

REST: RETRIEVE & SELF-TRAIN FOR GENERATIVE ACTION RECOGNITION

Anonymous authors

Paper under double-blind review

ABSTRACT

This work is on training a generative action/video recognition model whose output is a free-form action-specific caption describing the video (rather than an action class label). A generative approach has practical advantages like producing more fine-grained and human-readable output, and being naturally open-world. To this end, we propose to adapt a pre-trained *generative* Vision & Language (V&L) Foundation Model for video/action recognition. While recently there have been a few attempts to adapt V&L models trained with contrastive learning (e.g. CLIP) for video/action, to the best of our knowledge, we propose the very first method that sets out to accomplish this goal for a generative model. We firstly show that direct fine-tuning of a generative model to produce action classes suffers from severe overfitting. To alleviate this, we introduce REST, a training framework consisting of two key components: (a) an unsupervised method for adapting the generative model to action/video by means of pseudo-caption generation and *Self-training*, i.e. without using any action-specific labels; (b) a *Retrieval* approach based on CLIP for discovering a diverse set of pseudo-captions for each video to train the model. Importantly, we show that both components are necessary to obtain high accuracy. We evaluate REST on the problem of zero-shot action recognition where we show that our approach is very competitive when compared to contrastive learning-based methods. Code will be made available.

1 INTRODUCTION

Large-scale pre-training of Foundation Models (FM) for Vision-Language (V&L) understanding (Radford et al., 2021; Jia et al., 2021; Yuan et al., 2021; Yu et al., 2022; Alayrac et al., 2022; Li et al., 2022; Wang et al., 2022) has recently revolutionized several multi-modal understanding tasks including image captioning, VQA, text-based retrieval, and visual reasoning. Moreover, recently, such models, and especially CLIP (Radford et al., 2021), have been also applied to more traditional visual tasks including object detection (Gu et al., 2021; Du et al., 2022; Minderer et al., 2022) and video/action recognition (Wang et al., 2021; Ju et al., 2022; Castro & Heilbron, 2022; Lin et al., 2022c). This work is on adapting a V&L FM for video/action recognition but in contrast to all prior work that use a V&L model trained with contrastive learning, i.e. CLIP (Wang et al., 2021; Ju et al., 2022; Castro & Heilbron, 2022; Lin et al., 2022c), our focus is on adapting a *generative* V&L model, and specifically BLIP (Li et al., 2022), to the video domain with the goal of producing an action-specific caption in an autoregressive manner. Besides this being a challenging research goal on its own, a generative approach has practical advantages including producing more fine-grained (compared to an action class label) and more interpretable/human-readable output, and being naturally more open-world (compared to contrastive learning based approaches) in a sense that it does not require a priori definition of the action classes.

In prior work (Wang et al., 2021; Ju et al., 2022; Castro & Heilbron, 2022; Lin et al., 2022c), CLIP adaptation to action/video recognition was seamlessly done using the standard contrastive loss and hand-engineered or learned prompts encoding action class information. In contrast, as Table 1 shows the direct end-to-end fine-tuning of the model on class names suffers from severe overfitting and poor open-world (i.e. zero-shot) generalizability, a fundamental property of V&L models that we wish not to sacrifice when adapting the model to the video domain.

Table 1: Zero-shot generalization of our Generative Action Recognition (GAR) model trained with action class labels vs action-specific pseudo-caption as proposed in REST. We report generalized zero-shot classification results on Kinetics-220 (1-vs-620 setting; see Sec. 5 for more details). We observe that training with class labels results in very poor zero-shot generalization. Upon inspecting the model’s predictions, we observe that it does indeed restrict itself to the class names seen on the training set. Some examples of failure cases (expected - predicted): “massaging neck” - “massaging head”, “playing ocarina” - “playing harmonic”, “separating eggs” - “scrambling eggs”, “putting on lipstick” - “sticking tongue out” etc.

Method	pseudo-captions	Top-1	Top-5
GAR + class labels	✗	0.76 ± 0.4	37.8 ± 1.0
GAR + REST (this paper)	✓	29.51 ± 0.71	56.12 ± 0.37

To alleviate the aforementioned overfitting problem, we propose to entirely drop action labels during training and adopt an unsupervised adaptation approach. Specifically, starting from an image-based generative model (specifically BLIP (Li et al., 2022)), we describe a self-training procedure where the model is iteratively trained on action data, without any labels, using an auto-regressive objective where the training pseudo-captions are produced by the model itself. Not only the model, in this case, is able to generalize much better to unseen classes but also such an approach is label-free making it potentially suitable to be applied to any large-scale video dataset. However, our findings suggest that self-training solely is not sufficient to train a highly accurate model as the generated pseudo-captions are lacking diversity¹. To alleviate this, we further propose a retrieval approach which is integrated into self-training where every few epochs a CLIP model is used to retrieve, for each training video, additional pseudo-captions from similar videos (from the same dataset). These pseudo-captions are then used to enhance learning by means of increasing the diversity of the training data. Critically, we show that our retrieval approach drastically increases the accuracy and the quality of the trained models. In summary, our **main contributions** are:

- To the best of our knowledge, we propose the very first method for adapting a *generative* V&L FM for open-vocabulary action/video recognition.
- We introduce REtrieve & Self-Train (REST), a training framework consisting of two key components; (a) iterative *Self-training* without using any action-specific labels, and (b) CLIP-based *Retrieval* for discovering a diverse set of pseudo-captions to train the model. We show that both components are necessary to train a high quality model.
- We evaluate REST on zero-shot action recognition where we show that our approach is very competitive with respect to CLIP-based methods, matching and/or surpassing the state-of-the-art on standard evaluation benchmarks.

2 RELATED WORK

Vision-Language Foundation Models: Following CLIP and Align (Radford et al., 2021; Jia et al., 2021), a number of methods were proposed that train Vision & Language (V&L) Foundation Models with contrastive learning (Li et al., 2021; Yao et al., 2021; Yu et al., 2022; Zhai et al., 2022). More recently, a few attempts were made towards training generative (V&L) models that produce as output a caption describing the input image in an auto-regressive manner (Tsimpoukelli et al., 2021; Li et al., 2022; Alayrac et al., 2022; Wang et al., 2022). From these methods, we used BLIP (Li et al., 2022) as the start point of our framework, the goal of which is to produce a generative action recognition model that operates in the video domain in an unsupervised manner. BLIP was chosen on the basis that publicly available models are provided but, in practice, any available image-based generative model can be incorporated into our training pipeline.

