

Use What You Have: Video Retrieval Using Representations From Collaborative Experts

Yang Liu*

yangl@robots.ox.ac.uk

Samuel Albanie*

albanie@robots.ox.ac.uk

Arsha Nagrani*

arsha@robots.ox.ac.uk

Andrew Zisserman

az@robots.ox.ac.uk

Visual Geometry Group

University of Oxford

UK

Abstract

The rapid growth of video on the internet has made searching for video content using natural language queries a significant challenge. Human-generated queries for video datasets ‘in the wild’ vary a lot in terms of degree of specificity, with some queries describing ‘specific details’ such as the names of famous identities, content from speech, or text available on the screen. Our goal is to condense the multi-modal, extremely high dimensional information from videos into a single, compact video representation for the task of video retrieval using free-form text queries, where the degree of specificity is open-ended.

For this we exploit existing knowledge in the form of pre-trained semantic embeddings which include ‘general’ features such as motion, appearance, and scene features from visual content. We also explore the use of more ‘specific’ cues from ASR and OCR which are intermittently available for videos and find that these signals remain challenging to use effectively for retrieval. We propose a *collaborative experts* model to aggregate information from these different pre-trained experts and assess our approach empirically on five retrieval benchmarks: MSR-VTT, LSMDC, MSVD, DiDeMo, and ActivityNet. Code and data can be found at www.robots.ox.ac.uk/~vgg/research/collaborative-experts/. This paper contains a correction to results reported in the previous version.

1 Introduction

Videos capture the world in two important ways beyond a simple image: first, video contains temporal information – semantic concepts, actions and interactions evolve over time; Second, video may also contain information from multiple modalities, such as an accompanying audio track. This makes videos both richer and more informative, but also more challenging to represent. Our goal in this paper is to embed the information from multiple modalities and multiple time steps of a video segment into a compact fixed-length representation. Such a compact representation can then be used for a number of video understanding tasks, such as video retrieval, clustering and

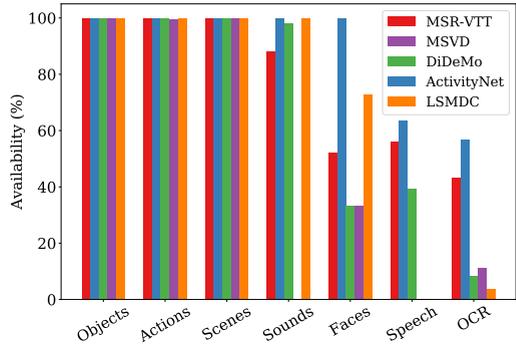


Figure 1: (Left): Unconstrained videos ‘in the wild’ convey information in various different ways, including (clockwise from upper-left), clues from distinctive speech, names of individuals on screen, other text clues embedded in the video and audio. (Right): For the five video datasets considered in this work, the chart portrays the video-level availability of “expert” embeddings from different domains (with potentially multiple experts per domain): certain generic embeddings can almost always be extracted via pretrained object/action/scene classification networks. Other features such as sounds, faces, speech and OCR are less consistently available and are more challenging to exploit (Sec. 4.3).

summarisation. In particular, we focus on retrieval; our objective is to be able to retrieve video clips using a free form text query that may contain both general and specific information.

Learning a robust and compact representation *tabula rasa* for this task is made extremely challenging by the high dimensionality of the sensory data contained in videos—to do so with discriminative training would require prohibitively expensive textual annotation of a vast number of videos. The primary hypothesis underpinning our approach is the following: *the discriminative content of the multi-modal video embedding can be well approximated by the set of semantic representations of the video data learnt by individual experts (in audio, scenes, actions, etc)*. In essence, this approximation enables us to exploit knowledge from existing individual sources where the cost of annotation is significantly reduced (e.g. classification labels for objects and scenes in images, labels for actions in videos etc.) and where consequently, there exist very large-scale labelled datasets. These large-scale datasets can then be used to train independent “experts” for different perception tasks, which in turn provide a robust, low-dimensional basis for the discriminative query-content approximation described above.

