

ProtoGAN: Towards Few Shot Learning for Action Recognition

Sai Kumar Dwivedi

Mercedes-Benz R&D India

saikumar.dwivedi@daimler.com

Vikram Gupta

Mercedes-Benz R&D India

vikram.gupta@daimler.com

Rahul Mitra

IIT Bombay

rmitter@cse.iitb.ac.in

Shuaib Ahmed

Mercedes-Benz R&D India

shuaib.ahmed@daimler.com

Arjun Jain

IIT Bombay, Axogyan AI

arjunjain@gmail.com

Abstract

Few-shot learning (FSL) for action recognition is a challenging task of recognizing novel action categories which are represented by few instances in the training data. In a more generalized FSL setting (G-FSL), both seen as well as novel action categories need to be recognized. Conventional classifiers suffer due to inadequate data in FSL setting and inherent bias towards seen action categories in G-FSL setting. In this paper, we address this problem by proposing a novel ProtoGAN framework which synthesizes additional examples for novel categories by conditioning a conditional generative adversarial network with class prototype vectors. These class prototype vectors are learnt using a Class Prototype Transfer Network (CPTN) from examples of seen categories. Our synthesized examples for a novel class are semantically similar to real examples belonging to that class and is used to train a model exhibiting better generalization towards novel classes. We support our claim by performing extensive experiments on three datasets: UCF101, HMDB51 and Olympic-Sports. To the best of our knowledge, we are the first to report the results for G-FSL and provide a strong benchmark for future research. We also outperform the state-of-the-art method in FSL for all the aforementioned datasets.

1. Introduction

Action recognition has been a long-standing and actively pursued problem in the computer vision community due to its practical applications in areas such as surveillance, semantic video retrieval and multimedia mining. Recently, Convolutional Neural Network (CNN) based methods [3, 22, 24] have achieved tremendous success in recognizing actions from videos in supervised learning paradigm. However, the performance of these methods deteriorates drastically [25] when recognizing action classes that are not

adequately represented in their training data (*novel* classes). This problem limits the deployment of these methods in real world applications where the number of action classes to be recognized increases rapidly with use cases. In certain cases, the number of samples for novel classes are quite few for even traditional data augmentation techniques [11] to work. Therefore, systems with more advanced learning paradigm like *Few Shot Learning* [8], which learns to recognize novel classes from only a few examples (*few-shots*), have come into prominence.

Few-shot learning problem can be classified into two broad settings [8] based on their evaluation protocols: standard FSL (*FSL*) and the more realistic Generalized-FSL (*G-FSL*). *FSL* focuses on recognition of only *novel* classes during evaluation whereas under *G-FSL*, a combination of both *seen* (adequately represented at training) and *novel* classes is considered. The presence of *seen* classes during evaluation can incorrectly bias a classifier to predict a *seen* class when the input belongs to a *novel* class. Hence, the *G-FSL* setting is considered more challenging. There have been several approaches [29, 19, 8] to tackle *FSL* for image classification. The most notable among them use *meta-learning*, *representation learning* and *generative modelling*. *Meta-learning* mimics the few-shot inference time scenario during training, *representation learning* tries to learn the similarity of new samples to existing few-shots and *generative modelling* augments the novel classes with synthetic data. However, similar approaches for action recognition remain quite under-explored.

Generative methods like Conditional GAN (CGAN) [13] synthesizes additional data for *novel* classes using a conditioning element. In image classification, Antoniou et al. [1] explored data augmentation GAN which uses the few-shot samples directly as a conditioning element to generate synthetic data. Zhang et al. [29] use statistics from the samples as an alternative for the conditioning element. However, when the number of shots is less, conditioning element in

