

Synthetic Sample Selection for Generalized Zero-Shot Learning

Shreyank N Gowda
University of Edinburgh

s.narayana-gowda@sms.ed.ac.uk

Abstract

Generalized Zero-Shot Learning (GZSL) has emerged as a pivotal research domain in computer vision, owing to its capability to recognize objects that have not been seen during training. Despite the significant progress achieved by generative techniques in converting traditional GZSL to fully supervised learning, they tend to generate a large number of synthetic features that are often redundant, thereby increasing training time and decreasing accuracy. To address this issue, this paper proposes a novel approach for synthetic feature selection using reinforcement learning. In particular, we propose a transformer-based selector that is trained through proximal policy optimization (PPO) to select synthetic features based on the validation classification accuracy of the seen classes, which serves as a reward. The proposed method is model-agnostic and data-agnostic, making it applicable to both images and videos and versatile for diverse applications. Our experimental results demonstrate the superiority of our approach over existing feature-generating methods, yielding improved overall performance on multiple benchmarks.

1. Introduction

In recent years, deep learning [17, 19, 21, 43] has made remarkable strides in recognition accuracy, approaching human levels. However, the practical implementation of deep learning models is limited by the need for a significant number of labeled samples and the high cost of large-scale datasets [4]. It is often infeasible to collect sufficient labeled data for all classes, presenting a significant challenge for traditional supervised learning methods. To address this challenge, several approaches have been proposed, including semi-supervised learning, transfer learning, and few-shot learning. Zero-shot learning (ZSL) [36] is a subset of these methods, which refers to tasks where there is no training data available for unseen classes and disjoint label sets between training and test data.

ZSL is a unique form of cross-modal retrieval learning that relies on knowledge transfer from known classes to un-

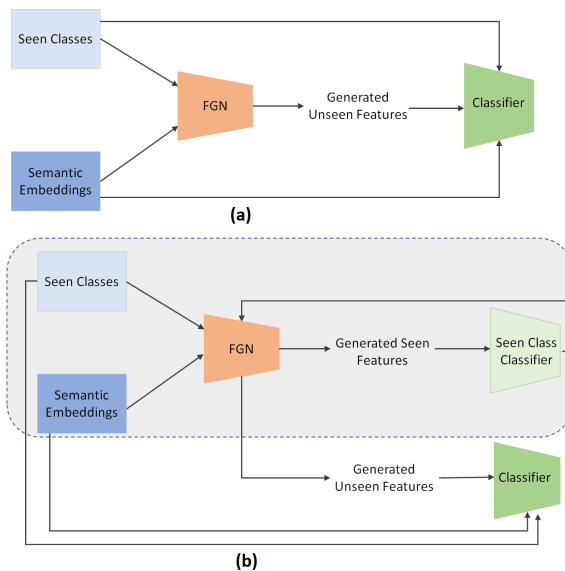


Figure 1. Comparison of feature-generating frameworks pipelines (a) standard pipeline where features that look “real” are used to train the generator (b) proposed pipeline where the generator is trained based on the performance of the seen class classifier.

known classes through attribute sharing. The most common ZSL techniques use an intermediate semantic representation, such as visual features or semantic word vectors [7, 24, 44, 61], shared between the labeled auxiliary dataset and the unlabeled target dataset. The projection from the low-level feature space to the semantic space is learned from the auxiliary dataset, without adapting to the target dataset. The Generalized ZSL scenario, where both seen and unseen class samples are available at test time, is considered to be more realistic than the standard ZSL setup [51].

In Generalized ZSL, classifiers tend to be biased towards seen classes, leading to the misclassification of test samples from unseen classes as seen classes. To address the problem of the lack of visual data for unseen classes, researchers have proposed the use of Generative Adversarial Networks (GAN) [12] to generate synthetic visual features

by leveraging attribute databases. However, while GANs have helped in zero-shot learning [6, 48, 52, 63], they do not explicitly learn to represent data in a structured way that is easily interpretable by humans or other models. They also suffer from the problems of mode collapse [22], class imbalance [1], and computational expense [18] when generating high-dimensional data. But what we care about most is that a lot of synthetic samples generated are used directly for training a classifier without studying if these samples actually help the classifier learn. Instead, these samples are chosen based on “realness”. Figure 1 shows a comparison of the standard pipeline and our proposed pipeline for feature-generating approaches.

To address the limitations of GANs in synthetic feature selection, we propose a novel reinforcement learning-based approach that automatically selects generated features that improve model performance. Specifically, we use a transformer model [47] for synthetic sample selection and use validation classification accuracy as the reward for RL training. We employ the proximal policy optimization (PPO) [42] algorithm to update the model weights during training. Our proposed approach aims to pick samples that help classification and not just generate real-looking samples. We dub our synthetic sample selection method as “**SPOT**” for Selection using Proximal policy **OpT**imization.