The two key aspects of this idea that we explore in this paper are: (i) *General and specific features*: in addition to using generic video descriptors (e.g. objects and actions) we investigate encodings of quite specific information from the clip, for example, text from overlaid captions and text from speech to provide effective coverage of the “queryable content” of the video (Fig. 1, left). While such features may be highly discriminative for humans, they may not always be available (Fig. 1, right) and as we show through experiments (Sec. 4.3), making good use of these cues is challenging. We therefore also propose (ii) *Collaborative experts*: a framework that seeks to make effective use of embeddings from different ‘experts’ (e.g. objects, actions, speech) by learning their combination in order to render them more discriminative. Each expert is filtered via a simple dynamic attention mechanism that considers its relation to all other experts to enable their collaboration. This pairwise approach enables, for instance, the sound of a dog barking to inform the modulation of the RGB features, selecting the features that have encoded the concept of the dog. As we demonstrate in the sequel, this idea yields improvements in the retrieval performance.

Concretely, we make the following three contributions: (i) We propose the *Collaborative*

Experts framework for learning a joint embedding of video and text by combining a collection of pretrained embeddings into a single, compact video representation. Our joint video embeddings are independent of the retrieval text-query and can be pre-computed offline and indexed for efficient retrieval; (ii) We explore the use of both *general* video features such as motion, image classification and audio features, and *specific* video features such as text embedded on screen and speech obtained using OCR and ASR respectively. We find that strong generic features deliver good performance, but that specific, rarely available features remain challenging to use for retrieval¹. (iii) We assess the performance of the representation produced by combining all available cues on a number of retrieval benchmarks, in several cases achieving an advance over prior work.

2 Related Work

Cross-Modal Embeddings: A range of prior work has proposed to jointly embed images and text into the same space [15, 16, 17, 27, 40], enabling cross-modal retrieval. More recently, several works have also focused on audio-visual cross-modal embeddings [2, 39], as well as audio-text embeddings [7]. Our goal in this work, however, is to embed videos and natural language sentences (sometimes multiple sentences) into the same semantic space, which is made more challenging by the high dimensional content of videos.

Video-Text Embeddings: While a large number of works [11, 41, 42, 50, 57] have focused on learning visual semantic embeddings for video and language, many of these existing approaches are based on image-text embedding methods by design and typically focus on single visual frames. Mithun et al. [38] observe that a simple adaptation of a state-of-the-art image-text embedding method [15] by mean-pooling features from video frames provides a better result than many prior video-text retrieval approaches [11, 41]. However, such methods do not take advantage of the rich and varied additional information present in videos, including motion dynamics, speech and other background sounds, which may influence the concepts in human captions to a considerable extent. Consequently, there has been a growing interest in fusing information from other modalities— [35, 38] utilise the audio stream (but do not exploit speech content) and use models pretrained for action recognition to extract motion features. These methods do not make use of speech-to-text or OCR for additional cues, which have nevertheless been used successfully to understand videos in other domains, particularly lecture retrieval [45, 58] (where the videos consist of slide shows) and news broadcast [20] retrieval, where a large fraction of the content is displayed on screen in the form of text. Our approach draws particular inspiration from the powerful joint embedding proposed by [35] (which in turn, builds on the classical Mixtures-of-Experts model [25]) and extends it to investigate additional cues (such as speech and text) and make more effective use of pretrained features via the robust collaborative gating mechanism described in Sec. 3.

Annotation scarcity: A key challenge for video-retrieval is the small size of existing training datasets, due to the high cost of annotating videos with natural language. We therefore propose to use the knowledge from existing embeddings pretrained on a wide variety of other tasks. This idea is not new: semantic projections of visual inputs in the form of ‘experts’ was used by [14] for the task of image retrieval and has also been central to modern video retrieval methods such as [35, 38]. More recently, alternative approaches to addressing the issue of annotation scarcity have been explored, which include self-supervised [49] and weakly-supervised [63] video-text models.

¹Note that this finding differs from the previous version of this paper (see suppl. material).