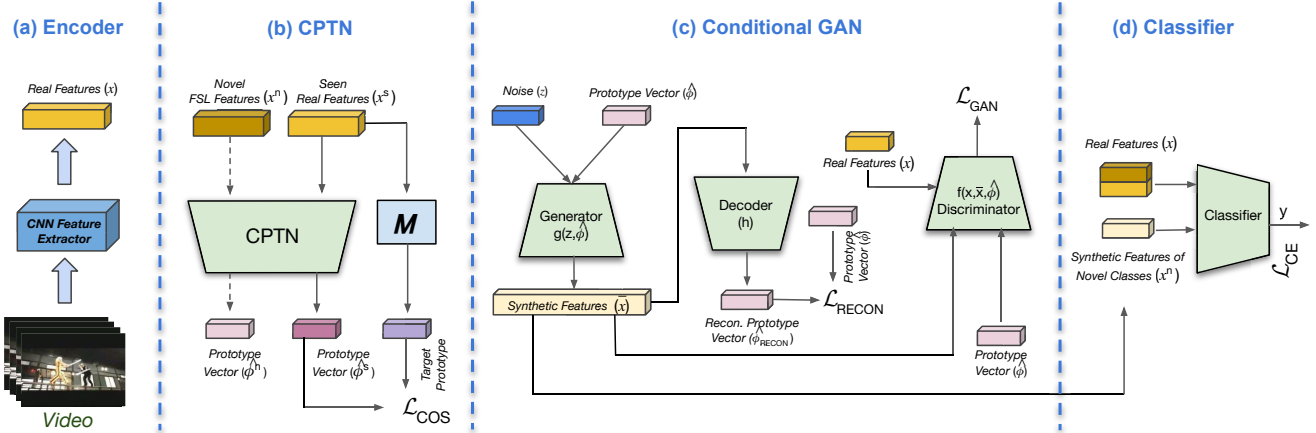


Figure 1: Illustration of our proposed *ProtoGAN* framework. (a) Encoder extracts video features using spatio-temporal CNN (b) *Class Prototype* vector ϕ^n for a novel class is learnt by mapping video features of seen data x^n in CPTN to target formed through feature aggregator function \mathcal{M} using cosine similarity loss L_{COS} . The solid line denotes the training time path and dotted line denotes the inference path. (c) Learnt class prototype vector is used as a conditioning element in CGAN to synthetically generate more samples using generator g and discriminator f using adversarial loss L_{GAN} . Decoder h is used to reconstruct the prototype vector from generator’s output using L_{RECON} to ensure discriminative properties. (d) Classifier takes real seen features, few-shot features and synthetic features of novel classes to train with L_{CE} loss

these methods does not represent the class semantics.

To this end, we propose the *ProtoGAN* framework which conditions CGAN on a *class-prototype* vector to synthesize additional video features for the action recognition classifier. Class-prototype vector is *learnt* through a feature aggregator network called *Class Prototype Transfer Network* (CPTN). The synthetic features generated using our learned conditioning element for a class are semantically similar to actual features of videos belonging to that class. We justify our claims by performing extensive experiments on three publicly available datasets namely *UCF101* [20], *HMDB* [12] and *Olympic-Sports* [16] under G-FSL and FSL settings. We focus the necessary ablations in our framework on G-FSL as it presents a more practical setting applicable for real world scenarios. The key contributions of our paper can be summarized as follows:

1. Introducing a *learned class-prototype* vector for videos capturing class level semantics which is subsequently used as conditioning element for a CGAN to generate relevant synthetic features for novel classes.
2. To the best of our knowledge, we are the first to demonstrate results for Generalized-Few Shot Learning (G-FSL) on publicly available action recognition datasets and hence provide a strong benchmark for future research.
3. Outperform state-of-the-art method under the standard FSL setting for the three aforementioned datasets and across different number of samples or shots.

2. Related Work

FSL approaches in image classification literature are predominantly based on meta-learning [18, 29, 6]. Under the umbrella of meta-learning, various approaches like metric learning [19, 23, 21], learning optimization [17] and learning to initialize and finetune [6] are proposed. Metric learning methods learn the similarity between images [23] to classify samples of novel classes at inference based on nearest neighbors with labelled samples. Learning optimization involves memory based networks like LSTM [26] or memory augmented networks [15] which aim to replace stochastic gradient descent optimizer with the help of external memory. Learning to initialize and fine-tune aims to make minimal changes in the network when adapting to new task of classifying novel classes with few examples. Although, all the works mentioned above attempt to solve the broad problem of few-shot learning, many of them do not consider the full potential of seen classes as they learn the models in k-shot n-way method [15]. In [8], the authors propose a method of representation learning of visual features from the seen classes and translate it to novel classes for images. Recently, GAN based synthesizing of novel classes is gaining importance [1] where the novel class is augmented with synthetic features. In [29], a generic framework to augment data with conditioning GAN on sample mean class-representative vector is explored. A detailed overview of existing methods of FSL in images is given in [4].

In contrary to FSL in images, video-based FSL methods

such as action recognition have got less attention. In [30], authors have proposed a compound memory architecture which transforms a variable length video sequence into a fixed length matrix helping in few shot classification. A method for action localization in FSL setting is explored in [27]. Attribute-based feature generation for unseen classes from GAN by using Fisher vector representation was explored in zero-shot learning in [28]. Authors in [14] used Gaussian based generative approach to augment data for novel classes, where each action is represented in the visual space as a probability distribution. Since, the experimental setup in [14] includes classifying all novel classes together in contrary to k-shot n-way setting, this forms a strong baseline to compare with our proposed framework in FSL setting.

Our approach: To avoid classifier’s bias towards seen classes, unlike the approaches mentioned above, we propose learning a prototype vector which captures the semantics of a class by using features from seen classes. A CGAN is trained to learn the mapping from prototype vector to visual features. Additional data for novel classes is synthesized by using the predicted prototype vector from CPTN as conditioning element in CGAN.

3. Proposed Method

An overview of our proposed framework can be seen in Fig. 1 which can be broadly segregated into four major blocks: (a) *Encoder* for extracting video features (b) *Class-Prototype Transfer Network* for creating class prototype vector, (c) *CGAN* for generating synthetic features when conditioned on class-prototype vector, and (d) *Classifier* for predicting correct action class.

