

End-to-end Generative Pretraining for Multimodal Video Captioning

Paul Hongsuck Seo Arsha Nagrani Anurag Arnab Cordelia Schmid
 Google Research

{phseo, anagrani, aarnab, cordelias}@google.com

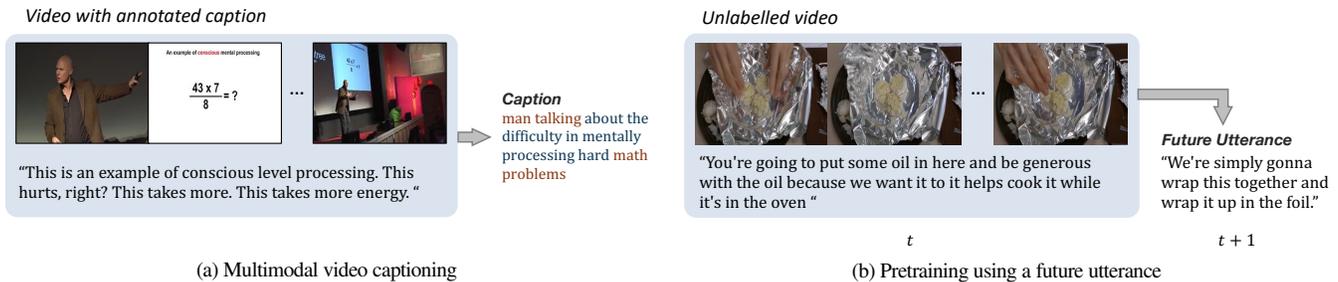


Figure 1. **Generative pretraining for Multimodal Video Captioning.** Multimodal Video Captioning takes visual frames and speech transcribed by ASR as inputs and predicts a caption. The example on the left (a) demonstrates that using both modalities jointly is beneficial to generate an accurate caption, *i.e.*, red words are present in the visual input whereas blue words correspond to the concepts in the ASR. Our new multimodal video generative pretraining (MV-GPT) uses a future utterance in time from the video stream as a captioning target (b). This objective can be applied to unlabeled data (*e.g.*, HowTo100M), which comes with ASR but no captions, and results in effective joint-pretraining for both a multimodal encoder and decoder.

Abstract

Recent video and language pretraining frameworks lack the ability to generate sentences. We present Multimodal Video Generative Pretraining (MV-GPT), a new pretraining framework for learning from unlabelled videos which can be effectively used for generative tasks such as multimodal video captioning. Unlike recent video-language pretraining frameworks, our framework trains both a multimodal video encoder and a sentence decoder jointly. To overcome the lack of captions in unlabelled videos, we leverage the future utterance as an additional text source and propose a bidirectional generation objective – we generate future utterances given the present multimodal context, and also the present utterance given future observations. With this objective, we train an encoder-decoder model end-to-end to generate a caption from raw pixels and transcribed speech directly. Our model achieves state-of-the-art performance for multimodal video captioning on four standard benchmarks, as well as for other video understanding tasks such as VideoQA, video retrieval and action classification.

1. Introduction

A long-standing goal of the AI community is the development of conversational multimodal systems that can both reliably perceive the world and effortlessly communicate with humans. An emerging benchmark of progress in this field is the task of multimodal video captioning [17, 32] - which tests

both abilities; a successful model must not only accurately understand ‘multimodal’ streams of input video (including the speech and the video frames), but also generate coherent natural language descriptions of the content.

Unsurprisingly, a major challenge in the field of vision and language learning is the lack of large-scale, manually annotated data. Annotating captions for videos is time intensive, expensive and subjective (with low inter-annotator agreement [17]) – this is in contrast to fields such as image classification where fully annotated datasets are orders of magnitude larger [15, 41, 56]. To overcome this limitation, there has been a flurry of recent works that pretrain their video-language models on instructional videos [32, 33, 42, 44, 45], a domain where the speech is particularly well aligned to visual content. Recently introduced datasets such as Cooking312K [45] and HowTo100M [34] leverage such instructional videos with associated captions from ASR (automatic speech recognition) to learn joint video-and-text embeddings [33, 44] or to train multimodal video encoders [27, 42]. However, the models in these works often do not contain a decoder, lacking the ability to generate sentences, and thus only the video encoder is transferred to the downstream tasks – indeed for the case of video captioning, the decoder is often learned from scratch [45, 47, 63]. While one can still initialize the decoder using independently pretrained weights such as those from a GPT-2 [37] model, we observe that this strategy is suboptimal and performance is significantly improved by optimizing the encoder and the decoder jointly.

For the task of multimodal video captioning, we require a

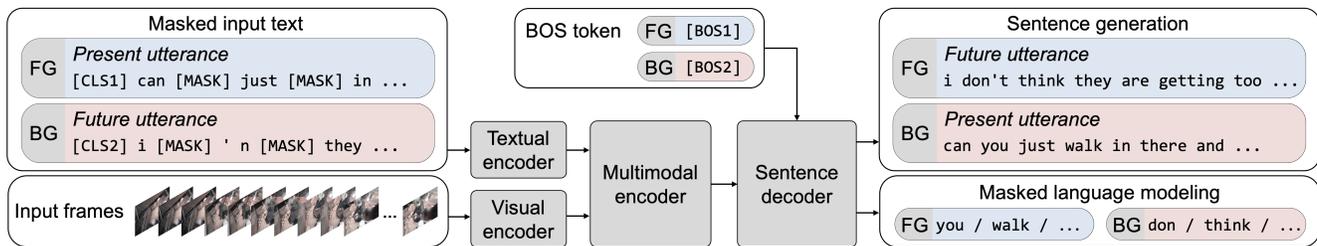


Figure 2. **Multimodal Video Generative Pretraining (MV-GPT) framework.** During pretraining, our network (which consists of modality specific encoders, a multimodal encoder and a sentence decoder) is trained with a new bi-directional objective. 1) Forward generation (FG, blue): Given input frames and present utterances from a video clip, we predict a future utterance and 2) Backward generation (BG, red): Given input frames and a future utterance, predict the current utterances. Both losses are applied to a triplet consisting of video frames, present utterances and a future utterance. To allow our model to recognise the different configurations, we attach distinct special classification tokens `CLS1` and `CLS2` to the input text for FG and BG respectively, as well as distinct `BOS1` and `BOS2` (beginning of sentence) tokens to the decoder for sentence generation.

model that can both encode multimodal videos (*i.e.* frames and textual inputs) and generate captions. Using multimodal information as input can greatly improve the quality of the generated captions (as illustrated in Figure 1a). However, learning such an encoder-decoder model jointly from unlabelled data is particularly challenging, as it requires two streams of textual data – naturally occurring transcribed speech accompanying the video for the encoder, and target sentences for the decoder – whereas unlabelled videos only come with a single stream of speech (Figure 1b). Recent works [17, 23, 32] have attempted to solve this problem with a denoising autoencoder - wherein the input speech to the model is artificially ‘noised’, *i.e.* random words are masked out [17, 23, 32]. The decoder is then tasked with simply reconstructing either the masked phrases or the original unmasked text, where the supervisory signals are provided only from the masked words. In these frameworks, additional losses are often required to strengthen the pretraining supervision, such as multimodal input alignment [32] and segment ordering [17].