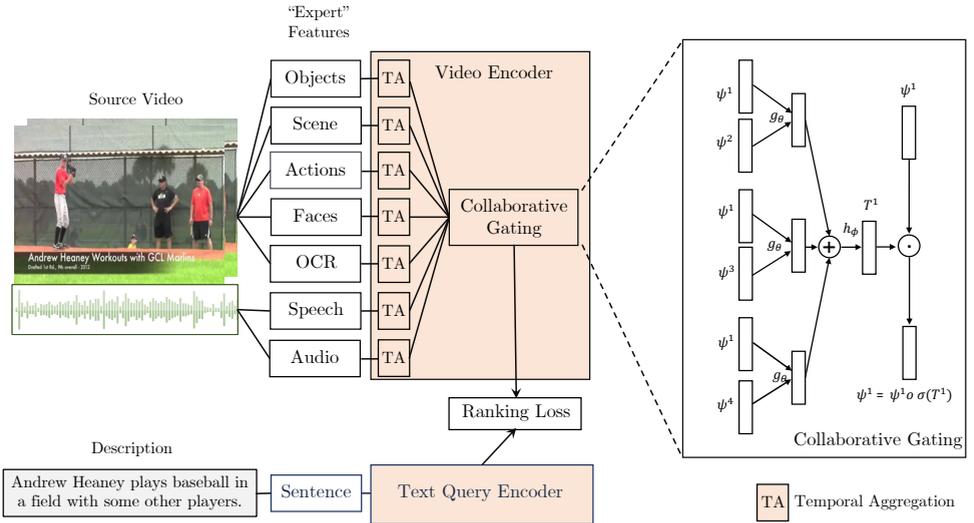


Figure 2: (Left): The proposed Collaborative Experts framework for learning a joint video-text embedding (coloured boxes denote learnable parameters). The information provided by each pretrained “expert” (potentially with multiple experts from a single domain) is temporally aggregated as it enters the video encoder and then refined through the use of a collaborative gating mechanism (right) to obtain the video-embedding (for visual clarity, we show the interaction of just a single expert with three others, though in practice all experts are used—see Sec. 3.1 for details). Note that to maintain retrieval efficiency, collaboration occurs only between video experts (the text-query and video embeddings are computed independently).

3 Collaborative Experts

Given a set of videos with corresponding text captions, we would like to create a pair of functions ϕ_v and ϕ_t that map sensory video data and text into a joint embedding space that respects this correspondence—embeddings for paired text and video should lie close together, while embeddings for text and video that do not match should lie far apart. We would also like ϕ_v and ϕ_t to be independent of each other to enable efficient retrieval: the process of querying then reduces to a distance comparison between the embedding of the query and the embeddings of the collection to be searched (which can be pre-computed offline). The proposed Collaborative Experts framework for learning these functions is illustrated in Fig. 2. In this work, we pay particular attention to the design of the video encoder ϕ_v and the process of combining information from different video modalities (Sec. 3.1). To complete the framework, we then discuss how the query text is encoded and the ranking loss used to learn the joint embedding space (Sec. 3.2).

3.1 Video Encoder

To construct the video encoder ϕ_v , we draw on a collection of pretrained, single-modality experts. These operate on the video sensory data \mathbf{v} and project it to a collection of n variable-length task-specific embeddings $\{\Psi_{\text{var}}^{(1)}(\mathbf{v}), \dots, \Psi_{\text{var}}^{(n)}(\mathbf{v})\}$. Here $\Psi_{\text{var}}^{(i)}$ represents the i^{th} expert (we use the “var” subscript to denote a variable-length output when applied to a sequence of frames)

whose parameters have been learned on a prior task such as object classification and then frozen. Each element of this collection is then aggregated along its temporal dimension to produce fixed-size, task-specific embeddings per video $\{\Psi^{(1)}(\mathbf{v}), \dots, \Psi^{(n)}(\mathbf{v})\}$. Any temporal aggregation function may be used here—in this work, we use simple average pooling to aggregate “slow” visual features such as objects and scenes, and NetVLAD [3] to aggregate more dynamic audio and word features (see Sec. 4.1 for further details). Next, to enable their combination, we apply linear projections to transform these task-specific embeddings to a common dimensionality. Our goal when fusing the resulting representations together into a single condensed video representation is to capture the valuable complementary information between task-specific projections while simultaneously filtering out irrelevant noise and resolving individual expert conflicts on a *per-sample basis*. To do so, this we propose a collaborative gating module, described next.