3.1. Preliminaries

Let $x, \bar{x}, y, \hat{\phi}$ represent real video features, synthetic video features, class labels and class prototype vectors, respectively. Let $\mathcal{S} = \{(x^s, y^s) \mid x^s \in \mathcal{X}^s, y^s \in \mathcal{Y}^s\}$ be the training set for seen classes where $x^s \in \mathbb{R}^{d_x}$ denotes the spatio-temporal features, y^s denotes the class labels in $\mathcal{Y}^s = \{y_1^s, \dots, y_S^s\}$ with S seen classes. The class prototype vector $\hat{\phi}^s \in \Phi^s \subset \mathbb{R}^{d_\phi}$ is calculated from x^s video features. Let the cardinality of each member in \mathcal{Y}^s be denoted by $\mathcal{K}^s = \{k_1^s, \dots, k_S^s\}$. Additionally, $\mathcal{N} = \{(x^n, y^n) \mid x^n \in \mathcal{X}^n, y^n \in \mathcal{Y}^n\}$ is available during training from novel categories where y^n is a class from a disjoint label set $\mathcal{Y}^n = \{y_1^n, \dots, y_N^n\}$ of N novel labels. Similarly, the cardinality of each member in \mathcal{Y}^n is denoted by $\mathcal{K}^n = \{k^1, \dots, k^N\}$ where $k^n \ll \min(K^s)$ and $\min()$ is defined as minimum value in the set. The class prototype vector of novel class $\hat{\phi}^n$ is inferred through CPTN. In GFSL, the task is to learn a classifier $Q_{gfsl} : \mathcal{X}^s \cup \mathcal{X}^n \rightarrow \mathcal{Y}^s \cup \mathcal{Y}^n$ and in FSL the task is to learn a classifier $Q_{fsl} : \mathcal{X}^n \rightarrow \mathcal{Y}^n$.

3.2. Class Prototype Transfer Network (CPTN)

Generation of synthetic samples for a class using a CGAN [13] requires a conditioning element capturing the semantics of that class. When the number of samples belonging to the class is high, statistical methods like *mean* of the samples tend to accurately represent those semantics [7]. However, in the case of *novel* classes with few examples, this doesn’t hold true as such methods are susceptible to capture noisy intricacies due to lack of sufficient data. To this end, we learn a mapping function *Class-Prototype Transfer Network* (CPTN) through *feature aggregation* where features of each video belonging to a particular class is mapped to a lower dimension embedding serving as the representative of that class. We term this vector embedding as the *class-prototype* vector $\hat{\phi}$. We use seen classes data containing large number of samples to model this mapping and then transfer the mapping to novel categories.

Feature aggregation \mathcal{M} is a two step process where the sample representation \mathcal{R} is followed by dimensionality reduction \mathcal{T} as per the following equation:

$$\mathcal{M}_{x^s \rightarrow \hat{\phi}^s} = \mathcal{T}(\mathcal{R}(x^s)); \quad \mathcal{R}(x^s) = \frac{1}{n} \sum_{i=1}^n x_{ci}^s \quad (1)$$

where n is the number of instances of class c , \mathcal{R} is the mean of the samples, \mathcal{T} is a spatial dimensionality reduction function like average/max pool which is applied on the output of \mathcal{R} and $\hat{\phi}^s$ is the prototype vector of seen class formed through process \mathcal{M} . We have used average-pool in our approach. The feature aggregation step ensures that meaningful information is retained while getting rid of intricate and noisy details specific to individual videos. Cosine similarity loss \mathcal{L}_{COS} is used as a loss function for training the CPTN and is given by the equation:

$$\mathcal{L}_{COS} = \frac{\phi^s \cdot \hat{\phi}^s}{\|\phi^s\| \|\hat{\phi}^s\|} \quad (2)$$

where, ϕ^s is the predicted class-prototype vector and $\hat{\phi}^s$ is the target formed by the Eq. 1 for seen classes. Once the network is trained, the *prototype* vector $\hat{\phi}_c^n$ for a sample video belonging to a novel class c is obtained by passing its input video features x_c^n through the CPTN. For subsequent stages, we use $\hat{\phi}^n$ and $\hat{\phi}^s$ as class prototype vectors for novel and seen classes, respectively. The training framework is shown in (b) block of Fig. 1.