Unsupervised & Semi-supervised Image Captioning: There are only very few methods that have attempted to train an image captioning model without full supervision. (Chen et al., 2017) propose a method that transfers a COCO model to other domains by means of adversarial training using unpaired data in the target domain. Similar in spirit approaches focusing on training an LSTM discriminator to distinguish real from generated captions were proposed in Feng et al. (2019),

¹Some sort of the so-called confirmation bias (Arazo et al., 2020)

Laina et al. (2019); Zhou et al. (2021). Moreover, semi-supervised approaches include Kim et al. (2019) which combines paired and unpaired data with adversarial training, and Chen et al. (2021) which performs iterative self-training using a mean teacher consisting of an ensemble of independently trained models. Compared to Chen et al. (2021), which is the closest work to ours from the above, our method (a) does not use a mean teacher but critically a *retrieval* component which is shown to greatly improve the generated captions, and (b) focuses on image-to-action/video captioning (rather than on image-to-image) which is significantly more challenging.

V&L Foundation Models for Action/Video Recognition: Following the development of CLIP, a number of very recent works have attempted to adapt it to the video domain. X-CLIP (Lin et al., 2022c) trains a lightweight transformer on top of CLIP image features for spatiotemporal fusion. Ju et al. (2022) uses soft prompt learning to adapt CLIP to the video domain, while Wang et al. (2021) performs standard end-to-end finetuning. Starting from CLIP, FitCLIP (Castro & Heilbron, 2022) proposes a teacher-student approach based on a small video-text labelled dataset and pseudo-labels generated on a large unlabelled dataset. Notably, their method uses an ensemble largely relying on the original CLIP model to produce high accuracy. The aforementioned works use contrastive learning and labelled data for CLIP-to-video adaptation. In contrast, our work is the first, to our knowledge, that attempts to adapt a generative V&L model to the action/video domain, which on its own poses significant difficulties, and notably, without using any action-specific labels.

Video Captioning: A large body of methods are trained in a fully supervised manner using caption annotations (which are costly), see for example Pan et al. (2020); Tang et al. (2021b;a); Lin et al. (2022b); Liu et al. (2022). In contrast, REST does not use any human annotations to train the models.

Zero-shot Action Recognition: Most works on zero-shot action recognition build on learning attribute-based *semantic embeddings* from video features in order to make them be close to the word embedding of the class names (Zhu et al., 2018; Brattoli et al., 2020; Bretti & Mettes, 2021; Mettes, 2022; Estevam et al., 2022; Pu et al., 2022; Luo et al., 2022), or use a text encoder which is learned or updated as part of the training (Ni et al., 2022; Ge et al., 2022; Qian et al., 2022; Lin et al., 2022a; Qian et al., 2022). Most recent approaches using word embeddings focus on the *alignment* problem i.e. learning visual representations that match the corresponding class semantic embeddings but also generalize to unseen class embeddings (Pu et al., 2022; Luo et al., 2022). All the aforementioned methods above operate in a discriminative setting where the class embeddings need to be computed. In contrast, our method is a generative one, being able to generate action-specific captions in an autoregressive manner without having to manually pre-define the classes of interest.

3 METHOD

3.1 OVERVIEW OF RETRIEVE & SELF-TRAIN (REST)

We are given an action/video recognition dataset D consisting of N video clips v_i , $i = 1, \dots, N$ *without* the class labels. We construct a generative action recognition model consisting of a video encoder $g_v(\cdot)$ and an autoregressive text decoder $g_t(\cdot)$, both instantiated from a pre-trained generative V&L model, specifically BLIP (Li et al., 2022) (see Sec. 3.2). Our objective is to iteratively train the generative action recognition model on pseudo-captions produced by the model itself (self-training) using a standard language modeling loss (see Sec. 3.3). The self-training process is greatly enhanced by increasing the diversity of the pseudo-captions by using a retrieval module based on CLIP (see Sec. 3.4). The integrated retrieve & self-train framework is described in Sec. 3.5.

3.2 GENERATIVE ACTION RECOGNITION MODEL

Both the video encoder and the autoregressive text decoder are instantiated from a BLIP model (Li et al., 2022) pretrained on a large scale image-text dataset. We largely maintain BLIP’s generative architecture by only slightly modifying its visual transformer to combine temporal information across frames using a temporal adapter. Fig. 1 depicts our generative action recognition network.

The **text decoder** $g_t(\cdot)$ consists of 12 transformer layers, each performing causal attention on the text tokens and cross-attention between the text and the visual tokens produced by the video encoder.

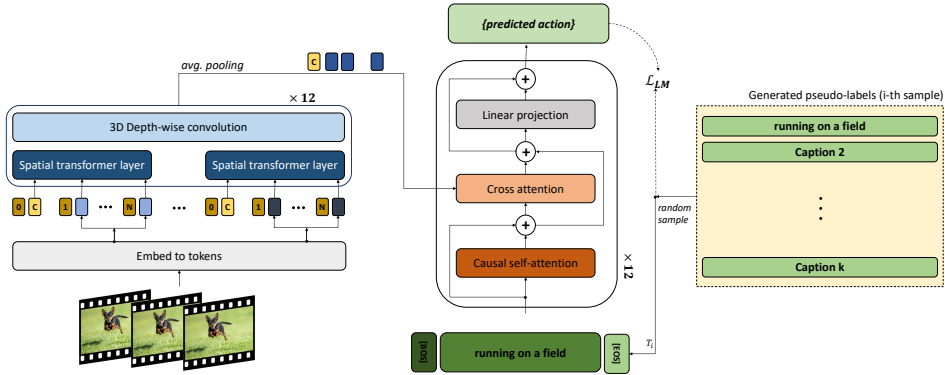


Figure 1: **Generative Action Recognition network.** We largely maintain BLIP’s generative architecture (Li et al., 2022) adding only a lightweight temporal adapter, in a form of a 3D depth-wise convolution, to combine information across frames.