In our framework, we introduce a novel stronger loss. We leverage future utterances as another source of textual data and train a model to generate these entirely unseen sentences as depicted in Figure 1b. To alleviate the problem that future utterances are not temporally aligned, we propose a backward generation objective where present aligned utterances are generated given future utterances. Experimental results show that a model pretrained with this bidirectional generation objective effectively transfers to multimodal video captioning and outperforms the state of the art by a margin.

We make the following contributions: (i) We propose a novel pretraining objective for multimodal video captioning that requires no manually annotated captions, and instead uses utterances sampled at different times in the same video. Our objective is bidirectional in time – *i.e.* we not only generate future utterances but also the present ones from the future; (ii) By using two sources of textual data, we are able to jointly train the entire encoder-decoder model. This is unlike previous works which pretrain only the (multimodal) encoder, thereby lacking the ability to generate captions [27, 42, 45]; (iii) Our encoder

is trained from raw pixels and words directly, in contrast with existing methods that rely on pre-extracted visual features limiting transfer to new domains [17, 23, 32]; (iv) We achieve state-of-the-art results on four video captioning benchmarks – YouCook2, ViTT, MSR-VTT and ActivityNet-Captions – consistently outperforming existing methods by significant margins; and finally (v) Our pretraining objective yields strong multimodal video representations, which achieve state-of-the-art performance on other video understanding tasks such as VideoQA, video retrieval and action classification.

2. Related Work

Video captioning. Early works in video captioning consisted of rule-based methods [10, 22], where subjects, verbs and objects (SVO-triplets) detected from the video were combined into sentence templates. Later work moved away from rule-based methods by framing captioning as a machine translation task [4, 39, 46], which developed the common encoder-decoder paradigm of today for the task – the encoder processes a set of video features and accumulates its hidden state, which is then passed to a decoder for producing a caption. Early works implemented the visual encoder as a 2D CNN (either frozen or finetuned) applied to video frames, which was then naturally extended to 3D CNNs [6, 53], to better capture motion dynamics, with temporal aggregation over the entire video typically performed using attention strategies [9]. Given the computational challenge of using expensive 3D CNNs applied to dense frame inputs (typically 30 fps), most of these works operated on pre-extracted features, only learning the fusion of features in the encoder. Unlike such works, we address this problem using a transformer-based encoder applied to raw pixels [3], sampled at a coarse rate to better capture long range context.

Pretraining with weakly paired data. Existing video captioning datasets [17, 55, 62] are orders of magnitude smaller than video classification datasets [20]. As a source of weakly paired video and language data, a number of works have used the visual frames and the Automatic Speed Recognition (ASR) transcripts of unlabelled videos to pretrain video

representations [27, 33, 42, 44, 45, 63]. These approaches learn multimodal representations by formulating proxy tasks such as masked language/frame modeling [42, 45], video-text matching [27, 33] or segment ordering [27]. While these studies show improvements on visual representation [33, 44, 45, 47] or multimodal video representation [27, 42, 63] learning, they are designed for discriminative tasks only, and lack the generation capability. Pretraining techniques for generative tasks such as ours, are fewer. While [23] use multimodal translation as a generative objective, their encoder is limited to accept visual inputs only. Works that use multimodal inputs to the encoder, train with masking losses – wherein words or phrases are masked and the objective is to reconstruct the original sentences [23, 32] or the masked targets [17] using an autoregressive generator. In contrast, we make use of utterances outside of the clip boundary, which are simply ignored in previous works. We leverage future utterances as a second source of textual data, and propose a bi-directional generation objective where the model generates the future utterance given the current utterance and vice versa. While we also use a masked language modelling loss, this is simply in addition to our primary generative bidirectional loss.

3. Method

Our objective is to pretrain a model that can effectively encode multimodal videos (visual frames and transcribed speech) as well as decode natural language sentences. This will allow us to use the model for multimodal captioning. In this section, we first describe the pretraining losses used to train the encoder and decoder jointly from unlabelled videos. We then describe our model, which consists of modality specific encoders, a multimodal encoder and a text decoder (Figure 2).

3.1. Pretraining Objectives and Losses

Our framework is designed to take advantage of unlabelled instructional video data, which consists of video frames and utterances often linked to the visual content [34]. As mentioned earlier, our framework requires two textual streams – an input to the encoder and a captioning target for the decoder. Because unlabelled videos do not have captioning targets, we instead propose a simple objective – our model is trained to generate a *future* utterance in the video given the current video context and current utterances (forward generation). This gives us two sources of textual supervision, the current utterance allows us to learn how to optimally fuse modalities in the video encoder, while the decoder is tasked with predicting a new utterance it has never seen before. However, our goal is video captioning, and not ‘predicting the future’. To enable our model to generate text corresponding to the present video context, we also add in an additional backward generation loss – where the model must generate the current utterance given the current video frames and a future utterance (backward generation). This encourages generated sentences to be temporally aligned (and hence more tightly coupled) with the visual inputs.

3.1.1 Bi-directional Utterance Generation

Given a large set of unlabelled videos, we extract short clips consisting of visual frames $F = \{f_1, \dots, f_{N_f}\}$ and transcribed speech utterances $U = \{u_1, \dots, u_{N_u}\}$ aligned with F . For each clip, we also consider the immediate future utterance $W = \{w_1, \dots, w_{N_w}\}$ where u_i and w_j are tokenized words in the transcribed utterances. Note that we use the term ‘utterance’ to refer to a single sentence of transcribed speech.