3.3. Conditional GAN

This module in our *ProtoGAN* framework involves a CGAN [13] to generate synthetic video features (\bar{x}^n) for novel classes by using the conditioning element $\hat{\phi}^n$ learned

UCF101									
	1-shot			3-shot			5-shot		
	<i>S</i>	<i>N</i>	<i>H</i>	<i>S</i>	<i>N</i>	<i>H</i>	<i>S</i>	<i>N</i>	<i>H</i>
<i>Base-Classifier</i>	82.7±0.6	38.4±2.0	52.4±1.9	83.1±5.4	61.2±3.0	70.3±1.2	88.2±0.9	68.8±1.9	77.3±0.9
<i>Heuristic-Proto</i>	88.0±1.0	46.3±2.2	60.7±2.0	92.0±0.7	62.1±1.9	74.1±1.2	94.0±1.4	68.7±1.5	79.3±0.7
<i>Sample-Proto</i>	88.0±1.9	45.9±2.4	60.3±2.2	92.0±0.9	61.9±1.9	74.0±1.3	94.4±0.8	68.4±1.3	79.3±0.8
<i>Learned-Proto</i>	75.3±1.3	52.3±2.2	61.7±1.6	87.7±0.8	64.9±1.7	74.6±1.0	90.5±0.9	71.3±1.2	79.7±0.8
HMDB									
Methods	1-shot			3-shot			5-shot		
	<i>S</i>	<i>N</i>	<i>H</i>	<i>S</i>	<i>N</i>	<i>H</i>	<i>S</i>	<i>N</i>	<i>H</i>
<i>Base-Classifier</i>	59.9±1.3	12.7±2.4	20.9±3.3	52.5±1.1	35.7±2.0	42.4±1.4	61.4±4.6	38.7±3.0	47.2±1.0
<i>Heuristic-Proto</i>	53.3±0.9	19.5±1.6	28.5±1.8	59.7±2.4	35.4±2.1	44.3±1.3	57.9±1.9	44.5±1.4	50.3±0.6
<i>Sample-Proto</i>	53.4±1.8	19.7±1.6	28.7±1.6	59.6±2.1	35.3±2.0	44.3±1.3	58.8±3.5	43.8±1.7	50.1±0.7
<i>Learned-Proto</i>	51.9±1.5	25.8±1.4	34.4±1.3	58.9±1.8	37.4±1.4	45.7±0.9	61.6±3.0	43.3±1.5	50.9±0.6
Olympic									
Methods	1-shot			3-shot			5-shot		
	<i>S</i>	<i>N</i>	<i>H</i>	<i>S</i>	<i>N</i>	<i>H</i>	<i>S</i>	<i>N</i>	<i>H</i>
<i>Base-Classifier</i>	94.5±3.1	16.9±3.1	28.6±4.4	93.2±3.2	41.1±2.9	56.9±2.9	92.0±3.3	54.9±3.2	68.7±2.0
<i>Heuristic-Proto</i>	95.5±3.2	18.5±2.9	30.9±4.0	95.3±3.5	41.2±2.8	57.5±2.6	95.0±3.7	54.8±3.2	69.4±2.1
<i>Sample-Proto</i>	95.5±3.7	18.6±2.9	31.0±4.8	95.1±3.4	41.9±2.8	58.1±2.6	94.8±3.9	55.7±3.4	70.0±2.3
<i>Learned-Proto</i>	94.8±3.4	20.9±2.6	34.1±4.7	93.5±3.0	46.0±2.7	61.5±2.2	93.2±3.4	59.2±3.4	72.2±2.0

Table 1: Comparison of Action Recognition accuracy (%) of our proposed framework with other baselines on UCF101, HMDB and Olympic datasets under G-FSL settings. *Base-Classifier* denotes vanilla classifier with standard augmentation. *Heuristic-Proto* and *Sample-Proto* denotes proposed baselines inspired from prototype-vectors used in image classification. *S*, *N*, *H* represents *seen*, *novel* and *harmonic mean*, respectively. We report mean accuracy from 20 different training runs and standard deviation is reported by \pm .

in the previous CPTN stage. The Wasserstein loss [2] is chosen over vanilla GAN loss as it provides more stability in training in low-data regime [28]. The Wasserstein loss in a CGAN between real features p_r and synthetically generated feature p_s is given by,

$$\mathcal{L}_{WGAN} = \mathbb{E}_{x \sim p_r} [f(x, \hat{\phi})] - \mathbb{E}_{z \sim p_z} [f(g(\bar{x}), \hat{\phi})] - \alpha \mathbb{E}[(\|\nabla_{\bar{x}} f(\hat{x}, \hat{\phi})\|_2 - 1)^2]; \quad (\bar{x} = g(\hat{\phi}, z)) \sim p_s \quad (3)$$

where x 's are *real* video features drawn from p_r . \hat{x} is a convex combination of x and \bar{x} , $\hat{\phi}$ is the conditioning element of a particular class, (f) is the discriminator, (g) is the generator, $z \sim \mathbf{N}(0, I)$ is the noise vector and α is the penalty coefficient. The first two terms in Equation. 3 approximate the Wasserstein distance and the third term is the penalty for constraining the gradient of (f) to have unit norm along the convex combination of real and generated pairs. The *Class-prototype* vector embedding $\hat{\phi}$ along with random noise z are sent as input the generator (g). It gener-

ates an output with a dimension same as that of input video features. The generated video features \bar{x} along with real video features x are passed to the discriminator (f) which is trained with an adversarial loss as per Eq. 3. Hence at equilibrium, the generator produces video features which are similar to the real features.