The **video encoder** $g_v(\cdot)$ is constructed by inserting a temporal adapter in-between the layers of BLIP’s spatial visual transformer (a ViT (Dosovitskiy et al., 2020)). The visual transformer layers continue to operate independently on each frame (i.e. spatial attention only) while the adapter will pool and combine information across frames, handling the temporal modeling aspect. Specifically, let $\mathbf{Z}^l \in \mathbb{R}^{T \times H \times W \times d}$ be the output of the spatial transformer’s layer l (the class token \mathbf{Z}_{cls}^l is excluded). The adapter operates on \mathbf{Z}^l as follows:

$$\mathbf{Z}^l = \mathbf{Z}^l + a_d(\mathbf{Z}^l), \quad (1)$$

where $a_d(\cdot)$ takes the form of a $3 \times 3 \times 3$ 3D depth-wise convolution. The adapter layer is initialized such that the temporal residual, introduced by it, will be of low magnitude initially, closely aligning the network’s output early on to that of the pre-trained image model. This helps stabilizing the training by avoiding large magnitude gradients caused by incoherent predictions.

Finally, \mathbf{Z}^l , along with the class tokens \mathbf{Z}_{cls}^l , are averaged over T , to compute video feature $\mathbf{F}^l \in \mathbb{R}^{(H \times W + 1) \times d}$ that interacts with the text decoder.

3.3 LANGUAGE MODELLING LOSS

The language modelling objective \mathcal{L}_{LM} is the negative log-likelihood of the self-generated pseudo-captions, and it is used as the main loss to update our model $g_v(\cdot)$ and $g_t(\cdot)$. To compute the loss, we compute the visual tokens $\mathbf{F} \in \mathbb{R}^{(H \times W + 1) \times d} = g_v(\mathbf{X})$ from the video frames \mathbf{X} , and we are given a pseudo-caption \tilde{w} . As we shall see in Sec. 3.5, \tilde{w} is sampled from a set of cached pseudo-captions initially generated by BLIP and then updated by $g_t(\cdot)$ over the course of training. We compute the text tokens as $y_{\tilde{w}} = \phi(\tilde{w}) \in \mathbb{N}^M$, with ϕ being a tokenizer function that maps (sub-)words into the one-hot vectors spanning the vocabulary size. Following standard practices in image captioning, we prepend to $y_{\tilde{w}}$ a $y_0 = [\text{BOS}]$ token as well as a prompt $y_p = \phi(\text{“A video of ”}) \in \mathbb{N}^P$. The input to the model’s decoder $g_t(\cdot)$ is set as $y = [y_0, y_p, y_{\tilde{w}}] \in \mathbb{N}^{1+P+M}$. The text decoder applies left-to-right masked attention (i.e. causal attention) and produces an output $o = [o_i]_{i=1}^{1+P+M}$ with $o_i = g_t(y_{i' < i} | \mathbf{F})$. The language modelling loss is then computed using standard cross-entropy (CE):

$$\mathcal{L}_{LM}(y_{\tilde{w}}) = \sum_{i=1+P}^{1+P+M} CE(y_i, o_i). \quad (2)$$

3.4 VIDEO-VIDEO & VIDEO-TEXT RETRIEVAL

A key component of REST is the integration of a retrieval module into the self-training process. Our framework uses video-video and video-text retrieval modules both instantiated from a pretrained CLIP model (Radford et al., 2021) that remains frozen over the training process. We selected CLIP due to its accuracy and strong generalizability properties (Radford et al., 2021; Alayrac et al., 2022).

Given a video \mathbf{X} and pseudo-caption w , we use the CLIP image $g_I^C(\cdot)$ and text $g_T^C(\cdot)$ encoders to compute video $\mathbf{f}_C = \sum_t g_I^C(\mathbf{X}_t) \in \mathbb{R}^d$ and text $\mathbf{t}_C = g_T^C(w) \in \mathbb{R}^d$ features, respectively. For a given pair of videos i, j , and a given pair of video-caption i, j , the video-video s_{vv} and video-text s_{vt} similarities are computed as:

$$s_{vv}^{i,j} = \mathbf{f}_C^i \cdot \mathbf{f}_C^j \quad (3) \quad s_{vt}^{i,j} = \mathbf{f}_C^i \cdot \mathbf{t}_C^j, \quad (4)$$

where we use the subscript C to refer to features computed by the frozen CLIP model.

3.5 RETRIEVE & SELF-TRAIN (REST)

This section describes our framework for training the model of Sec. 3.2 without using any human-annotated action classes/captions. Instead, the model is trained on pseudo-captions generated by the model itself in a self-training manner. We also introduce a retrieval module based on CLIP for increasing the diversity of the pseudo-captions used to supervise the training for each input video, which is shown to greatly enhance the learning process. REST is summarised in Algorithm 1.

Algorithm 1 REST Training

Require: $\{v_i\}$, $i = 1 : N$ clips, pre-trained CLIP and BLIP models.

- 1: Compute $\tilde{W}_i = [\tilde{w}_{i;t} = g_v^{BLIP}(\mathbf{X}_{i;t})]$ $\triangleright t = 1, \dots, T, i = 1, \dots, N$
 - 2: Compute $\mathbf{f}_C^i, \mathbf{t}_C^i = g_{I,T}^C(\mathbf{X}_i; \tilde{W}_i)$ (Sec. 3.4) $\triangleright i = 1, \dots, N$
 - 3: Compute $s_{vv}^{i,j}$ (Eq. 3) and $s_{vt}^{i,j}$ (Eq. 4) $\triangleright i = 1, \dots, N, j = 1, \dots, N$
 - 4: Update \tilde{W}_i (Eqs. 5-6) \triangleright Retrieve top- K captions from top- H similar videos
 - 5: **while** training **do**
 - 6: **for** R epochs **do**
 - 7: Sample batch v_i and $\tilde{w}_i \in \tilde{W}_i$
 - 8: Update g_t and g_v using Eq. 2
 - 9: **end for**
 - 10: Compute $\tilde{W}_i \leftarrow \tilde{W}_i \cup [\tilde{w}_{i;new} = g_t(g_v(\mathbf{X}_i))]$ $\triangleright i = 1, \dots, N$
 - 11: Update $\mathbf{t}_C^i = g_T^C(\tilde{W}_i)$ $\triangleright i = 1, \dots, N$
 - 12: Compute $s_{vv}^{i,j}$ (Eq. 3) and $s_{vt}^{i,j}$ (Eq. 4) $\triangleright i = 1, \dots, N, j = 1, \dots, N$
 - 13: Update \tilde{W}_i (Eqs. 5-6) \triangleright Retrieve top- K captions from top- H similar videos
 - 14: **end while**
-