Forward Generation: Our model is trained to generate a future utterance W given clip frames F and present utterances U . Formally speaking, we formulate our forward generation objective to minimize the negative log-likelihood of the true future utterance W , where the loss function given by the chain rule is $\mathcal{L}_{FG} = -\sum_{i=1}^{N_w} \log P(w_i | w_1, \dots, w_{i-1}, F, U)$. This loss encourages the pretrained model to effectively encode temporally aligned multimodal inputs to predict the future utterance.

Backward Generation: We now apply the same loss as above, albeit in the backward direction. Namely, the model is tasked with generating present utterances U aligned with video frames F , conditioned on future utterances W and F . As in the forward generation, we also minimize the negative log-likelihood of the true present utterance $\mathcal{L}_{BG} = -\sum_{i=1}^{N_u} \log P(u_i | u_1, \dots, u_{i-1}, F, W)$. Note that the visual input F is temporally aligned with the decoder output U . This loss function encourages the network to generate a caption related to the visual contents.

3.1.2 Masked Language Modeling

As an additional supplementary loss, we also train with a masked language modeling (MLM) loss [11] $\mathcal{L}_{MLM}(X)$ where X is the input utterance on which the masking is applied. We apply this loss on both the forward and backward input utterances, *i.e.* as $\mathcal{L}_{MLM}(U)$ and $\mathcal{L}_{MLM}(W)$. Note that these losses are computed independently from the above bidirectional generation losses.

Unlike UniVL [32] where the MLM loss is applied to the outputs of the encoder, we apply it to the outputs of the decoder. This encourages the self attention layers in the decoder to focus on further multimodal contextualization of the textual tokens (since each masked token prediction requires knowledge of neighbouring context). As we show in the experiments, this leads to performance gains.

3.2. Model

Our model consists entirely of transformer blocks, and is trained end-to-end directly from pixels and word tokens.

3.2.1 Modality Specific Encoders

Given a multimodal video input consisting of the visual frames $F = \{f_1, \dots, f_{N_f}\}$ and text inputs $X = \{x_1, \dots, x_{N_x}\}$, we first extract features from the individual modalities independently. Note here that the textual input X is the aligned utterance U

in general (for computing the forward generation loss and for downstream captioning tasks) but is set to W when computing the backward generation loss.

Textual Encoder: We extract N_x contextualized textual embeddings $E = \{e_i\}$ from the input text using a BERT [11] encoder.

Visual Encoder: Unlike previous approaches [17, 32, 42, 45] where visual features are pre-extracted by models pretrained on different datasets, we extract the visual features directly from pixels. We use the recent transformer-based video encoder ViViT [3], in particular, the tubelet embedding scheme and the factorized encoder architecture. For the tubelet embedding scheme we first extract spatio-temporal 3D tubes from the visual input volume resulting in $S \times T$ token embeddings where S and T correspond to the numbers of tokens in the spatial and temporal dimensions, respectively. Then, the spatial transformer first takes each group of S embeddings from the same temporal index with a special `CLS` token embedding, and the temporal transformer models interactions between the output `CLS` embeddings of the individual spatial groups with another `CLS` embedding resulting in $T+1$ visual features $V = \{v_j\}$ – see [3] for further details.

Unlike 3D CNN visual encoders which operate on consecutive frames extracted at high frame rates (30 fps), our visual encoder can operate on coarsely sampled frames (1 fps), thus significantly reducing compute. This allows us to train the visual encoder end-to-end, and helps adapt our features across the domain gaps between pretraining and downstream datasets. It also allows the easy adoption of off-the-shelf video augmentation directly to RGB frames, which is useful for small-scale downstream benchmarks.

3.2.2 Multimodal Encoder

Once the two sets of textual features E and visual features V are extracted, our multimodal encoder fuses multimodal information using the co-attentional transformer used in [31, 42]. Each layer consists of two streams where each stream is a stack of two transformer blocks. In the textual stream, we first contextualize the features E using a cross-attention transformer block attending to the visual features V . Then, the output features are further contextualized by another transformer block with self-attention. The first transformer block performs inter-modality contextualization through a cross-attention process whereas the second transformer block carries out intra-modality contextualization through a self-attention process. In the same way, the visual stream V attends to the textual stream. Our multimodal encoder repeats this process R times resulting in the output multimodal features \hat{E} and \hat{V} .

3.2.3 Sentence Decoder

As shown in Figure 2, given multimodal video features $C = \hat{E} \cup \hat{V}$ as context, we autoregressively generate the output sentence Y conditioned on this context using a transformer

decoder. To generate token y_i , we first encode the previous generated tokens $Y_i = \{y_0, \dots, y_{i-1}\}$ with a look-up table and a positional embedding to produce $H_i = \{h_0, \dots, h_{i-1}\}$. We then encode the context C and the previous embedded tokens H_i using a single transformer. The outputs of this transformer are $\tilde{C} \cup \tilde{H}_i$, where $\tilde{H}_i = \{\tilde{h}_0, \dots, \tilde{h}_{i-1}\}$. Note that \tilde{C} refers to the multimodal input embeddings obtained from the decoder and is used for computing the MLM loss as discussed in Section 3.1.2. We then predict the next token y_i from \tilde{h}_{i-1} by a linear projection with a softmax: $y_i = \text{argmax}(\text{softmax}(\Phi \tilde{h}_{i-1}))$, where $\Phi \in \mathbb{R}^{\nu \times d}$ is the linear projection matrix and ν is the vocabulary size. The first word h_0 is set using the special `BOS` (beginning of sentence) token, and tokens are generated until a special `EOS` (end of sentence) token is generated. In practice, each iteration requires only a single forward pass on the decoder transformer with the aid of causal masking introduced in [48].

3.2.4 Input and Output Configurations

Pretraining: Since our pretraining objective is bidirectional, each triplet (F, U, W) consisting of the visual frames F , the present utterances U and the future utterance W is processed by the network twice. For forward generation, the model takes F and U as inputs and generates W , and it generates U given F and W , in backward generation. To enable the model to recognize the different configurations, we attach distinct, special tokens `CLS1` and `CLS2` to the input text for the forward and backward generation losses respectively as illustrated in Figure 2. Similarly, we feed distinct `BOS1` and `BOS2` tokens to the decoder to initiate sentence generation.