Collaborative Gating: The collaborative gating module comprises two operations: (1) Prediction of attention vectors for every expert projection $T = \{T^{(1)}(\mathbf{v}), \dots, T^{(n)}(\mathbf{v})\}$; and (2) modulation of expert responses. Inspired by the relational reasoning module proposed by [47] for visual question answering, we define the attention vector of the i^{th} expert projection T_i as follows:

$$T^{(i)}(\mathbf{v}) = h_\phi \left(\sum_{j \neq i} g_\theta(\Psi^{(i)}(\mathbf{v}), \Psi^{(j)}(\mathbf{v})) \right), \quad (1)$$

where functions h_ϕ and g_θ are used to model the pairwise relationship between projection $\Psi^{(i)}$ and projection $\Psi^{(j)}$. Of these, g_θ is used to infer pairwise task relationships, while h_ϕ maps the sum of all pairwise relationships into a single attention vector. In this work, we instantiate both h_ϕ and g_θ as multi-layer perceptrons (MLPs). Note that the functional form of Equation (1) dictates that the attention vector of any expert projection should consider the potential relationships between all pairs associated with this expert. That is to say, the quality of each expert $\Psi^{(j)}$ should contribute in determining and selecting the information content from $\Psi^{(i)}$ in the final decision. It is also worth noting that the collaborative gating module uses the same functions g_θ and h_ϕ (shared weights) to compute all pairwise relationships. This mode of operation encourages greater generalisation, since g_θ and h_ϕ are encouraged not to over-fit to features of any particular pair of tasks. After the attention vectors $T = \{T^{(1)}(\mathbf{v}), \dots, T^{(n)}(\mathbf{v})\}$ have been computed, each expert projection is modulated follows:

$$\Psi^{(i)}(\mathbf{v}) = \Psi^{(i)}(\mathbf{v}) \circ \sigma(T^{(i)}(\mathbf{v})), \quad (2)$$

where σ is an element-wise sigmoid activation and \circ is the element-wise multiplication (Hadamard product). This gating function re-calibrates the strength of different activations of $\Psi^{(i)}(\mathbf{v})$ and selects which information is highlighted or suppressed, providing the model with a powerful mechanism for dynamically filtering content from different experts. A diagram of the mechanism is shown in Fig. 2 (right). The final video embedding is then obtained by passing the modulated responses of each expert through a Gated Embedding Module (GEM) [35] (note that this operation produces 12-normalized outputs) before concatenating the outputs together into a single fixed-length vector.

3.2 Text Query Encoder and Training Loss

To construct the text embeddings, query sentences are first mapped to a sequence of feature vectors with pretrained contextual word-level embeddings (see Sec. 4.1 for details)—as with the video experts, the parameters of this first stage are frozen. These are then aggregated, again using

NetVLAD [3]. Following aggregation, we follow the text encoding architecture proposed by [35], which projects the aggregated features to separate subspaces for each expert using GEMs (as with the video encoder, producing l2-normalized outputs). Each projection is then scaled by a mixture weight (one scalar weight per expert projection), which is computed by applying a single linear layer to the aggregated text-features, and passing the result through a softmax to ensure that the mixture weights sum to one (see [35] for further details). Finally, the scaled outputs are concatenated, producing a vector of dimensionality that matches that of the video embedding.

With the video encoder ϕ_v and text encoder ϕ_t as described, the similarity s_i^j of the i^{th} video, \mathbf{v}_i , and the j^{th} caption, \mathbf{t}_j , can then be directly computed as the cosine of the angle between their respective embeddings $\phi_v(\mathbf{v}_i)^T \phi_t(\mathbf{t}_j)$. During optimisation, the parameters of the video encoder (including the collaborative gating module) and text query encoder (the coloured regions of Fig. 2) are learned jointly. Training proceeds by sampling a sequence of minibatches of corresponding video-text pairs $\{\mathbf{v}_i, \mathbf{t}_i\}_{i=1}^{N_B}$ and minimising a *Bidirectional Max-margin Ranking Loss* [48]:

$$\mathcal{L}_r = \frac{1}{N_B} \sum_{i=1, j \neq i}^{N_B} \max(0, m + s_i^j - s_i^i) + \max(0, m + s_j^i - s_j^j) \quad (3)$$

where N_B is the batch size, and m is a fixed constant which is set as a hyperparameter. When assessing retrieval performance, at test time the embedding distances are simply computed via their inner product, as described above.