Self-train: The process is iterative and, at each training iteration, video v_i maintains a list of K associated pseudo-captions $\tilde{W}_i = [\tilde{w}_{i,k}]$, $k = 1, \dots, K$ describing its content. From this list, a pseudo-caption is randomly sampled and used as supervisory signal for v_i to train the model using Eq. 2. After training for R epochs, the model produces a new pseudo-caption for v_i denoted as $\tilde{w}_{i,new}$. This is added to the existing list resulting in $K + 1$ pseudo-captions $\tilde{W}_i \leftarrow \tilde{W}_i \cup \tilde{w}_{i,new}$.

Retrieve: The retrieval module is used to update the pseudo-caption list for each video v_i over the course of self-training. Specifically, for each v_i , we use the video-video similarity of Eq. 3 to compute $s_{vv}^i = s_{vv}^{i,j}$, $j = 1, \dots, N$, and then retrieve the corresponding $K + 1$ captions associated to the H most similar videos to v_i , creating a large list of $H(K + 1)$ captions as:

$$\Omega_i = \bigcup_{j \in \text{top-}H} \tilde{W}_j. \quad (5)$$

We then compute the CLIP-based text embeddings $\mathbf{t}_C^j = g_T^C(\tilde{w}_j)$ for $\tilde{w}_j \in \Omega_i$, and the $H(K + 1)$ video-text similarity scores $s_{vt}^{i,j} = \mathbf{f}_C^i \cdot \mathbf{t}_C^j$. Finally, we update \tilde{W}_i for video v_i by keeping the top- K most relevant captions from $\tilde{W}_i \cup \Omega_i$, i.e. we update the list of captions \tilde{W}_i for video v_i as:

$$\tilde{W}_i \leftarrow \bigcup_{i \in \text{top-}K} [\tilde{w}_i \in [\tilde{W}_i \cup \Omega_i]]. \quad (6)$$

Initialization: At the beginning of the training process, the videos v_i , $i = 1, \dots, N$ are populated with pseudo-captions by an off-the-shelf captioning model. To this end, we use BLIP’s text decoder to produce a caption $\tilde{w}_{i,t}$ for a set of T video frames $t = 1, \dots, T$. To encourage the

generation of action-specific outputs, we use a set of manual prompts, such as “a video of”, “a person is” or “someone is”, to initialize the text decoder’s output. **Training efficiency:** Besides training the model with the language modelling loss, the above self-training framework includes a retrieval step which can potentially render the training process slow. However, note that the video-video scores s_{vv}^i for each video v_i are computed using a frozen CLIP model, and, hence, all these scores can be pre-computed and re-used over the course of training. Moreover, the video feature used to compute the video-text scores s_{vt}^i can also be pre-computed for all videos, and only the text features corresponding to newly produced captions need be computed during training. Hence, s_{vt}^i can be also efficiently computed during training.

4 EVALUATION OF GENERATIVE ACTION RECOGNITION

While evaluating standard classification models is trivial, this is not the case for our generative action recognition model, given that its output is free unconstrained text. Direct character-by-character assessment is complicated as (a) there is more than one action caption that could describe the video, and (b) the same action can be expressed in multiple ways. Assessing the quality of generated text is an open research question, and various metrics such as CIDEr (Vedantam et al., 2015), BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) have been proposed to evaluate the correctness of the generated captions from the perspective of human judgment. However, such metrics are hard to correlate with the typical accuracy score expected in a classification problem and tend to penalize predictions outside the expected vocabulary disproportionately.

To alleviate this, we introduce a novel **CLIP-based Text Accuracy Metric** (CLIP-TAM), that capitalizes on the capacity of the CLIP text encoder to produce semantically distinctive features, ignoring elements of grammar completeness such as pronouns or adverbs, and being also invariant to permutations in the position of the class names. Given a set of class names $\mathcal{C} = \{\gamma_1, \dots, \gamma_{|\mathcal{C}|}\}$, we compute the class embeddings using the CLIP text encoder $\mathbf{w}_c = g_T^C(\gamma_c)$, compute the CLIP embedding for a caption \tilde{w}_i generated by our model as $\mathbf{t}_C^i = g_T^C(\tilde{w}_i)$, and select the target class from \mathcal{C} as $\hat{c} = \arg \max_c \mathbf{w}_c \cdot \mathbf{t}_C^i$. By associating a class to the predicted caption, we can directly compute an accuracy score, reducing the problem to a classic closed-set classification one.

User study: We evaluate the capacity of CLIP-TAM to correlate with human judgment by conducting a user study on 1,000 videos randomly selected from Kinetics-600. Following Levinboim et al. (2019), we formulate the correctness of a video-caption pair as the binomial probability $\hat{p} = P(CORRECT|video, caption)$ that can be estimated from the Bernoulli process. Each trial corresponds to a different human evaluator. The evaluators are shown an input video, which can be visualised multiple times, and are concomitantly asked the following question: *Does the text describe the action shown in the video?*. The raters are requested to choose between *yes* or *no*. To avoid inducing any bias, the interpretation of what represents a “correct” caption is left to the human evaluators to interpret and decide. While this could lead to unstable and inconsistent ratings, Levinboim et al. (2019) showed that with sufficient annotators (8-10) the results become stable and reproducible. The final accuracy score is produced by taking the average of the binary annotations of each annotator on a per-sample basis and thresholding it to 0.5.

Conclusions: The results of the above experiment showed that human annotators agree with the CLIP-TAM metric in 75.67% of the cases further validating the correctness of the latter to act as a proxy for classification. The accuracy measured by CLIP-TAM on the Kinetics-600 subset was 71.0% while the annotators’ accuracy was 73.80%.