Finetuning for captioning: In downstream video captioning datasets, video clips (consisting of frames F and aligned utterances U) are manually annotated with a natural language caption. During finetuning, we attach the `CLS1` token to U (as is done in forward generation), since U is an aligned utterance, but for generation we feed in the `BOS2` token (as is done in backward generation to predict the present utterance), so that we also generate a temporally aligned caption.

3.2.5 Implementation Details

For our text encoder, we adopt the BERT-Base architecture with uncased wordpiece tokenization [11]. Our visual encoder uses the corresponding ViViT-Base configuration with a 1-layer temporal transformer and a tubelet size of $16 \times 16 \times 4$ [3]. Our multimodal encoder consists of 2 layers following [42] and finally, the decoder is based on the GPT-2 (117M parameters) architecture [37] but we modify it to take both multimodal input context C and a `BOS` token allowing conditional generation (the original GPT starts generation immediately by taking the first word as its input and only conditions on text). We initialize the text encoder and the decoder with the standard BERT and GPT-2 weights respectively pretrained on large-scale unlabelled corpora [11, 37]. Similarly, we initialize our visual encoder

using the pretrained weights on Kinetics 400 in [3] unless otherwise specified. Our model is pretrained end-to-end using the Adam optimizer [21] for 1.5M iterations with the batch size of 2048. For more detailed hyperparameters and training strategies for pretraining and finetuning, please refer to the supplementary materials.

4. Experiments

In this section, we first demonstrate our results on four different benchmarks for multimodal video captioning. We then also show that our pretrained model has the ability to generalise to other video understanding tasks such as video question answering (VideoQA), video retrieval and action classification.

4.1. Multimodal Video Captioning

4.1.1 Datasets and Evaluation Protocols

We use HowTo100M [34] as our pretraining dataset, and evaluate on four downstream captioning benchmarks.

HowTo100M [34] consists of 1.2M instructional videos from YouTube. Transcribed speech is obtained using the YouTube ASR API [1]. Following [42], we extract 53M triplets of frames, current utterances and future utterances for pretraining.

YouCook2 [62] is the most widely adopted benchmark for multimodal video captioning and contains 2,000 cooking videos for 89 different dishes with 14K video clips. Each video clip is annotated with a single captioning sentence.

Video Timeline Tags (ViTT) [17] was created to better reflect the distribution of instructional videos in the wild. It consists of 8,169 videos, 5,840 of these videos for training and the remaining videos for validation and testing. Videos are divided into 7.1 segments on average, with each segment accompanied by a short timeline tag.

MSR-VTT [55] is a standard benchmark with 10K open domain video clips for video captioning. The duration of each video clip is between 10 and 30 seconds, and 20 natural language descriptions are manually annotated per clip.

ActivityNet-Captions [24] is a standard dense video captioning benchmark consisting of 100K temporally localized sentences for 20k videos. We follow the standard splits with 50/25/25% examples for training, validation and test sets. To evaluate our model’s ability to predict captions, we use ground truth temporal proposals following [24].

We pretrain a single model on HowTo100M, which is then transferred to all four captioning benchmarks through finetuning. We report results using the following established metrics: BLEU-4 (B-4) [35], CIDEr (C) [49], METEOR (M) [5] and ROUGE-L (R-L) [29]. For ViTT, we measure BLEU-1 (B-1) instead of BLEU-4 following [17].

| PT Losses | PT parts | B-4 | C | M | R-L |
|---------------|----------|--------------|-------------|--------------|--------------|
| No PT | – | 13.25 | 1.03 | 17.56 | 35.48 |
| Baseline PT | E | 16.13 | 1.46 | 21.76 | 41.50 |
| CoMVT [42] | E | 14.46 | 1.24 | 18.46 | 37.17 |
| M-MASS [17] | E+D | 19.03 | 1.88 | 24.00 | 45.10 |
| UniVL [32] | E+D | 19.95 | 1.98 | 25.27 | 46.81 |
| MV-GPT (Ours) | E+D | 21.26 | 2.14 | 26.36 | 48.58 |

Table 1. Comparisons to existing pretraining losses on YouCook2. **PT** stands for pretraining. **PT parts** indicates which part of the model are pretrained, encoder (E) or both encoder and decoder (E + D). We reimplement the loss functions of existing methods but use our model and training strategies for fair comparison.

| FG | BG | MLM-E | MLM-D | WD | B-4 | C | M | R-L |
|----|----|--------------|-------|----|--------------|-------------|--------------|--------------|
| | | <i>No PT</i> | | | 13.25 | 1.03 | 17.56 | 35.48 |
| | | ✓ | | | 16.13 | 1.46 | 21.76 | 41.50 |
| ✓ | | ✓ | | | 20.65 | 2.05 | 25.81 | 47.22 |
| ✓ | | | ✓ | | 20.77 | 2.09 | 25.90 | 47.41 |
| ✓ | ✓ | ✓ | | | 20.82 | 2.10 | 26.20 | 48.22 |
| ✓ | ✓ | | ✓ | | 20.89 | 2.11 | 26.42 | 48.30 |
| ✓ | ✓ | | ✓ | ✓ | 21.26 | 2.14 | 26.36 | 48.58 |

Table 2. Ablation on YouCook2 showing the effect of our different loss components in pretraining. **FG**: Forward Generation loss. **BG**: Backward Generation loss. **MLM-E/MLM-D**: Masked Language Modelling loss applied on encoder outputs (E) or decoder outputs (D). **WD**: Weight Decay. **No PT**: No pretraining with any of these losses.

4.1.2 Results

In this section we ablate some key design choices, in particular the backbone and objective functions used in MV-GPT, and explore the impact of the end-to-end training. Finally, we compare our model to the state of the art.

Pretraining Losses: We implement a simple baseline, which consists of a masked language modelling loss given visual frames and ASR as input (Baseline PT). We also reimplement three state-of-the-art pretraining losses: (i) CoMVT [42], (ii) UniVL [32] and (iii) M-MASS [17]. For a fair comparison, we use our model architecture for all experiments, varying the loss function only. For the methods which pretrain the encoder only, we initialise the decoder with public GPT-2 weights [37]. For ‘No PT’, the encoder is not pretrained either, but is initialized with public BERT and ViViT pretrained on ImageNet21k.

Table 1 compares these different losses. We can observe that pretraining the encoder only brings moderate gains over training from scratch, for all the losses investigated. This performance is greatly improved by pretraining both the encoder and decoder jointly. Finally, we observe that our approach MV-GPT outperforms existing joint pretraining losses.