The generated features of a particular class y_i should be similar to the real features of the that class and farther away from the features of other classes. As it is compute intensive to find the closest match for a generated feature from a set of real features of that class, generated and real features are grouped to form unmatched pairs (for different classes). Cosine embedding loss $\mathcal{L}_{EMD} = \max(0, \cos(x_i, \bar{x}_j))$ (for $j \neq i$) is used to compute the distance for unmatched pairs.

To ensure that the generated features contain class semantics for subsequent supervised classification, similar to [5], an decoder (h) is used. It generates ϕ_{RECON} as an output in an attempt to reconstruct back $\hat{\phi}$ from $g_{\theta}(\hat{\phi}, z)$. We use cosine similarity loss L_{RECON} for the reconstruc-

tion. Thus, the net loss of CGAN is given by,

$$\mathcal{L}_{total} = \min_g \max_f \mathcal{L}_{WGAN} + \lambda \mathcal{L}_{RECON} + \gamma \mathcal{L}_{EMD} \quad (4)$$

where λ is the hyper-parameter for weighting the reconstruction loss and γ for embedding loss. Training framework of the current stage is shown in (c) block of Fig. 1.

3.4. Classifier

We utilize the generator (g) learned in the previous stage to produce additional video features for *novel* classes given their corresponding *class-prototype* ϕ_c^n obtained through the CPTN network. The classifier is trained in a supervised learning paradigm with the real features of seen and novel classes and synthetic features of novel classes. We subsequently train action recognition classifiers Q_{gfs} and Q_{fsl} for G-FSL and FSL settings, respectively, with Cross-Entropy loss \mathcal{L}_{CE} .

We made an insightful observation that generated samples having high reconstruction error (\mathcal{L}_{RECON}) are not suitable for training the classifier as they are semantically different than its corresponding class. Hence, we use a pruning method to remove all the synthetic features with high reconstruction loss. Features with low reconstruction loss tend to be more discriminative. This is because the decoder tries to reconstruct the class-prototype vector which is the representative of the class.

4. Experimental Setup

Details about the experimental setup are discussed in this section. Implementation details such as network architecture and training procedures are explained in Section. 4.1, dataset used for evaluation are described in Section. 4.3 and evaluation protocol followed for experiments under G-FSL and FSL setting are highlighted in Section. 4.2.

4.1. Implementation Details

- **Encoder** - We extracted video features from *C3D* network [22] pre-trained on Sports-1M dataset [9] where each video was divided into non-overlapping chunks of 16 frames and the mean of the *fc6* layer outputs in [22], of size 4096, was taken.
- **Class-Prototype Transfer Network (CPTN)** - A two-layer MLP was used with an input size of 4096, a hidden layer of 1024 and output layer (*class-prototype* vector embedding $\hat{\phi}$) of 128 neurons with *Sigmoid* as activation function. The model was trained with cosine similarity loss with ADAM [10] optimizer initialized with a learning rate of 0.005 and weight-decay of 0.0005. The training was run for 50 epochs.

- **CGAN** - The generator (g) is a two layer fully connected network with input size of 256 ($128(\hat{\phi}_c) + 128(noise)$) and output size of 4096. The decoder (h) used in cyclic-reconstruction is a two-layer fully connected with input and output size being 4096 and 128, respectively. The discriminator (f) is a two layer fully connected neural network with input dimension of 4096 and output dimension of 1 deciding whether the features are real or fake. Learning rate of 0.001 and with a weight decay of 0.0001 was used to train the CGAN for 25 epochs using ADAM optimizer. *LeakyReLU* was used as activation function for (g), (f) and (h). We set λ and γ in Eq. 4 to be 0.01 and 0.1, respectively. The gradient penalty parameter α for WGAN is set to 10.

- **Classifier** - The final action recognition classifier is a single layer MLP with input as 4096 and output as the number of classes. Cross-entropy loss was used with ADAM optimizer with a learning rate of 0.001. For each novel class, we generated twice the maximum number of samples present in a seen class, sort them in ascending order of their reconstruction loss and pick the top 50%.

4.2. Evaluation Protocols

To highlight the efficacy of our proposed framework, we thoroughly evaluate our method under G-FSL and FSL setting for all the action recognition datasets. A brief description of the evaluation protocol are as follows:

- **G-FSL setting** - In this setting, the test set consists of *seen* S and *novel* N classes and the model is evaluated based on their classification accuracy. The accuracy of seen, novel and harmonic mean H as per Eq. 5 of the two are reported.

$$Acc_h = \frac{2 \cdot Acc_s \cdot Acc_n}{Acc_s + Acc_n} \quad (5)$$

where Acc_s and Acc_n are accuracy of *seen* and *novel* classes, respectively.