5 RESULTS

5.1 EXPERIMENTAL SETTING

Our models are trained on the training set of Kinetics-400 (Kay et al., 2017), without using any labels, and tested on HMDB-51 (Kuehne et al., 2011), UCF-101 (Soomro et al., 2012), Kinetics-600 (Carreira et al., 2018) and the validation set of Kinetics-400. For all the experiments, unless otherwise stated, we sample uniformly 8 frames at a resolution of 224×224 px. During inference we follow BLIP’s approach and use beam search to generate the captions.

Zero-shot experiments: We train our model following Algorithm 1 on the $\sim 200K$ videos comprising the training set of Kinetics-400, without using any labels. We initialize the algorithm with the following parameters: the number of nearest neighbours H is set to 2,000, the size of the pseudo-caption cache per video is $K = 3$, and the number of epochs R between retrievals is set to 10. The model is trained for 60 epochs. The algorithm and the models were implemented using PyTorch (Paszke et al., 2019). The full list of hyper-parameters is reported in the appendix. In all cases, we evaluate the accuracy of our method using our newly introduced CLIP-TAM.

Evaluation protocol: On HMDB-51 and UCF-101, following Ni et al. (2022), we report the average top-1 accuracy and standard deviation computed across each of the three test splits. Similarly, on Kinetics-600, we perform evaluations on the three testing splits introduced in Chen & Huang (2021). Each split contains videos belonging to 160 classes different from the ones found in Kinetics-400, with a total of 220 unique classes among the three splits. Note that Chen & Huang (2021) re-assigned the labels to ensure that there is no overlap between the defined 220 classes and the 400 ones from Kinetics-400. Further to this standard setting, we also report results for the more challenging generalized zero-shot setting. Finally, we also report accuracy on the validation set of Kinetics-400.

Few-shot experiments: We also validate our model for few-shot recognition tasks. We finetune the trained models directly using the class names (instead of action-specific captions). Depending on the number of samples per class available at train time, we finetune the model between 50 (for 16-shot) and 200 epochs (for 2-shot) using the same hyper-parameters used as for the REST training.

Evaluation protocol: We apply the standard M -shot setting with $M = \{2, 4, 8, 16\}$ training videos per class. We report the average Top-1 and Top-5 accuracy on the UCF-101 and HMDB-51 datasets, after training and evaluating on each of the 3 train-test splits.

5.2 COMPARISON WITH STATE-OF-THE-ART

Evaluation of different types of supervision: In Table 2, we compare three types of supervision on the validation set of Kinetics-400. CLIP and BLIP are trained on very large scale datasets with image-language supervision. For BLIP, we evaluate its performance in standard generative mode but also in *CLIP mode*. A6 (Ju et al., 2022) and X-CLIP (Ni et al., 2022) are methods based on CLIP adaptation trained in a fully supervised manner on Kinetics-400. Our model was also trained on Kinetics-400 but without any labels in an unsupervised manner. As expected A6 and X-CLIP perform the best. Notably, our method significantly outperforms both CLIP (Radford et al., 2021) and BLIP (Li et al., 2022), illustrating the impact of the unsupervised adaptation on the target dataset.

Another interesting observation is that BLIP, in both generative and CLIP mode, performs significantly worse than CLIP. This shows that REST, initialized with BLIP, is in a disadvantage when compared with other methods based on CLIP adaptation (e.g. Ju et al. (2022); Ni et al. (2022)).

Table 2: Evaluation of different types of supervision.

Method	Dataset	Top-1	Top-5
<i>Web image-language supervision</i>			
CLIP (Radford et al., 2021)	CLIP-400M	53.6	79.9
BLIP (CLIP mode) (Li et al., 2022)	LAION-115M	47.3	66.1
BLIP (Li et al., 2022)	LAION-115M	23.5	41.7
<i>Full supervision</i>			
A6 (Ju et al., 2022)	Kinetics-400	76.9	93.5
X-CLIP (Ni et al., 2022)	Kinetics-400	83.8	96.7
<i>Unsupervised adaptation</i>			
REST (Ours)	Kinetics-400	63.9	81.0

Evaluation on HMDB-51/UCF-101: Zero-shot classification results are shown in Table 3, where we compare the performance of our approach against current state-of-the-art methods, including the concurrent works of Ni et al. (2022); Ge et al. (2022). Despite the generative nature of our approach, our method sets a new state-of-the-art result on HMDB-51, and delivers competitive results on UCF-101 where the best method is Ni et al. (2022) trained in a supervised setting.

Evaluation on Kinetics-220: We report the results of REST against state-of-the-art methods in Tables 4 and 5 under *two different scenarios*. In Table 4, we report the results obtained under the standard setting of *novel, unseen classes* where the evaluation is done over the restricted set of novel classes only, i.e. where the models have to perform 1-vs-160 classification². Under this setting, our model is only surpassed by the concurrent work of (Ni et al., 2022).

Table 5 shows results for the more challenging *generalized zero-shot* setting, i.e. where the model is evaluated under a 1-vs-620 classification setting. For this setting, not surprisingly, CLIP, due to being trained on a very large dataset, performs the best. Notably, our method largely outperforms Ni et al. (2022), showcasing the benefits of generative modelling for zero-shot action recognition.

Few-shot experiments: Results are reported in Table 6. Notably, our approach gets a large boost with minimum additional training, surpassing the concurrent work of Ni et al. (2022) when $M = 2$ and $M = 4$, as well as getting on par results for $M = 8$ and $M = 16$.

Table 3: Zero-shot classification results on HMDB-51 and UCF-101 in terms of top-1 accuracy.