Effect of each Loss Term in MV-GPT: Table 2 shows the effect of each term in our loss function. The forward generation (FG) loss already provides strong supervision. When applying the masked language modelling loss on the decoder outputs (MLM-

| Arch. | Weights from / Trained on | E2E | | B-4 | C | M | R-L |
|-----------------|---------------------------|-----|----|--------------|-------------|--------------|--------------|
| | | PT | FT | | | | |
| <i>YouCook2</i> | | | | | | | |
| S3D | S3D [53] / Kinetics | | | 19.65 | 1.93 | 24.47 | 45.79 |
| S3D | MIL-NCE [33] / HowTo100M | | | 20.02 | 1.96 | 24.98 | 46.65 |
| ViViT | ViViT [3] / Kinetics | | | 19.54 | 1.93 | 24.42 | 45.93 |
| ViViT | MV-GPT / HowTo100M | ✓ | | 21.77 | 2.20 | 26.97 | 49.29 |
| ViViT | MV-GPT / HowTo100M | ✓ | ✓ | 21.26 | 2.14 | 26.36 | 48.58 |
| ViViT | MV-GPT / HowTo100M | ✓ | ✓† | 21.88 | 2.21 | 27.09 | 49.38 |
| <i>MSR-VTT</i> | | | | | | | |
| ViViT | MV-GPT / HowTo100M | ✓ | | 47.04 | 0.55 | 36.80 | 62.99 |
| ViViT | MV-GPT / HowTo100M | ✓ | ✓ | 48.92 | 0.60 | 38.66 | 64.00 |

Table 3. Ablation on YouCook2 with different visual encoder configurations. **E2E**: End-to-end training including the visual encoder. **PT**: Pretraining. **FT**: Finetuning. † Freeze the visual encoder at the beginning and tune end-to-end once converged during finetuning.

D) instead of the encoder outputs (MLM-E), performance is slightly improved due to the additional input contextualization provided by the decoder. Adding the backward generation (BG) loss provides a boost across all metrics. Additionally, we observe that adding weight decay (WD) [25] brings additional gains, and we report our scores in this full setting for the rest of the paper.

Visual Encoder and End-to-end Training: In Table 3, we first compare the ViViT [3] encoder to commonly used S3D features [53]. When both encoders are trained on Kinetics and fixed for multimodal pretraining and finetuning, they show comparable scores despite the large complexity of S3D due to the high frame rate required (30 fps vs. 1 fps for ViViT). Using HowTo100M to train a visual encoder, we observe large gains with both architectures as expected given the similarity in the domains – HowTo100M and YouCook2 are both instructional video datasets. However, we observe larger gains with ViViT where the visual encoder is optimized for generative losses within our framework and jointly trained with the other components thanks to the low complexity of the ViViT encoder. These results show the benefits of end-to-end pretraining.

We further investigate the effects of end-to-end training for finetuning. For YouCook2, we observe slight performance degradation when naively finetuning the network end-to-end from the beginning (row 4 to 5). This degradation is overcome by initially freezing the visual encoder and starting end-to-end training after convergence, which gives us a minor gain (row 6). These results indicate that our pretrained visual encoder already captures strong representations for inputs in a similar domain, and end-to-end finetuning is less critical in this case. However, we observe more significant gains on MSR-VTT since end-to-end finetuning becomes crucial given a larger domain gap (row 7 to 8).

Pretraining with Random Initialization: We also investigate the ability of the model to learn from scratch. We initialize the model either entirely randomly or using pretrained BERT, ViViT and GPT-2 weights. Table 4 shows that with random

| Initialization | MV-GPT Pretraining | B-4 | C | M | R-L |
|----------------|--------------------|--------------|-------------|--------------|--------------|
| Random | | 10.93 | 64.56 | 12.88 | 29.03 |
| Random | ✓ | 20.78 | 2.09 | 25.83 | 47.76 |
| Public weights | | 13.25 | 1.03 | 17.56 | 35.48 |
| Public weights | ✓ | 21.26 | 2.14 | 26.36 | 48.58 |

Table 4. Ablations on YouCook2 showing the effect of initialization and pretraining. **Public Weights**: Initialization with public BERT, GPT-2 and ViViT weights.

| Method | PT parts | Inputs | B-4 | C | M | R-L |
|----------------------|------------|------------|--------------|-------------|--------------|--------------|
| VideoBERT [45] | E | V | 4.04 | 0.49 | 11.01 | 27.50 |
| ActBERT [63] | E | V | 5.41 | 0.65 | 13.30 | 30.56 |
| MART [26] | – | V | 8.00 | 0.36 | 15.90 | – |
| AT [14] | – | T | 8.55 | 1.06 | 16.93 | 35.54 |
| DPC [43] | – | V+T | 2.76 | – | 18.08 | – |
| AT+Video [14] | – | V+T | 9.01 | 1.12 | 17.77 | 36.65 |
| DECEMBART [47] | E | V+T | 11.92 | 0.58 | 20.01 | 40.22 |
| VideoAsMT [23] | E+D | V | 5.30 | – | 13.40 | – |
| M-MASS [17] | E+D | V+T | 12.04 | 1.23 | 18.32 | 39.03 |
| UniVL [32] | E+D | V+T | 17.35 | 1.81 | 22.35 | 46.52 |
| MV-GPT (Ours) | E+D | V | 16.71 | 1.53 | 21.43 | 41.56 |
| MV-GPT (Ours) | E+D | T | 16.71 | 1.56 | 20.88 | 40.19 |
| MV-GPT (Ours) | E+D | V+T | 21.88 | 2.21 | 27.09 | 49.38 |

Table 5. Comparison to SOTA on YouCook2 for video captioning.

| Method | PT parts | Inputs | B-1 | C | M | R-L |
|----------------------|------------|------------|--------------|-------------|--------------|--------------|
| M-MASS [17] | E+D | V+T | 22.37 | 0.82 | 11.00 | 31.40 |
| MV-GPT (Ours) | E+D | V+T | 37.89 | 1.04 | 26.75 | 34.76 |

Table 6. Comparison to SOTA on ViTT for video captioning.

initialization, our method still performs very well (row 2), outperforming the model initialized with public BERT, GPT-2 and ViViT weights (row 3). Note that the pretrained ViViT weights were obtained from training on the fully supervised dataset Kinetics. Also, pretraining entirely from scratch even approaches the case where all parts of the model are initialized using public weights and pretrained (row 4).