3.2.1 Missing Experts

When a set of expert features are missing, such as when there is no speech in the audio track, we simply zero-pad the missing experts when estimating the similarity score. To compensate for the implicit scaling introduced by missing experts (the similarity is effectively computed between shorter embeddings), we follow the elegant approach proposed by [35] and simply remove the mixture weights for missing experts, then renormalise the remaining weights such that they sum to one.

4 Experiments

In this section, we evaluate our model on five benchmarks for video retrieval tasks. The description of datasets, implementation details and evaluation metric are provided in Sec. 4.1. A comprehensive comparison on general video retrieval benchmarks is reported in Sec. 4.2. We present an ablation study in Sec. 4.3 to explore how the performance of the proposed method is affected by different model configurations, including the aggregation methods, importance of different experts and number of captions in training.

4.1 Datasets, Implementation Details and Metrics

Datasets: We perform experiments on five video datasets: MSR-VTT [56], LSMDC [46], MSVD [8], DiDeMo [1] and ActivityNet-captions [29], covering a challenging set of domains which include videos from YouTube, personal collections and movies.

Expert Features: In order to capture the rich content of a video, we draw on existing powerful representations for a number of different semantic tasks. These are first extracted at a frame-level, then aggregated to produce a single feature vector per modality per video. *RGB “object”* frame-level embeddings of the visual data are generated with two models: an SENet-154 model [23]

(pretrained on ImageNet for the task of image classification), and a ResNext-101 [55] pretrained on Instagram hashtags [33]. *Motion* embeddings are generated using the I3D inception model [6] and a 34-layer R(2+1)D model [51] trained on IG-65m [18]. *Face* embeddings are extracted in two stages: (1) Each frame is passed through an SSD face detector [4, 31] to extract bounding boxes; (2) The image region of each box is passed through a ResNet50 [21] that has been trained for the task of face classification on the VGGFace2 dataset [5]. *Audio* embeddings are obtained with a VGGish model, trained for audio classification on the YouTube-8m dataset [22]. *Speech-to-Text* features are extracted using the Google Cloud speech API, to extract word tokens from the audio stream, which are then encoded via pretrained word2vec embeddings [37]. *Optical Character Recognition* is done in two stages: (1) Each frame is passed through the Pixel Link [9] text detection model to extract bounding boxes for text; (2) The image region of each box is passed through a model [32] that has been trained for scene text recognition on the Synth90K dataset [24]. The text is then encoded via a pretrained word2vec embedding model [37].

Temporal Aggregation: We adopt a simple approach to aggregating the features described above. For appearance, motion, scene and face embeddings, we average frame-level features along the temporal dimension to produce a single feature vector per video (we found max-pooling to perform similarly). For speech, audio and OCR features, we adopt the NetVLAD mechanism proposed by [3], which has proven effective in the retrieval setting [34]. As noted in Sec. 3.1, all aggregated features are projected to a common size (768 dimensions).

Text: Each word is encoded using pretrained word2vec word embeddings [37] and then passed through a pretrained OpenAI-GPT model [44] to extract contextual word embeddings. Finally, the word embeddings in each sentence are aggregated using NetVLAD.

Dataset-specific details: Except where noted otherwise for ablation purposes, we use each of the embeddings described above for the MSR-VTT, ActivityNet and DiDeMo datasets. For MSVD, we extract the subset of features which do not require an audio stream (since no audio is available with the dataset). For LSMDC, we re-use the existing face, text and audio features made available by [35], and combine them with the remaining features described above.

Training Details: The CE framework is implemented with PyTorch [43]. Optimisation is performed with the Lookahead solver [26] in combination with RAdam [30] (implementation by [54]). Optimisation settings and the hyperparameter selection procedure is described in the appendix.

Evaluation Metrics: We follow prior work (e.g. [11, 35, 38, 60, 61]) and report standard retrieval metrics (where existing work enables comparison) including median rank (lower is better), mean rank (lower is better) and R@K (recall at rank K—higher is better). When computing video-to-sentence metrics for datasets with multiple independent sentences per video (MSR-VTT and MSVD), we follow the evaluation protocol used in prior work [12, 13, 38] which corresponds to reporting the minimum rank among all valid text descriptions for a given video query. For each benchmark, we report the mean and standard deviation of three randomly seeded runs.