Furthermore, our proposed approach is model-agnostic and data-agnostic, as we evaluate our method on multiple benchmark datasets in images and videos and various feature-generating models. Our comprehensive experiments demonstrate that our approach consistently improves model performance across different datasets and models, highlighting the effectiveness and versatility of our proposed method. By leveraging RL-based synthetic feature selection, we can more effectively generate synthetic data that captures the underlying structure of the data, improving the generalization performance of downstream models.

2. Related Work

Zero-Shot Learning in Images Zero-shot learning (ZSL) is a challenging problem in computer vision, where the task is to recognize object categories without any training examples for them. Various approaches have been proposed to solve this problem in images. One of the early works [26] in this field used attributes, such as color and shape, to describe the object categories and mapped them to a visual space. They then used a nearest-neighbor classifier to recognize unseen object categories. However, this approach suffers from the semantic gap problem, where the attributes do not always correlate well with the visual features.

To address this problem, more recent works have explored the use of deep learning techniques to learn a joint embedding space for the visual and semantic features. One such approach is proposed by Frome et al. (2013) [7], where

they used a deep neural network to learn a joint embedding space for the visual and textual features of the objects. They then used a nearest-neighbor classifier to recognize unseen object categories. Another approach is proposed by Socher et al. (2013) [44], where they used a recursive neural network to learn a compositional representation of the textual descriptions of the object categories.

More recently, there has been a trend towards using generative models to solve the ZSL problem. One such approach is proposed by Xian et al. [52], where they used a generative adversarial network (GAN) to generate visual features for unseen object categories. They then used a joint embedding space to match the generated features with the semantic features and recognize unseen object categories. Another approach is proposed by Schonfeld et al. (2019) [40], where they used a GAN to generate visual features conditioned on the textual descriptions of the object categories. ZeroGen [59] uses pre-trained language models to synthesize a dataset for a zero-shot task and then trains a small task model on it. NereNet [29] generates unseen samples by combining noise from similar seen classes with unseen class attributes using GAN. CMC-GAN [58], performs data hallucination of unseen classes by performing semantics-guided intra-category knowledge transfer across image categories.

We add our proposed module of synthetic sample selection to multiple feature-generating frameworks and show that this leads to improved performance for all these models across all datasets.

Zero-Shot Learning in Videos The initial study by Rohrbach et al. [39] utilized script data from cooking activities to facilitate the transfer of knowledge to unseen categories. Gan et al. [10] considered each action class as a domain and tackled the problem of identifying semantic representations as a multisource domain generalization task. To extract semantic embeddings of class labels, popular approaches employ label embeddings such as word2vec [31], which solely requires class names. Several methods have used a shared embedding space between video features and class labels [56, 57], error-correcting codes [38], pairwise relationships between classes [8], interclass relationships [9], out-of-distribution detectors [30], synthetic features [32, 33] and graph neural networks [11].

Recently, it has been observed that clustering joint visual-semantic features results in better representations for zero-shot action recognition [15]. Similar to CLASTER, ReST [28] jointly encodes video data and textual labels for zero-shot action recognition. In ReST, transformers are utilized to conduct modality-specific attention. On the other hand, JigSawNet [37] models visual and textual features jointly but disassembles videos into atomic actions in an unsupervised manner and establishes group-to-group relation-

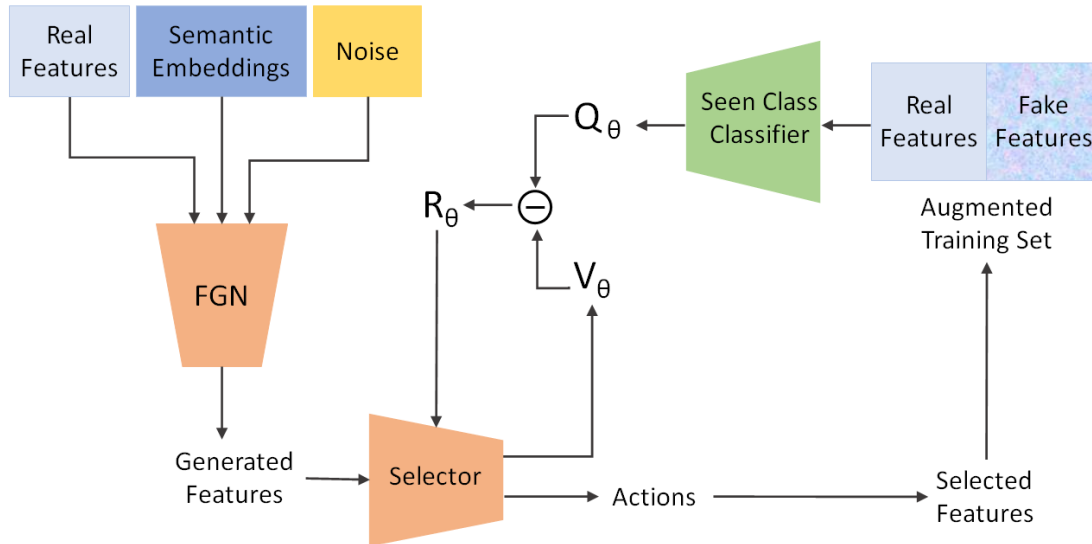


Figure 2. Overall pipeline of our proposed SPOT. The feature generator generates features that the selector module ranks based on the seen class classifier’s performance. The selector is updated based on the performance of the classifier on the selected features. The proposed pipeline is model and data-agnostic.

ships between visual and semantic representations instead of the one-to-one relationships that CLASTER and ReST establish.