Multimodal vs. Single Modality: In Table 5, we show results with text only and visual only inputs (we only feed the CLS token for the omitted modality). It is clear that both modalities are complementary and performance is best when combining both. Additionally, to assess the contribution of the visual modality, we test a model pretrained with text inputs only. Even when this pretrained model is finetuned with both modalities, the performance is significantly lower compared to a pretrained multimodal model (last row in Table 2): there is a 25% relative drop on all 4 metrics (e.g., 1.43 vs. 2.14 in CIDEr). When finetuned with text inputs only, the scores drop further (e.g., to 1.20 in CIDEr). These results confirm the importance of the visual inputs during pretraining.

Comparisons to the State of the Art: Finally, we compare



Figure 3. Qualitative example on YouCook2 (first row) and MSR-VTT (last two rows) including a failure case (last row). GT: Ground-truth caption. No-PT: No multimodal pretraining. MV-GPT: Our model pretrained on HowTo100M.

| Method | PT parts | Inputs | B-4 | C | M | R-L |
|----------------------|------------|------------|--------------|-------------|--------------|--------------|
| OA-BTG [60] | - | V | 41.40 | 0.47 | 28.20 | - |
| MGSA [9] | - | V | 42.40 | 0.48 | 27.60 | - |
| POS+CG [50] | - | V | 42.00 | 0.49 | 28.20 | 61.60 |
| POS+VCT [16] | - | V | 42.30 | 0.49 | 29.70 | 62.80 |
| SAM-SS [8] | - | V | 43.80 | 0.51 | 28.90 | 62.40 |
| ORG-TRL [61] | - | V | 43.60 | 0.51 | 28.80 | 62.80 |
| VNS-GRU [7] | - | V | 45.30 | 0.53 | 29.90 | 63.40 |
| DECĒMBERT [47] | E | V | 45.20 | 0.52 | 29.70 | 64.70 |
| VideoAsMT [23] | E+D | V | 41.70 | - | 28.50 | - |
| UniVL [32] | E+D | V+T | 41.79 | 0.50 | 28.94 | 60.78 |
| MV-GPT (Ours) | E+D | V+T | 48.92 | 0.60 | 38.66 | 64.00 |

Table 7. Comparison to SOTA on MSR-VTT for video captioning.

| Method | B-4 | M |
|----------------------|-------------|--------------|
| DCEV [24] | 1.60 | 8.88 |
| DVC [28] | 1.71 | 9.31 |
| Bi-SST [51] | - | 10.89 |
| HACA [52] | 2.71 | 11.16 |
| MWSDEC [38] | 1.46 | 7.23 |
| MDVC [19] | 1.46 | 7.23 |
| BMT [18] | 1.99 | 10.90 |
| MV-GPT (Ours) | 6.84 | 12.31 |

Table 8. Comparison to SOTA on ActivityNet-Captions for video captioning with ground-truth action proposals.

MV-GPT to existing methods on all four datasets. Table 5 compares our method to the state of the art on YouCook2, where we outperform all prior work including works pretrained on HowTo100M. On ViTT (Table 6), the gap is even larger, with our model advancing the state-of-the-art by 15% (absolute) compared to M-MASS in B-1 and M scores.

Despite the domain gap between instructional videos in

HowTo100M and general online videos in MSR-VTT, our model outperforms all existing work as shown in Table 7. Although UniVL also pretrains both the encoder and the decoder on HowTo100M, our method achieves relative improvements of over 31% thanks to our end-to-end training. Similarly, Table 8 shows that our pretraining method achieves state-of-the-art performance on ActivityNet-Captions despite the significant domain gap.

Qualitative Results: We show examples from YouCook2 and MSR-VTT in Figure 3. The first example illustrates that our model can use the visual modality to infer the term ‘sauce’ despite the ASR error ‘source’ and further recognizes its name ‘sriracha’. Similarly, the second example illustrates that our approach manages to take into account both modalities jointly. Finally, we show a failure case in the last row in which our model fails to capture the concept ‘ski lift’. A possible explanation is that the concept of a ski lift may be rarely seen in the pretraining dataset, a problem which may be alleviated by collecting more diverse pretraining videos, or incorporating external object knowledge through the use of pre-trained object detectors.

4.2. Non-generative Video Understanding Tasks

Although MV-GPT is a generative model and is particularly designed for multimodal video captioning, we also find that our pretraining technique learns a powerful multimodal video encoder that can be transferred easily to multiple video understanding tasks. In particular, we show results on VideoQA, video retrieval and action classification. For details on each task please refer to the supplementary materials.

VideoQA: We use MV-GPT as an encoder (no BOS token is fed to the decoder so it only contextualizes the input tokens; see supplementary materials for details) and the average pooled input embedding is fed to a two-layered MLP classifier to

predict the answer. The question is simply concatenated to the ASR inputs. Following the standard protocols in [42, 57], we measure the answer prediction accuracy on MSRVTT-QA [54] and ActivityNet-QA [59].

Table 9 compares the accuracy of MV-GPT to existing methods that are pretrained on HowTo100M [34]. Even though MV-GPT is not designed for this particular task, our model slightly outperforms the previous state-of-the-art VQA-T [57] (which is specifically designed for VideoQA) on both datasets.

Video Retrieval: The common practice for retrieval is to train a video-text joint embedding using *discriminative* losses only, typically in the form of a standard NCE loss [13], where each video clip has a single corresponding textual caption. Here we investigate whether our generative pretraining loss can provide a boost to performance. Since each example forms two inputs-target triplets in our bidirectional framework, we apply NCE losses on both (Bi-NCE). We then add our generative pretraining loss to this framework and report results in Table 10. We evaluate our model with and without ASR to compare fairly to existing works. We report recall at $k = \{1, 5, 10\}$ ($R@k$) and median rank (MdR) on MSR-VTT [55] following the standard 9K retrieval splits [58].

Our first observation is that our Bi-NCE serves as a strong baseline pretraining method for retrieval. We show that adding our generative losses further improves performance by a relative 6.3% in $R@1$, yielding state-of-the-art performance. Finally, adding ASR to our multimodal encoder further improves performance by a significant margin (+ 4%).

