., Schiele, B.: Latent embeddings for zero-shot classification. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 69–77 (2016)
24. 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)
25. 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)
26. Xian, Y., Schiele, B., Akata, Z.: Zero-shot learning-the good, the bad and the ugly. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4582–4591 (2017)
27. Zhang, H., Zhang, L., Qi, X., Li, H., Torr, P.H., Koniusz, P.: Few-shot action recognition with permutation-invariant attention. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer (2020)