However, since we only evaluate on feature generating approaches, we compare directly to OD [30], CLSWGAN [52], GGM [33] and Bi-dir GAN [32] and show that using our synthetic feature selection approach all methods can be improved significantly.

Reinforcement Learning for Data Valuation The quantification of data value for a specific machine learning task, known as data valuation [14, 23, 50], has numerous applications such as domain adaptation, corrupted sample detection, and robust learning. Various techniques have been proposed to estimate data values based on different criteria, including influence functions, Shapley values, leave-one-out errors, and data deletion. However, these methods are computationally expensive, necessitate model perturbations or retraining, and do not consider the interactions among data points. Recently, an adaptive approach to data valuation using reinforcement learning [62], in which data values are jointly learned with the predictor model using a data value estimator that is trained using a reinforcement signal reflecting task performance. Similarly, Learn2Augment [13] performs data valuation of augmented samples created by combining foreground and background videos using reinforcement learning to quantify the value of an augmented sample.

Synthetic sample selection [60] for medical image segmentation is an under-investigated research area that fo-

cuses on the quality control of synthetic images for data augmentation purposes. Synthetic images are not always realistic and may contain misleading features that distort data distribution when mixed with real images. As a result, the effectiveness of synthetic images in medical image recognition systems cannot be ensured when they are randomly added without quality assurance. A reinforcement learning-based synthetic sample selection approach is proposed in which synthetic images containing reliable and informative features are chosen.

However, none of the above approaches consider the extreme case of zero-shot learning. In the case of zero-shot learning, synthetic features are often biased towards seen classes, and the generated synthetic features do not represent the true distribution. Training a model to generate realistic-looking features will produce features similar to the training distribution without any guarantee on the effect on zero-shot evaluation. Therefore, we propose a data valuation method for synthetic features based on classification performance rather than their realism.

3. Methodology

The overall framework of the proposed method is visually depicted in Figure 2, which provides an illustrative overview of the various components employed. In this section, we delve deeper into the individual constituents of the model and explore in detail the novel SPOT selector that has been put forth. It is imperative to note that the proposed pipeline is model and data-agnostic, which implies that the

choice of classifier model and network backbone is dependent solely on the feature-generating framework itself.

The development of the SPOT selector draws significant inspiration from the synthetic sample selector introduced in [60]. However, our approach seeks to tackle the more challenging task of zero-shot learning, where data from unseen classes is limited. Furthermore, we demonstrate that training the selector using data from seen classes can enhance the selection of better features for data from unseen classes. This approach represents a significant contribution and serves to bridge the gap between the seen and unseen class data.

3.1. Feature Generating Network (FGN)

As previously highlighted, it is worth noting that the proposed pipeline is entirely independent of the feature-generating approach itself. This aspect renders the framework highly versatile and adaptable to a diverse range of feature-generating models, which may be employed in place of the FGN utilized in this study.

Examples of alternate feature-generating models that could be employed include the WGAN [52] and Cycle-WGAN [6]. The utilization of such models would permit the proposed pipeline to be seamlessly integrated into a broader range of applications and extend the reach and scope of the framework. The flexibility afforded by this design decision represents a critical feature of the proposed pipeline and enables the framework to be readily adapted to a range of diverse use cases as shown with consistent improvements in multiple image and video datasets.

3.2. Selector

The reasoning behind selecting the particular selector is rooted in the interdependence of features among the potential images. We hypothesise that the sequence in which the features are generated is not completely autonomous, as the later additions must differentiate themselves from the earlier ones in order to ensure diversity across the entire set of augmented training data. We use a transformer-based architecture [5] to be our selector. The input is a feature vector of a dimension dependent on the FGN used. The goal of the selector is to tell us if the generated feature vector is good for classification performance. The selector takes in a feature vector and outputs a score that tells us how good that feature vector is for classification. To do this, the selector outputs a binary action: select or not select. However, we do not have ground truth to tell us how good the generated feature is and hence optimizing the selector is not trivial.