- **FSL setting** - In this setting, the test split consists of only *novel* classes N .

In our experiments, under k -shot setting, k random samples were chosen from a set of available examples for each *novel* class. To eliminate bias towards any given set of samples and (*seen*, *novel*) class split, we repeat our experiments 20 times, with randomly chosen splits and samples. We report the mean results. For instance, in the 3-shot learning setting, for each of the 20 training runs over different class splits, 3 random samples were drawn for each *novel* class.

	UCF101		HMDB		Olympic	
Method	1-shot	5-shot	1-shot	5-Shot	1-shot	5-shot
<i>Sample-Proto</i>	0.248	0.117	0.341	0.205	0.297	0.193
<i>Learned-Proto</i>	0.219	0.106	0.323	0.198	0.272	0.191

Table 2: Cosine distance between mean of synthetic features and mean of test data for novel classes. Results are reported for all the datasets in 1-shot and 5-shot setting. Lower number is better.

	UCF101		HMDB	
\mathcal{T} Function	1-shot	5-shot	1-shot	5-Shot
<i>No-Pool</i>	57.8±1.7	79.8±1.1	23.3±2.3	51.2±0.8
<i>Max-Pool</i>	60.6±1.3	79.6±0.9	24.3±2.5	50.7±0.8
<i>Average-Pool</i>	61.7±1.6	79.7±0.8	34.4±1.3	50.9±0.6

Table 3: Comparison of Action Recognition accuracy (%) of our proposed framework with different dimensionality reduction function \mathcal{T} on UCF101 and HMDB datasets under G-FSL settings. *No-Pool* denotes use of no dimensionality reduction function, *Max-Pool* denotes max-pooling and *Average-Pool* denotes average-pooling. We report harmonic mean accuracy from 20 different training runs and standard deviation is reported by \pm .

4.3. Datasets

We evaluate our method on three publicly available datasets for action recognition,

- UCF101 [20]: contains 13320 videos spanned over 101 classes. For our experiments we randomly split the data into 51 *seen* and 50 *novel* classes.
- HMDB51 [12]: contains 6766 videos spanned over 51 classes. For our experiments we randomly split the data into 26 *seen* and 25 *novel* classes.
- Olympic Sports [16]: contains 783 videos spanned over 16 classes. We refer Olympic Sports as Olympic in subsequent sections. For our experiments we randomly split the data into 8 *seen* and 8 *novel* classes.

5. Experimental Results

Details about our experimental results are discussed in this section. Results for our proposed framework under G-FSL setting are compared with different baselines in Section. 5.1. Section. 5.2 describes ablation studies highlighting the importance of various components in our framework. Comparison with state-of-the art method under FSL setting is discussed in Section. 5.3.

5.1. Generalized-Few Shot Learning (G-FSL)

As the results under G-FSL on the aforementioned action recognition datasets are reported for the first time, we designed the baseline methods by taking inspiration from the *Few-Shot Learning* methods proposed in image classification tasks. Similar to [29], sample mean of the examples from a class was taken as its class-prototype and

we term this method as *Heuristic-Proto* in our experiments. In *Sample-Proto* the video features were directly taken as the class-prototype as described in [1]. Our proposed class prototype vectors generated through CPTN are termed as *Learned-Proto* in the experiments. To evaluate the performance of the entire framework, a vanilla classifier - *Base-Classifier* was directly trained on *C3D* video features with standard augmentation mentioned in [11]. The results for all the above approaches on UCF101, HMDB and Olympic datasets are reported in Table. 1.

As can be seen in Table. 1, *Heuristic-Proto*, *Sample-Proto* and *Learned-Proto* outperform the *Base-Classifier* with a huge margin for all the shots in all the datasets. This establishes that, the addition of generated features removes classifier’s bias towards *seen* classes which is depicted by the gain of accuracy for novel classes. Taking mean of the samples for novel classes, decreases the performance of *Heuristic-Proto* for 1-shot while giving competitive results for 5-shots. This can be attributed to poor representation of the conditional element by taking mean of single example. *Sample-Proto* performs better than other two baselines but shows higher standard deviation. This is because its performance depends on the quality of chosen samples as the class prototype is represented by the samples themselves without any aggregation. In contrast, transferred statistics through CPTN in our proposed *ProtoGAN* framework reduces this adverse effect. *Learned-Proto* outperforms baseline methods by a maximum of 1.4%, 5.7% and 2.9% on UCF101, HMDB and Olympics, respectively, for the 1-shot setting.