Action Classification: We test the visual encoder of MV-GPT on action classification following [3]. We evaluate models using top-1 classification accuracy on Kinetics 400 and 600 [20]. Note that we adopt the ViViT-Base architecture with factorized encoder following [3], however we use a tubelet size of $16 \times 16 \times 4$ instead of $16 \times 16 \times 2$ to reduce complexity. We compare our model with two different initializations for the visual encoder: random and pretrained weights on ImageNet21k. The baseline models are finetuned on the evaluation benchmarks immediately from these initializations whereas we first post-pretrain models in our MV-GPT framework and finetune for action classification.

Table 11 demonstrates that MV-GPT is an effective pretraining strategy for the visual encoder. High-capacity transformer models like ViViT are challenging to train from scratch, and overfit easily as shown in the first row. However, ViViT initialized from an MV-GPT visual encoder trained from scratch performs substantially better, obtaining absolute improvements of 24% on Kinetics-400 (a standard video classification benchmark). This number is close to the performance of ViViT initialized with ImageNet-21K pretraining, as done by the original authors [3] (note that ImageNet-21K was created with high manual annotation cost, while we used no labels at all during pretraining). Finally, initialising the MV-GPT visual encoder with these same ImageNet-21K weights, and then pretraining the MV-GPT visual encoder weights on HowTo100M achieves the best results,

| Method | MSRVTT-QA | ActivityNet-QA |
|----------------------|-------------|----------------|
| SSML [2] | 35.1 | – |
| MAR-VQA [64] | – | 34.6 |
| DECEMBER [47] | 37.4 | – |
| CoMVT [42] | 39.5 | 38.8 |
| VQA-T [57] | 41.5 | 38.9 |
| MV-GPT (Ours) | 41.7 | 39.1 |

Table 9. Comparison to SOTA on MSRVTT-QA and ActivityNet-QA for video question answering. Our method is comparable to other works, even those designed specifically for the task of VideoQA. We compare models pretrained on HowTo100M.

| Method | With ASR | R@1 | R@5 | R@10 | MdR |
|----------------------|----------|-------------|-------------|-------------|----------|
| UniVL [32] | | 21.2 | 49.6 | 63.1 | 6 |
| MMT [12] | | 26.6 | 57.1 | 69.6 | 4 |
| AVLnet [40] | | 27.1 | 55.6 | 66.6 | 4 |
| SSB [36] | | 30.1 | 58.5 | 69.3 | 3 |
| HiT [30] | | 30.7 | 60.9 | 73.2 | – |
| No PT | | 3.5 | 8.0 | 12.1 | 114 |
| Bi-NCE | | 31.6 | 59.0 | 70.2 | 3 |
| MV-GPT (Ours) | | 33.6 | 61.2 | 73.6 | 3 |
| No PT | ✓ | 5.6 | 13.3 | 18.4 | 92 |
| Bi-NCE | ✓ | 33.7 | 61.6 | 73.0 | 3 |
| MV-GPT (Ours) | ✓ | 37.3 | 65.5 | 75.1 | 2 |

Table 10. Comparison to SOTA on MSR-VTT for video retrieval. We compare models pretrained on HowTo100M. $R@k$: Recall at k . **MdR**: Median rank.

| ViViT initialization | Kinetics-400 | Kinetics-600 |
|-----------------------|--------------|--------------|
| Scratch | 50.14 | 55.47 |
| MV-GPT† | 74.20 | 77.10 |
| ImageNet21k [3] | 78.90 | 80.62 |
| ImageNet21k + MV-GPT† | 80.40 | 82.42 |

Table 11. Action classification results on Kinetics with different ViViT initializations. MV-GPT† refers to a model initialised with our MV-GPT pretraining on HowTo100M with *no manually annotated* labels. We use a factorized encoder ViViT-Base following [3], but use a tubelet size of $16 \times 16 \times 4$ instead of $16 \times 16 \times 2$.

improving upon the initialisation of [3] by 1.5% and 1.8% on Kinetics-400 and Kinetics-600 respectively, which is the current state of the art on this dataset with this particular architecture.

5. Conclusion

We present a novel generative pretraining framework for multimodal video captioning. Our bi-directional generative objective jointly trains an encoder for multimodal inputs and a decoder to generate meaningful captions, by using utterances sampled at different times in unlabelled videos. The model is trained end-to-end both during pretraining and finetuning, and achieves state-of-the-art results on multiple video captioning benchmarks as well as on other video understanding tasks, namely VideoQA, video retrieval and action classification.