To address the possibility of a relationship between augmented images without relying heavily on sequential assumptions, we utilize the self-attention mechanism through the implementation of the transformer [5] model as our selector. The transformer architecture eliminates all recurrent

structures, requiring feature vectors to be combined with their positional embeddings using sinusoidal functions prior to being input into the encoder layer of the transformer. The primary component of the transformer encoder is the multi-head attention block, comprised of n self-attention layers, where n denotes the number of heads. In each self-attention layer, input features are projected to three separate feature spaces - query Q , key K , and value V - by multiplying learnable weight matrices. The resulting attention map is obtained through the following process:

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (1)$$

Each head of the multi-head attention block represents a distinct projected feature space for the input, achieved by multiplying the same input embedding with different weight matrices. These separate outputs are then concatenated to form the final attention map, which is expressed as:

$$MultiHead(Q, K, V) = Concat(head_1, \dots, head_h)W^O \quad (2)$$

$$head_i = Attention(QW_i^Q, KW_i^K, VW_i^V) \quad (3)$$

Once we obtain the attention map, the context vector is then fed to the feed-forward layer as follows:

$$F(x) = max(((0, xW_1 + b_1)W_2 + b_2)W_3 + b_3 \dots)W_n + b_n \quad (4)$$

Given that the objective of the selector is to produce a binary action for every input feature vector, the decoder for the transformer model is a linear layer that functions as the policy network. Overall, the use of the transformer as the selector within our image selection framework, based on reinforcement learning, is beneficial due to its self-attention mechanism that effectively captures the interdependencies among the input feature vectors. We conducted a thorough ablation study regarding the selection of the selector and this can be seen in Section 4.2.

To optimize the selector, we turn to reinforcement learning as this is a common solution [13, 62]. In particular, we use proximal policy optimization [42]. Details are explained in the next section.

3.3. Proximal Policy Optimization

As previously stated, we turn to reinforcement learning approach to update the selector model. A proficient policy gradient method is fundamental to effectively utilize reward feedback as input to the selector in the reinforcement learning process. Among various policy gradient algorithms, Proximal Policy Optimization (PPO) [42] has

gained popularity due to its computational efficiency and satisfactory performance, surpassing previous approaches like TRPO [41]. Additionally, PPO alleviates the instability encountered during RL training. PPO achieves comparable performance with reduced complexity by replacing the KL convergence constraint enforced in TRPO with a clipped probability ratio between the current and previous policy within a small interval around 1. At each time step t , with A_θ representing the advantage function, the objective function is defined as follows:

$$L(\theta) = E[\min(\gamma\theta(t)A_\theta(s_t, a_t), \text{clip}(\gamma\theta(t), 1 - \epsilon, 1 + \epsilon)A_\theta(s_t, a_t))] \quad (5)$$

Here, $A_\theta(s_t, a_t) = Q_\theta(s_t, a_t) - V_\theta(s_t, a_t)$. As a component of the transformer output, the learned state-value $V_\theta(s_t, a_t)$ serves as a baseline for the q-value to mitigate the variance of rewards during the training process. The probability of actions is denoted by π . The q-value at time t , $Q_\theta(s_t, a_t)$, is defined as a smoothed version of the maximum validation accuracy observed among the last five epochs in the classification task. As our target tasks are trained on the seen class data and needs to generalize to the unseen class data, it is crucial to obtain a robust estimation of the reward’s changing pattern. To achieve this, we employ the Exponential Moving Average (EMA) algorithm to smooth the original reward curve. Thus, the final reward at time t is obtained as follows:

$$\hat{Q}_\theta(s_t, a_t) = \begin{cases} Q_\theta(s_t, a_t), & t = 1 \\ \alpha\hat{Q}_\theta(s_{t-1}, a_{t-1}) + (1 - \alpha)Q_\theta(s_t, a_t), & t > 1 \end{cases} \quad (6)$$

Drawing inspiration from the concept of importance sampling, the weight assigned to the current policy is influenced by earlier policies. The probability ratio between the previous and current policies, denoted by $\gamma_\theta(t)$, is mathematically defined as:

$$\gamma_\theta(t) = \frac{\pi_\theta(a_t | s_t)}{\pi_\theta(a_{t-1} | s_{t-1})} \quad (7)$$

Here, $a_t \in \mathbb{R}^{N \times 2}$ refers to the number of synthetic samples in the candidate pool. If at any given timestep t , $a_i(t) = 0$ then i is discarded. Else, it is added to the original training set.

4. Experimental Analysis

4.1. Implementation Details

Since we propose a plug-and-play component to feature-generating networks, the backbone and technical details




Figure 3. Ablation comparison of different combinations of RL algorithms with selector choices (GRU or Transformer). ‘G’ = GRU, ‘Tr’ = Transformer, ‘R’ = REINFORCE, ‘T’ = TRPO and ‘P’ = PPO.

follow the exact same implementation that the feature-generating model uses. Here, we talk about the technical details of running SPOT.

The candidate features generated by the feature-generating framework are passed into the selector network which is a 8-layer encoder having a 8-head multi-head attention block. The output is a vector A_θ which is the action vector and a value vector V_θ that is used together to calculate the reward for the policy gradient algorithm.

The classifier network depends on the feature-generating network being used since our proposed method is model agnostic. Similar to [13, 60] we use the EMA-smoothed validation accuracy obtained from the last 5 epochs as the reward with $\alpha=0.5$ when using the classifier on the validation set. As long as this average is increasing, we continue updating our policy. The policy function π is obtained from the softmax layer of the seen-class classifier model.

ϵ in Eq. 4 is set to 0.15 (see Ablations for empirical comparison), this helps to set an upper and lower bound at the current time step t and previous one $t - 1$ for the ratio of the policy function. The number of selected synthetic features are dependent on the selector and varies according to model and dataset. However, we set the learning rate to be fixed for the PPO at $2e - 04$.