4.2 Comparison to Prior State-of-the-Art

We first compare the proposed method with the existing state-of-the-art on the MSR-VTT benchmark for the tasks of sentence-to-video and video-to-sentence retrieval Tab. 1. Driven by strong expert features, we observe that Collaborative Experts (CE) consistently improves retrieval performance for both sentence and video queries. We next evaluate the performance of the CE framework on the LSMDC benchmark for sentence-to-video retrieval (Tab. 2, left) and observe that CE matches or outperforms all prior work, including the prior state-of-the-art method [35] which incorporates additional training images and captions from the COCO benchmark during training, but uses fewer experts. We observe similar trends in the results for the MSVD retrieval benchmark (Tab. 2, right).

Method	Test-set	Text \Rightarrow Video					Video \Rightarrow Text				
		R@1	R@5	R@10	MdR	MnR	R@1	R@5	R@10	MdR	MnR
JSFusion [60]	1k-A	10.2	31.2	43.2	13	-	-	-	-	-	-
CE	1k-A	20.9\pm1.2	48.8\pm0.6	62.4\pm0.8	6\pm0	28.2\pm0.8	20.6\pm0.6	50.3\pm0.5	64.0\pm0.2	5.3\pm0.6	25.1\pm0.8
MoEE [35]	1k-B	13.6	37.9	51.0	10	-	-	-	-	-	-
MoEE _{coco} [35]	1k-B	14.2	39.2	53.8	9	-	-	-	-	-	-
CE	1k-B	18.2\pm0.7	46.0\pm0.4	60.7\pm0.2	7\pm0	35.3\pm1.1	18.0\pm0.8	46.0\pm0.5	60.3\pm0.5	6.5\pm0.5	30.6\pm1.2
VSE [38]	Full	5.0	16.4	24.6	47	215.1	7.7	20.3	31.2	28	185.8
VSE++ [38]	Full	5.7	17.1	24.8	65	300.8	10.2	25.4	35.1	25	228.1
Mithun et al. [38]	Full	7.0	20.9	29.7	38	213.8	12.5	32.1	42.4	16	134.0
W2VV [12]	Full	6.1	18.7	27.5	45	-	11.8	28.9	39.1	21	-
Dual Enc. [13]	Full	7.7	22.0	31.8	32	-	13.0	30.8	43.3	15	-
E2E [36]	Full	9.9	24.0	32.4	29.5	-	-	-	-	-	-
CE	Full	10.0\pm0.1	29.0\pm0.3	41.2\pm0.2	16\pm0	86.8\pm0.3	15.6\pm0.3	40.9\pm1.4	55.2\pm1.0	8.3\pm0.6	38.1\pm1.8

Table 1: Retrieval with sentences and videos on the MSR-VTT dataset. R@k denotes recall@k (higher is better), MdR and MnR denote median rank and mean rank resp. (lower is better). Standard deviations are reported from three randomly seeded runs. 1k-A and 1k-B denote test sets of 1000 randomly sampled text-video pairs used by [60] and [35] resp.

In Tab. 4, we compare with prior work on the ActivityNet paragraph-video retrieval benchmark (note that we compare to methods which use the same level of annotation as our approach i.e. video-level annotation), and see that CE is competitive. Finally, in Tab. 3 we provide a comparison with previously reported numbers on the DiDeMo benchmark and see that CE again outperforms prior work.

Method	Text \Rightarrow Video				Text \Rightarrow Video					
	R@1	R@5	R@10	MdR	R@1	R@5	R@10	MdR	MnR	
Yu et al. [59]†	3.6	14.7	23.9	50	-	-	-	-	245.3	
CCA [28] (rep. by [35])	7.5	21.7	31.0	33	-	-	-	-	236.3	
JSFusion [60]‡	9.1	21.2	34.1	36	VSE [27] ([38])	12.3	30.1	42.3	14	57.7
MoEE [35]	9.3	25.1	33.4	27	VSE++ [15] ([38])	15.4	39.6	53.0	9	43.8
MoEE _{coco} [35]	10.1	25.6	34.6	27	Multi. Cues [38]	20.3	47.8	61.1	6	28.3
CE	11.2\pm0.4	26.9\pm1.1	34.8\pm2.0	25.3\pm3.1	CE	19.8\pm0.3	49.0\pm0.3	63.8\pm0.1	6\pm0	23.1\pm0.3

Table 2: Text-to-Video retrieval results on the LSMDC dataset (left) and the MSVD dataset (right). †, ‡ denote the winners of the 2016 and 2017 LSMDC challenges, respectively.