Method	HMDB-51	UCF-101
Discriminative approaches		
E2E (Gao et al., 2019)	32.7	48
TS-GCN (Brattoli et al., 2020)	23.2 ± 3.0	34.2 ± 3.1
ER-ZSAR (Chen & Huang, 2021)	35.3 ± 4.6	51.8 ± 2.9
CLIP (Radford et al., 2021)	46.2	73.0
MUFI (Qiu et al., 2021)	31.0	60.9
ActionCLIP (Wang et al., 2021)	40.8 ± 5.4	58.3 ± 3.4
ClipBert (Lei et al., 2021)	21.4 ± 1.0	27.8 ± 0.8
Frozen (Bain et al., 2021)	27.8 ± 0.3	45.9 ± 1.3
ViSET-96 (Doshi & Yilmaz, 2022)	40.2	68.3
BridgeFormer (Ge et al., 2022)	37.7 ± 1.2	53.1 ± 1.4
AURL (Pu et al., 2022)	40.4	60.9
ResT_101 (Lin et al., 2022a)	41.1 ± 3.7	58.7 ± 3.3
X-CLIP (Ni et al., 2022)	44.6 ± 5.2	72 ± 2.3
X-Florence (Ni et al., 2022)	48.4 ± 4.9	73.2 ± 4.2
Generative approaches		
REST (Ours)	49.7 ± 1.14	69.1 ± 0.62

Table 4: Zero-shot classification results on Kinetics-220 (1-vs-160 setting).

Method	Top-1	Top-5
Discriminative approaches		
DEM (Zhang et al., 2017)	23.6 ± 0.6	49.5 ± 0.4
GCN (Ghosh et al., 2020)	22.3 ± 0.6	49.7 ± 0.6
ER-ZSAR (Chen & Huang, 2021)	42.1 ± 1.4	73.1 ± 0.3
X-CLIP (Ni et al., 2022)	65.2 ± 0.4	86.1 ± 0.8
X-Florence (Ni et al., 2022)	68.8 ± 0.9	88.4 ± 0.6
Generative approaches		
REST (Ours)	51.7 ± 1.1	75.2 ± 0.4

5.3 ABLATION STUDIES

Effect of retrieval step: A key component of REST is integrating retrieval into self-training. We analyze its importance by conducting an experiment where only self-training is used (i.e. $H = 1$).

²While there are a total of 220 unique classes in Kinetics-220 each split contains 160 only.

Table 5: **Generalized** zero-shot classification results on Kinetics-220 (1-vs-620 setting).

Method	Top-1	Top-5
Discriminative approaches		
CLIP (Radford et al., 2021)	47.03	74.4
X-CLIP (Ni et al., 2022)	14.76 \pm 0.51	60.93 \pm 0.25
Generative approaches		
REST (Ours)	29.51 \pm 0.71	56.12 \pm 0.37

Table 6: Few-shot classification results on HMDB-51 and UCF-101 in terms of top-1 accuracy.

Method	HMDB-51				UCF-101			
	2	4	8	16	2	4	8	16
Discriminative approaches								
TSM (Lin et al., 2019)	17.5	20.9	18.4	31.0	25.3	47.0	64.4	61.0
TimeSformer (Bertasius et al., 2021)	19.6	40.6	49.4	55.4	48.5	75.6	83.7	89.4
Swin-B (Liu et al., 2022)	20.9	41.3	47.9	56.1	53.3	74.1	85.8	88.7
X-CLIP (Ni et al., 2022)	53.0	57.3	62.8	64.0	76.4	83.4	88.3	91.4
X-Florence (Ni et al., 2022)	51.6	57.8	64.1	64.2	84.0	88.5	92.5	94.8
Generative approaches								
REST (Ours)	54.0	59.1	62.1	64.0	88.2	90.2	92.6	93.5

We observe that the results on the validation set of Kinetics-400 drop massively from 63.9% to 44.0%, clearly validating the importance of retrieval in REST.

Effect of number of cached captions: We evaluate the impact of K in Algorithm 1 (the number of generated captions cached per video) on Kinetics-400. The results shown in Table 7a suggest that increasing K from 1 to 3 offers a significant improvement in terms of Top-1 accuracy. However, going beyond $K = 3$ does not offer further gains indicating that samples ranked lower are less likely to be representative of the video.

Table 7: Analysis on Kinetics-400.

(a) Effect of number of cached captions.

	K = 1	K = 3	K = 5
Ours	62.60	63.93	63.98

(b) Effect of number of retrieval steps.

	$N_I = 1$	$N_I = 2$	$N_I = 3$	$N_I = 4$
Ours	59.16	62.75	63.93	63.95

Effect of number of retrieval steps: We evaluate the effect of the frequency of updating the pseudo-captions as shown in 1.10-13 of Algorithm 1. Table 7b shows that updating the captions results in better performance, showing that, over the course of training, they become more semantically meaningful. Generally, up to a certain point, more retrieval steps correlate with better performance.

6 CONCLUSION

In this paper, we proposed REST, the very first generative method for video/action recognition trained without human supervision. Departing from an image-based generative model, an iterative algorithm that alternates between self-training and retrieving is proposed, leading to a model that can produce captions with a strong semantic correlation with action classification. Remarkably, our model can operate in a zero-shot setting without the need of manually defining a target set, and sets a strong baseline for adaptation with limited data. We believe our method sets a new path on training generative models for zero-shot action recognition in a truly open-set scenario. Code and models will be released to encourage further analysis.

ETHICS CONSIDERATIONS

Generative language models exhibit various forms of biases learned from the data, such as occupational biases that link certain activities with some groups of individuals (Kirk et al., 2021). Our work derives its pseudo-labels from the representation learned by BLIP and CLIP models, and hence, the resulting network may have inherited the bias present in the source. Therefore, before deployment, the models should undergo in-depth checks and considerations.

REPRODUCIBILITY STATEMENT

We detail in Section 5.1 the training settings and evaluation protocols, while in appendix, Table 10 we list the augmentations and training hyper-parameters. We will also release the code to ensure the reproducibility of our approach.