4.2. Ablation Study

We have made a few choices with regards to the hyperparameter selection and choice of RL optimization algorithms and in this section, we show empirical reasons why the choices were made. Figure 3 shows the performance differences when using different RL optimization algorithms to modify the selector. Similar to [60], we also consider alternative choices such as GRU and LSTMs for the selector. In terms of RL algorithms, we compare to REINFORCE [46] and TRPO [41].

Model	CUB	AWA	SUN	FLO
WGAN	57.3	68.2	60.8	67.2
WGAN + SPOT	60.7	71.1	63.3	69.9
Cycle-WGAN	57.8	65.6	59.7	68.6
Cycle-WGAN + SPOT	61.1	69.7	62.5	70.9
f-VAEGAN	61.0	71.1	64.7	67.7
f-VAEGAN + SPOT	62.8	72.7	66.0	69.2
CMC-GAN	61.4	71.4	63.7	69.8
CMC-GAN + SPOT	62.9	73.1	65.1	71.9

Table 1. Results on zero-shot image classification using recent feature-generating frameworks.

4.3. Images

4.3.1 Datasets and Evaluation Protocol

Our method is evaluated on four challenging benchmark datasets, namely AWA [27], CUB (Caltech UCSD Birds 200) [49], SUN (SUN Attribute) [55], and FLO [35]. CUB and SUN are fine-grained datasets, while AWA and FLO are coarse-grained datasets. We adopt the same seen/unseen splits and class embeddings to ensure consistency with previous work as in [53]. AWA1 contains 30,475 instances across 50 categories. CUB comprises 11,788 images of 200 bird classes (150/50 for seen/unseen classes) with 312 attributes. SUN contains 14,340 images from 717 scene classes (645/72 for seen/unseen classes) with 102 attributes. FLO consists of 8,189 images from 102 flower classes with an 82/20 class split for seen and unseen classes, respectively. These datasets are widely used in the literature, enabling a direct comparison of our results with those of previous studies.

4.3.2 Zero-Shot Learning

We compare strictly with recent state-of-the-art feature-generating approaches and as such compare to WGAN [52], Cycle-WGAN [6], f-VAEGAN [54], NereNet [29] and CMC-GAN [58]. Table 1 shows the results. We see consistent gains with increase of up to 4.1% when we add the proposed SPOT to any of the models.

4.3.3 Generalized Zero-Shot Learning

We perform a much more extensive comparison in the generalized setting as this is where most feature generating frameworks perform experiments. We compare against WGAN [52], Cycle-WGAN [6], f-VAEGAN [54], CMC-GAN [58], NereNet [29], FREE [3] and DAA [64]. We see consistent improvements on the unseen class accuracies as this is where selected features make a difference. We see gains of up to 3.3% on the unseen class accuracies. As a result, there is consistent improvement on the harmonic mean

of the seen and unseen class accuracies as well.

4.4. Videos

4.4.1 Datasets and Evaluation Protocol

For videos, we use the widely adopted Olympic Sports [34], HMDB-51 [25], and UCF-101 [45] datasets to evaluate our method for zero-shot action recognition and compare it against recent state-of-the-art feature generating models [20, 30, 52]. The aforementioned datasets comprise 783, 6766, and 13320 videos and are associated with 16, 51, and 101 classes, respectively. To enable comparison with existing works [20, 30, 32, 33, 52], we adopt the widely used 50/50 splits proposed by Xu et al. [56], where half of the classes are considered as seen and the other half as unseen. We report the average accuracy and standard deviation over 10 independent runs, following previous approaches.

Moreover, we extend our evaluation to include TruZe [16], which was recently introduced to address the issue of overlapping classes between the pre-training dataset (Kinetics [2]) and the unseen classes in zero-shot settings. The TruZe split acknowledges the presence of such overlapping classes, which contradicts the fundamental assumption that the unseen classes have not been previously seen.

4.4.2 Zero-Shot Learning

Table 3 shows the effect of using our proposed SPOT selector to enhance the performance of state-of-the-art feature-generating frameworks on the zero-shot setting. We compare with the most recent best-performing methods, which include the Bi-Dir GAN [32], GGM [33], OD [30], WGAN [52] and FFG [20] (fine-grained feature generation framework).

We observe that the proposed method consistently outperforms all approaches across all datasets by gains of up to 4.5%.

4.4.3 Generalized Zero-Shot Learning

We evaluate SPOT on the generalized setting where at test time both seen and unseen class samples are used. Table 4 shows the results, with the harmonic mean of the seen and unseen class accuracies. The proposed SPOT selector consistently improves all approaches across all datasets by gains of up to 4.2%.