Figure. 2 illustrates the comparison of per class mean accuracy of our proposed *ProtoGAN* framework with other

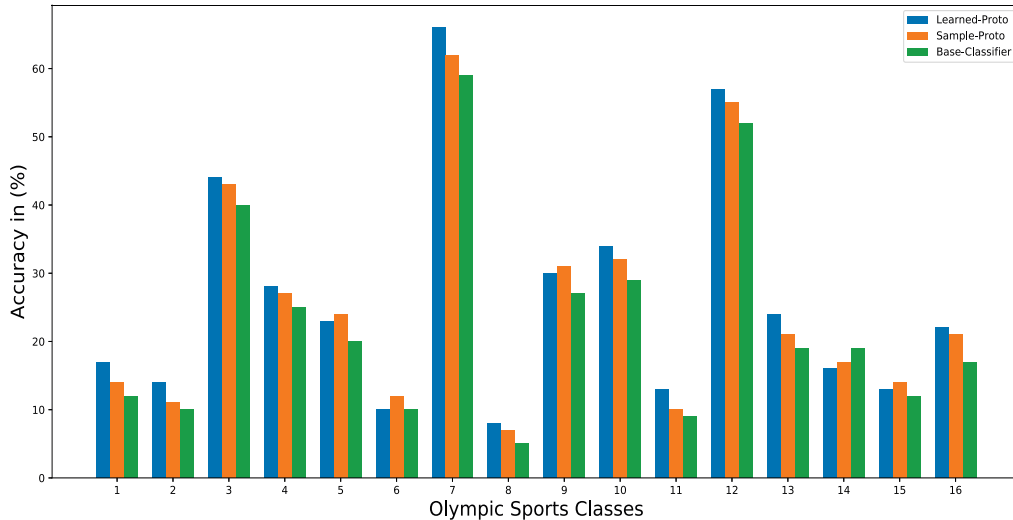


Figure 2: Mean classification accuracy (in %) of our proposed *ProtoGAN* framework with other baselines for all classes (when present as novel classes in different experimental runs) for olympic sports dataset. Classes are numbered according to alphabetical order of names. We report mean accuracy from 10 different runs. Best viewed in color.

	UCF101		HMDB	
Method	1-shot	5-shot	1-shot	5-Shot
<i>Without-Pruning</i>	60.8±1.9	79.5±0.9	23.7±2.6	50.7±0.8
<i>With-Pruning</i>	61.7±1.6	79.7±0.8	34.4±1.3	50.9±0.6

Table 4: Comparison of action recognition accuracy (%) of our proposed framework with and without pruning the synthetic features under G-FSL setting. We report harmonic mean accuracy from 20 different training runs and standard deviation is reported by \pm .

baselines. Mean accuracy is taken for all classes when present as novel classes during 10 different experimental runs. As can be seen from the figure, our proposed framework performs better in 11 out of 16 classes in Olympic Sports dataset in the range of 1-4%. However, the margin of inaccuracies is less significant as compared to the accurate classes. This re-establishes the superiority of the *ProtoGAN* framework.

5.2. Ablation Studies

5.2.1 Quality of Synthetic Features

To verify and compare the quality of the synthesized features of *novel* classes using different conditional elements, we try to quantify their similarity with real video features available from the test set. Specifically, we take mean of all the synthesized and real features for a *novel* class separately and compute cosine distance between them. The mean cosine distance over all the *novel* classes for 1 and 5-shots are reported in Table. 2. Results are reported for UCF101, HMDB and Olympic datasets. One can observe, examples

generated using our *Learned-Proto* for 1-shot matches the distribution of real features much more accurately than that of *Sample-Proto*.

5.2.2 Dimensionality Reduction

Dimensionality reduction function (\mathcal{T}) plays a crucial role in getting rid of intricate details and preserving class semantics. To demonstrate the effect of (\mathcal{T}), a model was trained without applying \mathcal{T} and another with Max-Pool instead of Average-Pool. As can be in seen in Table. 3, the recognition accuracy drops significantly when no (\mathcal{T}) is used and thus restating that dropping intricacies helps in creating better *class-prototype* vector. Max-Pool is slightly inferior in performance compared to Average-Pool. A reason for this is that features obtained after average-pool has more aggregate information than max-pool, hence provides more stable features.

	Methods	k-shots				
		1	2	3	4	5
UCF101	[14]	-	68.7±3.3	73.5±2.2	76.5±2.1	78.6 ±1.8
	<i>Learned-Proto</i>	57.8±3.0	71.1±3.0	75.3±2.7	78.0±1.8	80.2±1.3
HMDB	[14]	-	42.1±3.6	47.5±3.3	50.3±3.4	52.5±3.1
	<i>Learned-Proto</i>	34.7±9.2	43.8±5.4	49.1±5.1	52.3±4.2	54.0±3.9
Olympic	[14]	-	73.2±7.4	75.3±7.3	80.2±7.2	83.8±7.1
	<i>Learned-Proto</i>	71.6±9.4	75.0±7.4	78.4±6.2	82.1±5.6	86.3±5.1

Table 5: Comparison of action recognition accuracy (in %) of our proposed *ProtoGAN* framework against current state-of-the-art on UCF101, HMDB and Olympic dataset under FSL settings. We report mean accuracy from 20 different training runs and standard deviation is reported by \pm . Recognition accuracy for 1-shot is missing for [14].