REFERENCES

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022.
- Eric Arazo, Diego Ortego, Paul Albert, Noel E O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *International Joint Conference on Neural Networks*, 2020.
- Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *IEEE International Conference on Computer Vision*, 2021.
- Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *International Conference on Machine Learning*, 2021.
- Biagio Brattoli, Joseph Tighe, Fedor Zhdanov, Pietro Perona, and Krzysztof Chalupka. Rethinking zero-shot video classification: End-to-end training for realistic applications. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- Carlo Bretti and Pascal Mettes. Zero-shot action recognition from diverse object-scene compositions. *arXiv preprint arXiv:2110.13479*, 2021.
- Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018.
- Santiago Castro and Fabian Caba Heilbron. Fitclip: Refining large-scale pretrained image-text models for zero-shot video understanding tasks. *arXiv preprint arXiv:2203.13371*, 2022.
- Shizhe Chen and Dong Huang. Elaborative rehearsal for zero-shot action recognition. In *IEEE International Conference on Computer Vision*, 2021.
- Tseng-Hung Chen, Yuan-Hong Liao, Ching-Yao Chuang, Wan-Ting Hsu, Jianlong Fu, and Min Sun. Show, adapt and tell: Adversarial training of cross-domain image captioner. In *IEEE International Conference on Computer Vision*, 2017.
- Xianyu Chen, Ming Jiang, and Qi Zhao. Self-distillation for few-shot image captioning. In *Winter Conference on Applications of Computer Vision*, 2021.
- Keval Doshi and Yasin Yilmaz. Zero-shot action recognition with transformer-based video semantic embedding. *arXiv preprint arXiv:2203.05156*, 2022.
- 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.
- Yu Du, Fangyun Wei, Zihe Zhang, Miaoqing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- Valter Estevam, Rayson Laroca, Helio Pedrini, and David Menotti. Global semantic descriptors for zero-shot action recognition. *IEEE Signal Processing Letters*, 2022.
- Yang Feng, Lin Ma, Wei Liu, and Jiebo Luo. Unsupervised image captioning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- 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 *AAAI Conference on Artificial Intelligence*, 2019.
- Yuying Ge, Yixiao Ge, Xihui Liu, Dian Li, Ying Shan, Xiaohu Qie, and Ping Luo. Bridging video-text retrieval with multiple choice questions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

- Pallabi Ghosh, Nirat Saini, Larry S Davis, and Abhinav Shrivastava. All about knowledge graphs for actions. *arXiv preprint arXiv:2008.12432*, 2020.
- Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, 2021.
- Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. *European Conference on Computer Vision*, 2022.
- Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The Kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. Image captioning with very scarce supervised data: Adversarial semi-supervised learning approach. *arXiv preprint arXiv:1909.02201*, 2019.
- Hannah Rose Kirk, Yennie Jun, Filippo Volpin, Haider Iqbal, Elias Benussi, Frederic Dreyer, Aleksandar Shtedritski, and Yuki Asano. Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models. *Advances on Neural Information Processing Systems*, 2021.
- Hildegard Kuehne, Hueihan Jhuang, Estibaliz Garrote, Tomaso Poggio, and Thomas Serre. HMDB: a large video database for human motion recognition. In *IEEE International Conference on Computer Vision*, 2011.
- Iro Laina, Christian Rupprecht, and Nassir Navab. Towards unsupervised image captioning with shared multimodal embeddings. In *IEEE International Conference on Computer Vision*, 2019.
- Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- Tomer Levinboim, Ashish V Thapliyal, Piyush Sharma, and Radu Soricut. Quality estimation for image captions based on large-scale human evaluations. *arXiv preprint arXiv:1909.03396*, 2019.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*, 2022.
- Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021.
- Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pp. 74–81, 2004.
- Chung-Ching Lin, Kevin Lin, Lijuan Wang, Zicheng Liu, and Linjie Li. Cross-modal representation learning for zero-shot action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022a.
- Ji Lin, Chuang Gan, and Song Han. TSM: Temporal shift module for efficient video understanding. In *IEEE International Conference on Computer Vision*, 2019.
- Kevin Lin, Linjie Li, Chung-Ching Lin, Faisal Ahmed, Zhe Gan, Zicheng Liu, Yumao Lu, and Lijuan Wang. SwinBERT: End-to-end transformers with sparse attention for video captioning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022b.
- Ziyi Lin, Shijie Geng, Renrui Zhang, Peng Gao, Gerard de Melo, Xiaogang Wang, Jifeng Dai, Yu Qiao, and Hongsheng Li. Frozen CLIP models are efficient video learners. *arXiv preprint arXiv:2208.03550*, 2022c.

- Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Zhekun Luo, Shalini Ghosh, Devin Guillory, Keizo Kato, Trevor Darrell, and Huijuan Xu. Disentangled action recognition with knowledge bases. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2022.
- Pascal Mettes. Universal prototype transport for zero-shot action recognition and localization. *arXiv preprint arXiv:2203.03971*, 2022.
- Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection with vision transformers. *arXiv preprint arXiv:2205.06230*, 2022.
- Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding language-image pretrained models for general video recognition. *arXiv preprint arXiv:2208.02816*, 2022.
- Boxiao Pan, Haoye Cai, De-An Huang, Kuan-Hui Lee, Adrien Gaidon, Ehsan Adeli, and Juan Carlos Niebles. Spatio-temporal graph for video captioning with knowledge distillation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Annual meeting of the Association for Computational Linguistics*, 2002.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances on Neural Information Processing Systems*, 2019.
- Shi Pu, Kaili Zhao, and Mao Zheng. Alignment-uniformity aware representation learning for zero-shot video classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- Rui Qian, Yeqing Li, Zheng Xu, Ming-Hsuan Yang, Serge Belongie, and Yin Cui. Multimodal open-vocabulary video classification via pre-trained vision and language models. *arXiv preprint arXiv:2207.07646*, 2022.
- Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xiao-Ping Zhang, Dong Wu, and Tao Mei. Boosting video representation learning with multi-faceted integration. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021.
- 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.
- Mingkang Tang, Zhanyu Wang, Zhenhua LIU, Fengyun Rao, Dian Li, and Xiu Li. CLIP4Caption: Clip for video caption. In *ACM International Conference on Multimedia*, 2021a.
- Mingkang Tang, Zhanyu Wang, Zhenhua Liu, Fengyun Rao, Dian Li, and Xiu Li. CLIP4Caption ++: Multi-clip for video caption. *arXiv preprint arXiv:2208.02816*, 2021b.
- Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *Advances on Neural Information Processing Systems*, 2021.

- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. GIT: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022.
- Mengmeng Wang, Jiazheng Xing, and Yong Liu. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021.
- Shen Yan, Xuehan Xiong, Anurag Arnab, Zhichao Lu, Mi Zhang, Chen Sun, and Cordelia Schmid. Multiview transformers for video recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021.
- Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- Li Zhang, Tao Xiang, and Shaogang Gong. Learning a deep embedding model for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Yucheng Zhou, Wei Tao, and Wenqiang Zhang. Triple sequence generative adversarial nets for unsupervised image captioning. In *International Conference on Acoustics, Speech, and Signal Processing*, 2021.
- Yi Zhu, Yang Long, Yu Guan, Shawn Newsam, and Ling Shao. Towards universal representation for unseen action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

A ADDITIONAL ABLATION STUDIES

Fully supervised training with action class names: While the main focus of our work is learning generative zero-shot and few-shot classifiers, nevertheless, herein we train and evaluate the generative action recognition (GAR) model trained with the original class names on Kinetics-400. The training is performed using the same hyper-parameters as the ones used for REST (see Table 10). As the results from Table 8 show, while competitive, GAR lags behind the current CLIP-based discriminative methods. Based on our prior results, we hypothesise that this is due to: (a) the BLIP model used in our method being weaker than CLIP, and (b) certain augmentations performed for classification being harder to apply for a generative model (e.g.: mixup, cutout etc.).

Table 8: Fully supervised classification results on Kinetics-400.

Method	Top-1	Top-5	Throughput
Discriminative approaches			
Action-CLIP (Wang et al., 2021)	83.8	96.2	-
A6 (Ju et al., 2022)	76.9	93.5	-
MVT-H (Yan et al., 2022)	89.1	98.2	-
X-Florence (Ni et al., 2022)	86.2	96.6	6
X-CLIP (Ni et al., 2022)	83.8	96.7	33
Generative approaches			
GAR	79.0	87.8	11

Effect of temporal adapter: In Table 9, we evaluate the effect of the video adapter introduced in Sec. 3.2 by training a model that applies only temporal average pooling in the last layer.

Table 9: Effect of temporal adapter.

Method	Top-1	Top-5
Ours (w/o) adapter	61.5	79.5
Ours	63.93	81.04

Training hyper-parameters are listed in Table 10.

Table 10: Training hyper-parameters. K is the number of samples per class for few-shot learning.

	REST training	Few-shot finetuning
Optimization		
Optimizer	AdamW (Loshchilov & Hutter, 2017)	
Optimizer betas	(0.9, 0.98)	
Batch size	64	
Weight decay	0.001	
Learning rate scheduler	cosine	
Initial learning rate	$2e - 5$	
Minimal learning rate	$2e - 8$	
Epochs	60	max(400/K, 30)
Data augmentation		
Random Flip	0.5	
Multi Scale Crop	(0.66, 0.75, 0.875, 1.0)	
Color Jitter	0.8	
Gray Scale	0.2	
Label smoothing	0.2	

B QUALITATIVE RESULTS




Figure 2: **Qualitative results:** Each row shows 5 frames sampled equidistantly from the video. Below the frames we print out the predictions made by our model.



a little girl eating a piece of chocolate
Failure case: Miss detecting the object (chocolate vs lipstick).



a man moving a mattress down a flight of stairs
Failure case: Wrong temporal action flow (here the mattress is moving up).



a little girl talking to the camera with the words subscribe on it
Failure case: Bias from the pre-trained model (the words subscribe on it).



the president of the united states speaking at a senate committee meeting
Failure case: Hallucinating context-based details (i.e. that the person is the president of US).



a person using a drill to drill a hole in a piece of wood
Failure case: Hallucinating plausible actions not shown in the video.

Figure 3: **Failure cases.** Each row shows 5 frames sampled equidistantly from the video. Below each set of frames we print out the predictions made by our model and a possible explanation associated to the particular failure case.

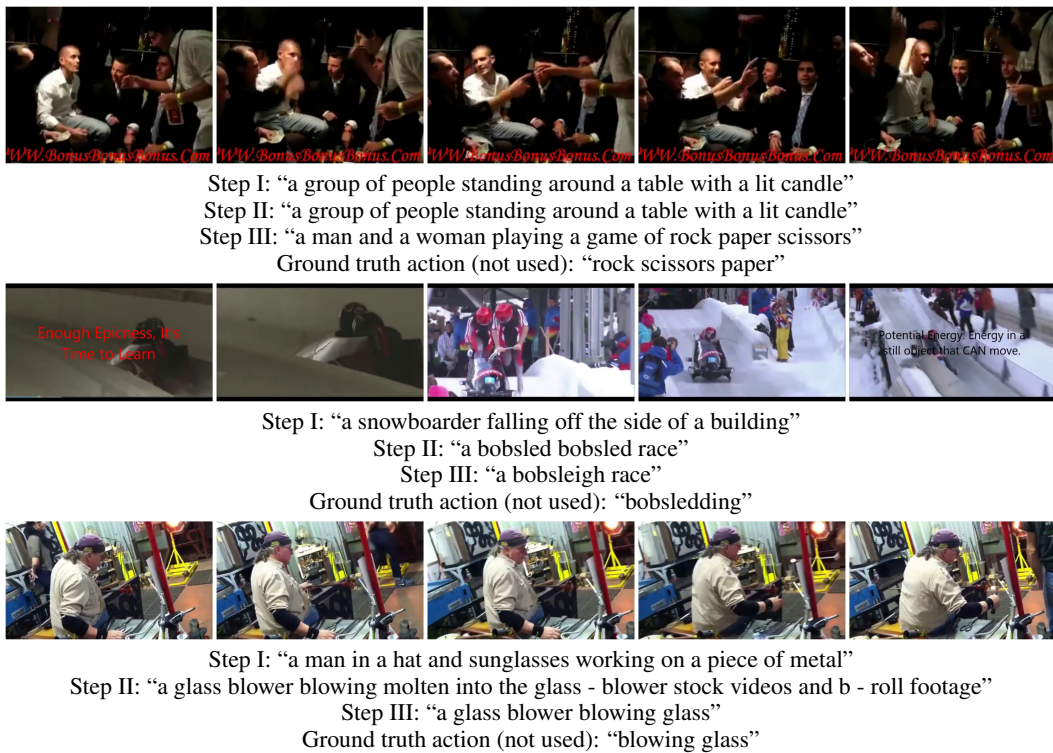


Figure 4: **Qualitative results** showcasing the evolution of the most relevant pseudo-caption for each video after each retrieval step.