4.5. Results on TruZe

We also evaluate on the stricter TruZe [16] split that ensures no overlap between the pre-trained model and test classes. Results are shown in Table 5. We only evaluate on OD and WGAN as these are the two feature-generating approaches that have results reported on the TruZe split. Again, we see that using SPOT for selection consistently

Model	CUB			AWA1			SUN			FLO		
	S	U	H	S	U	H	S	U	H	S	U	H
WGAN	43.7	57.7	49.7	57.9	61.4	59.6	42.6	36.6	39.4	59.0	73.8	65.6
WGAN+SPOT	44.1	60.9	51.1	58.6	64.9	61.6	42.8	39.1	40.9	59.3	75.9	66.6
Cycle-WGAN	46.0	60.3	52.2	56.4	63.5	59.7	48.3	33.1	39.2	59.1	71.1	64.5
Cycle-WGAN+SPOT	46.5	62.9	53.5	56.9	66.1	61.1	48.1	36.2	41.3	59.4	74.4	66.1
f-VAEGAN	48.4	60.1	53.6	57.6	70.6	63.5	45.1	38.0	41.3	56.8	74.9	64.6
f-VAEGAN+SPOT	48.8	62.8	54.9	57.9	73.3	64.7	45.5	41.1	43.2	57.0	77.2	65.6
CMC-GAN	52.6	65.1	58.2	63.2	70.6	66.7	48.2	40.8	44.2	64.5	80.2	71.5
CMC-GAN+SPOT	53.1	66.7	59.1	63.3	73.8	68.1	48.9	44.1	46.4	64.6	82.8	72.6
NereNET	51.0	56.5	53.6	-	-	-	45.7	38.1	41.6	-	-	-
NereNET+SPOT	51.3	58.4	54.6	-	-	-	45.9	40.4	43.0	-	-	-
FREE	55.7	59.9	57.7	62.9	69.4	66.0	47.4	37.2	41.7	67.4	84.5	75.0
FREE+SPOT	55.5	62.2	58.6	63.1	72.1	67.3	47.8	39.9	43.5	67.8	86.3	75.9
DAA	66.1	65.5	65.8	64.3	76.6	69.9	47.8	38.7	42.8	-	-	-
DAA+SPOT	66.3	67.7	67.0	64.6	77.9	70.6	48.1	40.3	43.8	-	-	-

Table 2. Results on generalized zero-shot image classification on 4 challenging benchmarks.

Method	Olympics	HMDB51	UCF101
Bi-Dir GAN [32]	53.2 ± 10.5	21.3 ± 3.2	24.7 ± 3.7
Bi-Dir GAN [32] + SPOT	56.6 ± 10.1	25.1 ± 3.4	27.7 ± 3.5
GGM [33]	57.9 ± 14.1	20.7 ± 3.1	24.5 ± 2.9
GGM [33] + SPOT	62.4 ± 12.4	25.1 ± 2.8	27.4 ± 2.5
OD [30]	65.9 ± 8.1	30.2 ± 2.7	38.3 ± 3.0
OD [30] + SPOT	68.7 ± 7.5	34.4 ± 2.2	40.9 ± 2.6
WGAN [52]	64.7 ± 7.5	29.1 ± 3.8	37.5 ± 3.1
WGAN [52] + SPOT	68.1 ± 7.1	33.8 ± 2.4	40.6 ± 2.4
FFG [20]	-	32.4 ± 2.3	27.6 ± 2.4
FFG [20] + SPOT	-	35.9 ± 2.5	30.9 ± 2.2

Table 3. Results on zero-shot action recognition on the Olympics, HMDB51 and UCF101 datasets.

Method	Olympics	HMDB51	UCF101
Bi-Dir GAN [32]	44.2 ± 11.2	7.5 ± 2.4	22.7 ± 2.5
Bi-Dir GAN [32] + SPOT	48.4 ± 10.3	10.9 ± 3.1	25.1 ± 3.8
GGM [33]	52.4 ± 12.2	20.1 ± 2.1	23.7 ± 1.2
GGM [33] + SPOT	55.3 ± 11.9	23.4 ± 2.3	27.1 ± 3.1
WGAN [52]	59.9 ± 5.3	32.7 ± 3.4	44.4 ± 3.0
WGAN [52] + SPOT	62.4 ± 5.5	34.4 ± 2.9	46.2 ± 2.6
OD [30]	66.2 ± 6.3	36.1 ± 2.2	49.4 ± 2.4
OD [30] + SPOT	69.1 ± 6.5	38.2 ± 2.5	51.8 ± 2.5
FFG [20]	-	37.4 ± 1.9	40.4 ± 2.2
FFG [20] + SPOT	-	39.8 ± 1.4	42.8 ± 1.7

Table 4. Results on generalized zero-shot setting. Reported results are the harmonic mean of the seen and unseen class accuracies.

improves the performance of the feature-generating framework.

Method	UCF101		HMDB51	
	ZSL	GZSL	ZSL	GZSL
WGAN	22.5	36.3	21.1	31.8
WGAN + SPOT	25.3	39.1	23.8	33.3
OD	22.9	42.4	21.7	35.5
OD + SPOT	25.5	44.1	24.0	37.1

Table 5. Results on TruZe. We report the mean class accuracy for zero-shot; for generalized zero-shot, we report the harmonic mean of seen and unseen class accuracies.