5.2.3 Pruning

To validate the effect of pruning synthetic features in increasing order of their reconstruction loss, the proposed *ProtoGAN* framework was trained without it and the results are reported in Table. 4. The superior performance of the classifier trained on data after pruning validates the efficacy of the reconstruction loss in creation and subsequent selection of meaningful examples. However, the gain is higher for HMDB as compared to UCF101. This can be attributed to less number of seen examples which affects the quality of class-prototype vectors and hence the generator.

5.3. Few Shot Learning (FSL)

A comparison of our *ProtoGAN* framework against the current state-of-the-art approach [14] under FSL setting is presented in Table. 5. The authors of [14] have also used *C3D* [22] video features in their evaluation. Our method outperforms [14] in all k-shots with similar or lower standard deviation. Note that the results for 1-shot are not reported in [14]. The improved performance can be attributed to the usage of a *learned* prototype vector. The *learned* vector computed through a non-linear function via a network provides a better alternative to a vector formed by a linear combination of basis vectors as mentioned in [14]. This establishes the wide applicability of our approach in both G-FSL and conventional FSL settings.

6. Conclusion

In this paper, we present a novel *ProtoGAN* framework which synthesizes video features for novel categories by conditioning a CGAN with a *class-prototype* vector embedding to address the problem of *Few-Shot Learning* for action recognition. Class-prototype vector is learnt through a feature aggregator network called Class Prototype Transfer Network (CPTN). The performance of the proposed framework was evaluated on three publicly

available datasets for both seen and novel classes under G-FSL setting for the first time. We obtained encouraging results showing the efficacy of the proposed framework under G-FSL settings on action recognition and established a strong benchmark for future research. Under standard FSL setting, we outperform state-of-the-art method on all the datasets across different shots.

Acknowledgement: We gratefully acknowledge Bri-jesh Pillai and Partha Bhattacharya at Mercedes-Benz R&D India, Bangalore for providing the funding and infrastructure for this work.

References

- [1] A. Antoniou, A. Storkey, and H. Edwards. Data augmentation generative adversarial networks. In *arXiv:1711.04340*, 2018.
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. In *arXiv:1701.07875*, 2017.
- [3] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [4] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang. A closer look at few-shot classification. In *ICLR*, 2019.
- [5] R. Felix, B. V. Kumar, I. Reid, and G. Carneiro. Multi-modal cycle-consistent generalized zero-shot learning. In *ECCV*, 2018.
- [6] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [7] S. Guerriero, B. Caputo, and T. Mensink. Deepncm: Deep nearest class mean classifiers. In *ICLRw*, 2018.
- [8] B. Hariharan and R. B. Girshick. Low-shot visual object recognition. In *ICCV*, 2017.
- [9] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.
- [10] D. Kingma and J. Ba. Adam: a method for stochastic optimization. In *arXiv:1412.6980*, 2015.

- [11] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *ACM Communication*, 2017.
- [12] H. Kuehne, H. Jhuang, R. Stiefelhagen, and T. Serre. A large video database for human motion recognition. In *HPCSE*, 2013.
- [13] M. Mirza and S. Osindero. Conditional generative adversarial nets. In *arXiv:1411.1784*, 2014.
- [14] A. Mishra, V. K. Verma, M. S. K. Reddy, A. Subramaniam, P. Rai, and A. Mittal. A generative approach to zero-shot and few-shot action recognition. In *WACV*, 2018.
- [15] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel. A simple neural attentive meta-learner. In *ICLR*, 2018.
- [16] J. C. Niebles, C.-W. Chen, and L. Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *ECCV*, 2010.
- [17] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2016.
- [18] M. Ren, S. Ravi, E. Triantafillou, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel. Meta-learning for semi-supervised few-shot classification. In *ICLR*, 2018.
- [19] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017.
- [20] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. In *arXiv:1212.0402*, 2012.
- [21] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018.
- [22] D. Tran, L. D. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [23] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, et al. Matching networks for one shot learning. In *NIPS*, 2016.
- [24] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. In *PAMI*, 2018.
- [25] B. Xu, H. Ye, Y. Zheng, H. Wang, T. Luwang, and Y.-G. Jiang. Dense dilated network for few shot action recognition. In *ICMR*, 2018.
- [26] Z. Xu, L. Zhu, and Y. Yang. Few-shot object recognition from machine-labeled web images. In *CVPR*, 2017.
- [27] H. Yang, X. He, and F. Porikli. One-shot action localization by learning sequence matching network. In *CVPR*, 2018.
- [28] C. Zhang and Y. Peng. Visual data synthesis via gan for zero-shot video classification. In *IJCAI*, 2018.
- [29] R. Zhang, T. Che, Z. Ghahramani, Y. Bengio, and Y. Song. Metagan: An adversarial approach to few-shot learning. In *NIPS*. 2018.
- [30] L. Zhu and Y. Yang. Compound memory networks for few-shot video classification. In *ECCV*, 2018.