References

- [1] YouTube Data API. <https://developers.google.com/youtube/v3/docs/captions>. 5
- [2] Elad Amrani, Rami Ben-Ari, Daniel Rotman, and Alex Bronstein. Noise estimation using density estimation for self-supervised multimodal learning. In *AAAI*, 2021. 8
- [3] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViViT: A video vision transformer. In *ICCV*, 2021. 2, 4, 5, 6, 8
- [4] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015. 2
- [5] Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005. 5
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *CVPR*, 2017. 2
- [7] Haoran Chen, Jianmin Li, and Xiaolin Hu. Delving deeper into the decoder for video captioning. In *ECAI*, 2020. 7
- [8] Haoran Chen, Ke Lin, Alexander Maye, Jianmin Li, and Xiaolin Hu. A semantics-assisted video captioning model trained with scheduled sampling. *Frontiers in Robotics and AI*, 7, 2020. 7
- [9] Shaoxiang Chen and Yu-Gang Jiang. Motion guided spatial attention for video captioning. In *AAAI*, 2019. 2, 7
- [10] Pradipto Das, Chenliang Xu, Richard F Doell, and Jason J Corso. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In *CVPR*, 2013. 2
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. 3, 4
- [12] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, 2020. 8
- [13] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, 2010. 8
- [14] Jack Hessel, Bo Pang, Zhenhai Zhu, and Radu Soricut. A case study on combining asr and visual features for generating instructional video captions. In *CoNLL*, 2019. 6
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. 1
- [16] Jingyi Hou, Xinxiao Wu, Wentian Zhao, Jiebo Luo, and Yunde Jia. Joint syntax representation learning and visual cue translation for video captioning. In *ICCV*, 2019. 7
- [17] Gabriel Huang, Bo Pang, Zhenhai Zhu, Clara Rivera, and Radu Soricut. Multimodal pretraining for dense video captioning. In *AAACL*, 2020. 1, 2, 3, 4, 5, 6
- [18] Vladimir Iashin and Esa Rahtu. A better use of audio-visual cues: Dense video captioning with bi-modal transformer. In *BMVC*, 2020. 7
- [19] Vladimir Iashin and Esa Rahtu. Multi-modal dense video captioning. In *CVPRW*, 2020. 7
- [20] 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. 2, 8
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [22] Atsuhiko Kojima, Takeshi Tamura, and Kunio Fukunaga. Natural language description of human activities from video images based on concept hierarchy of actions. *International Journal of Computer Vision*, 50(2):171–184, 2002. 2
- [23] Bruno Korbar, Fabio Petroni, Rohit Girdhar, and Lorenzo Torresani. Video understanding as machine translation. *arXiv preprint arXiv:2006.07203*, 2020. 2, 3, 6, 7
- [24] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *CVPR*, 2017. 5, 7
- [25] Anders Krogh and John A Hertz. A simple weight decay can improve generalization. In *NeurIPS*, 1992. 6
- [26] Jie Lei, Liwei Wang, Yelong Shen, Dong Yu, Tamara Berg, and Mohit Bansal. MART: Memory-augmented recurrent transformer for coherent video paragraph captioning. In *ACL*, 2020. 6
- [27] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. HERO: Hierarchical encoder for video+ language omni-representation pre-training. In *EMNLP*, 2020. 1, 2, 3
- [28] Yehao Li, Ting Yao, Yingwei Pan, Hongyang Chao, and Tao Mei. Jointly localizing and describing events for dense video captioning. In *CVPR*, 2018. 7
- [29] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 2004. 5
- [30] Song Liu, Haoqi Fan, Shengsheng Qian, Yiru Chen, Wenkui Ding, and Zhongyuan Wang. HiT: Hierarchical transformer with momentum contrast for video-text retrieval. *arXiv preprint arXiv:2103.15049*, 2021. 8
- [31] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019. 4
- [32] Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Jason Li, Taroon Bharti, and Ming Zhou. UniVL: A unified video and language pre-training model for multimodal understanding and generation. *arXiv e-prints*, 2020. 1, 2, 3, 4, 5, 6, 7, 8
- [33] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, 2020. 1, 3, 6
- [34] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *ICCV*, 2019. 1, 3, 5, 8
- [35] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: a method for automatic evaluation of machine translation. In *ACL*, 2002. 5
- [36] Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metzger, Alexander G Hauptmann, Joao F. Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. In *ICLR*, 2021. 8

- [37] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *Technical Report*, 2019. 1, 4, 5
- [38] Tanzila Rahman, Bicheng Xu, and Leonid Sigal. Watch, listen and tell: Multi-modal weakly supervised dense event captioning. In *CVPR*, 2019. 7
- [39] Marcus Rohrbach, Wei Qiu, Ivan Titov, Stefan Thater, Manfred Pinkal, and Bernt Schiele. Translating video content to natural language descriptions. In *ICCV*, 2013. 2
- [40] Andrew Rouditchenko, Angie Boggust, David Harwath, Dhiraj Joshi, Samuel Thomas, Kartik Audhkhasi, Rogerio Feris, Brian Kingsbury, Michael Picheny, Antonio Torralba, et al. Avlnet: Learning audio-visual language representations from instructional videos. In *Interspeech*, 2021. 8
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 1
- [42] Paul Hongsuck Seo, Arsha Nagrani, and Cordelia Schmid. Look before you speak: Visually contextualized utterances. In *CVPR*, 2021. 1, 2, 3, 4, 5, 8
- [43] Botian Shi, Lei Ji, Yaobo Liang, Nan Duan, Peng Chen, Zhendong Niu, and Ming Zhou. Dense procedure captioning in narrated instructional videos. In *ACL*, 2019. 6
- [44] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer. *arXiv preprint arXiv:1906.05743*, 2019. 1, 3
- [45] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. VideoBERT: A joint model for video and language representation learning. In *ICCV*, 2019. 1, 2, 3, 4, 6
- [46] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *NeurIPS*, 2014. 2
- [47] Zineng Tang, Jie Lei, and Mohit Bansal. DECEMBER: Learning from noisy instructional videos via dense captions and entropy minimization. In *NAACL*, 2021. 1, 3, 6, 7, 8
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 4
- [49] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. CIDER: Consensus-based image description evaluation. In *CVPR*, 2015. 5
- [50] Bairui Wang, Lin Ma, Wei Zhang, Wenhao Jiang, Jingwen Wang, and Wei Liu. Controllable video captioning with pos sequence guidance based on gated fusion network. In *ICCV*, 2019. 7
- [51] Jingwen Wang, Wenhao Jiang, Lin Ma, Wei Liu, and Yong Xu. Bidirectional attentive fusion with context gating for dense video captioning. In *CVPR*, 2018. 7
- [52] Xin Wang, Yuan-Fang Wang, and William Yang Wang. Watch, listen, and describe: Globally and locally aligned cross-modal attentions for video captioning. In *NAACL-HLT*, 2018. 7
- [53] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning for video understanding. In *ECCV*, 2018. 2, 6
- [54] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In *ACM MM*, 2017. 8
- [55] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, 2016. 2, 5, 8
- [56] I Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification. *arXiv preprint arXiv:1905.00546*, 2019. 1
- [57] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just Ask: Learning to answer questions from millions of narrated videos. In *ICCV*, 2021. 8
- [58] Youngjae Yu, Jongseok Kim, and Gunhee Kim. A joint sequence fusion model for video question answering and retrieval. In *ECCV*, 2018. 8
- [59] Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *AAAI*, 2019. 8
- [60] Junchao Zhang and Yuxin Peng. Object-aware aggregation with bidirectional temporal graph for video captioning. In *CVPR*, 2019. 7
- [61] Ziqi Zhang, Yaya Shi, Chunfeng Yuan, Bing Li, Peijin Wang, Weiming Hu, and Zheng-Jun Zha. Object relational graph with teacher-recommended learning for video captioning. In *CVPR*, 2020. 7
- [62] Luwei Zhou, Chenliang Xu, and Jason Corso. Towards automatic learning of procedures from web instructional videos. In *AAAI*, 2018. 2, 5
- [63] Linchao Zhu and Yi Yang. ActBERT: Learning global-local video-text representations. In *CVPR*, 2020. 1, 3, 6
- [64] Yueting Zhuang, Dejing Xu, Xin Yan, Wenzhuo Cheng, Zhou Zhao, Shiliang Pu, and Jun Xiao. Multichannel attention refinement for video question answering. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 16(1s):1–23, 2020. 8