5. Conclusion

In conclusion, although generative techniques have made significant progress in transforming traditional GZSL to fully supervised learning, they often generate redundant synthetic features, which can lead to reduced accuracy. To overcome this limitation, we have proposed an approach for synthetic feature selection using reinforcement learning, which involves training a transformer-based selector using proximal policy optimization (PPO) to select synthetic features based on the validation classification accuracy of seen classes as the reward. Our proposed method is model-agnostic and data-agnostic and hence is suitable for images and videos. The experimental results of our approach demonstrate its superiority over existing feature-generating methods, with improved overall performance observed across multiple benchmarks. Overall, our approach represents a significant contribution towards addressing the issue of synthetic feature redundancy in GZSL, and we believe that it has the potential to be widely applied in real-world scenarios.

References

- [1] Sanjeev Arora and Yi Zhang. Do gans actually learn the distribution? an empirical study. *arXiv preprint arXiv:1706.08224*, 2017. [2](#)
- [2] J. Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. [6](#)
- [3] Shiming Chen, Wenjie Wang, Beihao Xia, Qinmu Peng, Xinge You, Feng Zheng, and Ling Shao. Free: Feature refinement for generalized zero-shot learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 122–131, 2021. [6](#)
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [1](#)
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [4](#)
- [6] Rafael Felix, Ian Reid, Gustavo Carneiro, et al. Multi-modal cycle-consistent generalized zero-shot learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 21–37, 2018. [2](#), [4](#), [6](#)
- [7] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26, 2013. [1](#), [2](#)
- [8] Chuang Gan, Ming Lin, Yi Yang, Gerard De Melo, and Alexander G Hauptmann. Concepts not alone: Exploring pairwise relationships for zero-shot video activity recognition. In *Thirtieth AAAI conference on artificial intelligence*, 2016. [2](#)
- [9] Chuang Gan, Ming Lin, Yi Yang, Yueting Zhuang, and Alexander G Hauptmann. Exploring semantic inter-class relationships (sir) for zero-shot action recognition. In *Proceedings of the National Conference on Artificial Intelligence*, 2015. [2](#)
- [10] Chuang Gan, Tianbao Yang, and Boqing Gong. Learning attributes equals multi-source domain generalization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 87–97, 2016. [2](#)
- [11] Junyu Gao, Tianzhu Zhang, and Changsheng Xu. 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*, volume 33, pages 8303–8311, 2019. [2](#)
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [1](#)
- [13] Shreyank N Gowda, Marcus Rohrbach, Frank Keller, and Laura Sevilla-Lara. Learn2augment: Learning to composite videos for data augmentation in action recognition. In *European Conference on Computer Vision*, pages 242–259. Springer, 2022. [3](#), [4](#), [5](#)
- [14] Shreyank N Gowda, Marcus Rohrbach, and Laura Sevilla-Lara. Smart frame selection for action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1451–1459, 2021. [3](#)
- [15] Shreyank N Gowda, Laura Sevilla-Lara, Frank Keller, and Marcus Rohrbach. Cluster: clustering with reinforcement learning for zero-shot action recognition. In *European Conference on Computer Vision*, pages 187–203. Springer, 2022. [2](#)
- [16] Shreyank N Gowda, Laura Sevilla-Lara, Kiyoon Kim, Frank Keller, and Marcus Rohrbach. A new split for evaluating true zero-shot action recognition. *arXiv preprint arXiv:2107.13029*, 2021. [6](#)
- [17] Shreyank N Gowda and Chun Yuan. Colornet: Investigating the importance of color spaces for image classification. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part IV 14*, pages 581–596. Springer, 2019. [1](#)
- [18] Wenzhong Guo, Jianwen Wang, and Shiping Wang. Deep multimodal representation learning: A survey. *IEEE Access*, 7:63373–63394, 2019. [2](#)
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#)
- [20] Mingyao Hong, Xinfeng Zhang, Guorong Li, and Qingming Huang. Fine-grained feature generation for generalized zero-shot video classification. *IEEE Transactions on Image Processing*, 2023. [6](#), [7](#)
- [21] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. [1](#)
- [22] Ali Jahanian, Lucy Chai, and Phillip Isola. On the” steerability” of generative adversarial networks. In *International Conference on Learning Representations*. [2](#)
- [23] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J Spanos. Towards efficient data valuation based on the shapley value. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1167–1176. PMLR, 2019. [3](#)
- [24] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3174–3183, 2017. [1](#)
- [25] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 International Conference on Computer Vision*, pages 2556–2563. IEEE, 2011. [6](#)
- [26] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *2009 IEEE Conference on Com-*

- puter Vision and Pattern Recognition*, pages 951–958. IEEE, 2009. 2
- [27] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):453–465, 2013. 6
- [28] Chung-Ching Lin, Kevin Lin, Lijuan Wang, Zicheng Liu, and Linjie Li. Cross-modal representation learning for zero-shot action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19978–19988, 2022. 2
- [29] Jingren Liu, Haoyue Bai, Haofeng Zhang, and Li Liu. Near-real feature generative network for generalized zero-shot learning. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2021. 2, 6
- [30] Devraj Mandal, Sanath Narayan, Sai Kumar Dwivedi, Vikram Gupta, Shuaib Ahmed, Fahad Shahbaz Khan, and Ling Shao. Out-of-distribution detection for generalized zero-shot action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9985–9993, 2019. 2, 3, 6, 7
- [31] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013. 2
- [32] Ashish Mishra, Anubha Pandey, and Hema A Murthy. Zero-shot learning for action recognition using synthesized features. *Neurocomputing*, 390:117–130, 2020. 2, 3, 6, 7
- [33] Ashish Mishra, Vinay Kumar Verma, M Shiva Krishna Reddy, S Arulkumar, Piyush Rai, and Anurag Mittal. A generative approach to zero-shot and few-shot action recognition. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 372–380. IEEE, 2018. 2, 3, 6, 7
- [34] Juan Carlos Niebles, Chih-Wei Chen, and Li Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *European conference on computer vision*, pages 392–405. Springer, 2010. 6
- [35] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008. 6
- [36] Mark Palatucci, Dean Pomerleau, Geoffrey E Hinton, and Tom M Mitchell. Zero-shot learning with semantic output codes. *Advances in neural information processing systems*, 22, 2009. 1
- [37] Yijun Qian, Lijun Yu, Wenhe Liu, and Alexander G Hauptmann. Rethinking zero-shot action recognition: Learning from latent atomic actions. In *European Conference on Computer Vision*, pages 104–120. Springer, 2022. 2
- [38] Jie Qin, Li Liu, Ling Shao, Fumin Shen, Bingbing Ni, Jiaxin Chen, and Yunhong Wang. Zero-shot action recognition with error-correcting output codes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2833–2842, 2017. 2
- [39] Marcus Rohrbach, Michaela Regneri, Mykhaylo Andriluka, Sikandar Amin, Manfred Pinkal, and Bernt Schiele. Script data for attribute-based recognition of composite activities. In *ECCV*, 2012. 2
- [40] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8247–8255, 2019. 2
- [41] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015. 5
- [42] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 2, 4
- [43] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1
- [44] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. *Advances in neural information processing systems*, 26, 2013. 1, 2
- [45] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 6
- [46] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999. 5
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [48] Vinay Kumar Verma, Gundeep Arora, Ashish Mishra, and Piyush Rai. Generalized zero-shot learning via synthesized examples. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4281–4289, 2018. 2
- [49] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 6
- [50] Zhaoxuan Wu, Yao Shu, and Bryan Kian Hsiang Low. Davinz: Data valuation using deep neural networks at initialization. In *International Conference on Machine Learning*, pages 24150–24176. PMLR, 2022. 3
- [51] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2251–2265, 2018. 1
- [52] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5542–5551, 2018. 2, 3, 4, 6, 7
- [53] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning—the good, the bad and the ugly. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4582–4591, 2017. 6

- [54] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata. f-vaegan-d2: A feature generating framework for any-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10275–10284, 2019. 6
- [55] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010. 6
- [56] Xun Xu, Timothy Hospedales, and Shaogang Gong. Transductive zero-shot action recognition by word-vector embedding. *International Journal of Computer Vision*, 123(3):309–333, 2017. 2, 6
- [57] Xun Xu, Timothy M Hospedales, and Shaogang Gong. Multi-task zero-shot action recognition with prioritised data augmentation. In *European Conference on Computer Vision*, pages 343–359. Springer, 2016. 2
- [58] Fu-En Yang, Yuan-Hao Lee, Chia-Ching Lin, and Yu-Chiang Frank Wang. Semantics-guided intra-category knowledge transfer for generalized zero-shot learning. *International Journal of Computer Vision*, pages 1–15, 2023. 2, 6
- [59] Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. Zerogen: Efficient zero-shot learning via dataset generation. *arXiv preprint arXiv:2202.07922*, 2022. 2
- [60] Jiarong Ye, Yuan Xue, L Rodney Long, Sameer Antani, Zhiyun Xue, Keith C Cheng, and Xiaolei Huang. Synthetic sample selection via reinforcement learning. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I 23*, pages 53–63. Springer, 2020. 3, 4, 5
- [61] Meng Ye and Yuhong Guo. Zero-shot classification with discriminative semantic representation learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7140–7148, 2017. 1
- [62] Jinsung Yoon, Serkan Arik, and Tomas Pfister. Data valuation using reinforcement learning. In *International Conference on Machine Learning*, pages 10842–10851. PMLR, 2020. 3, 4
- [63] Yunlong Yu, Zhong Ji, Jungong Han, and Zhongfei Zhang. Episode-based prototype generating network for zero-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14035–14044, 2020. 2
- [64] Xiaojie Zhao, Yuming Shen, Shidong Wang, and Haofeng Zhang. Generating diverse augmented attributes for generalized zero shot learning. *Pattern Recognition Letters*, 2023. 6