Method	Text \Rightarrow Video					Video \Rightarrow Text				
	R@1	R@5	R@50	MdR	MnR	R@1	R@5	R@50	MdR	MnR
S2VT [52] ([61])	11.9	33.6	76.5	13	-	13.2	33.6	76.5	15	-
FSE [61]	13.9 \pm 0.7	36 \pm 0.8	78.9 \pm 1.6	11	-	13.1 \pm 0.5	33.9 \pm 0.4	78.0 \pm 0.8	12	-
CE	16.1\pm1.4	41.1\pm0.4	82.7\pm0.3	8.3\pm0.6	43.7\pm3.6	15.6\pm1.3	40.9\pm0.4	82.2\pm1.3	8.2\pm0.3	42.4\pm3.3

Table 3: Comparison of paragraph-video retrieval methods trained with video-level information on the DiDeMo dataset.

4.3 Ablation Studies

In this section, we provide ablation studies to empirically assess: (1) the effectiveness of the proposed collaborative experts framework vs other aggregation strategies; (2) the importance of using of a diverse range of experts with differing levels of specificity; (3) the relative value of using experts in comparison to simply having additional annotated training data.

Method	Text \implies Video					Video \implies Text				
	R@1	R@5	R@50	MdR	MnR	R@1	R@5	R@50	MdR	MnR
LSTM-YT [53] ((61))	0.0	4.0	24.0	102	-	0.0	7.0	38.0	98	-
NOCTXT [52] ((61))	5.0	14.0	32.0	78	-	7.0	18.0	45.0	56	-
DENSE [29]	14.0	32.0	65.0	34	-	18.0	36.0	74.0	32	-
FSE [61]	18.2 \pm 0.2	44.8 \pm 0.4	89.1 \pm 0.3	7	-	16.7 \pm 0.8	43.1 \pm 1.1	88.4 \pm 0.3	7	-
HSE(4SEGS) [61] [†]	20.5	49.3	-	-	-	18.7	48.1	-	-	-
CE	18.2 \pm 0.3	47.7 \pm 0.6	91.4 \pm 0.4	6 \pm 0	23.1 \pm 0.5	17.7 \pm 0.6	46.6 \pm 0.7	90.9 \pm 0.2	6 \pm 0	24.4 \pm 0.5

Table 4: Comparison of paragraph-video retrieval methods trained with video-level information on the ActivityNet-captions dataset (vall test-split).

Aggregation method: We compare the use of collaborative experts with several other baselines (with access to the same experts) for embedding aggregation including: (1) simple expert concatenation; (2) CE without projecting to a common dimension, without mixture weights and without the collaborative gating module described in Sec. 3.1; (3) the state of the art MoEE [35] method (equivalent to CE without the common projection and collaborative gating) and (4) CE without collaborative gating. The results, presented in Tab. 6 (left), demonstrate the contribution of collaborative gating which improves performance and leads to a more efficient parameterisation than the prior state of the art.

Importance of different experts: The value of different experts is assessed in Tab. 5 (note that since several experts are not present in all videos, we combine them with features produced by a “scene” expert pretrained on Places365 [62]—the expert with the lowest performance that is consistently available as a baseline to enable a more meaningful comparison). There is considerable variance in the effect produced by different choices of expert. Using stronger features within a given modality (pretraining on Instagram [33] rather than Kinetics [6] (resp. ImageNet) [10] for actions (resp. object) experts can yield a significant boost in performance). The cues from scarce features (such as speech, face and OCR) which are often missing from videos (see Fig. 1, right) provide significantly weaker cues and bring a limited improvement to performance when used in combination.

Experts	Text \implies Video					Experts	Text \implies Video				
	R@1	R@5	R@10	MdR	MnR		R@1	R@5	R@10	MdR	MnR
Scene	4.0 \pm 0.1	14.1 \pm 0.1	22.4 \pm 0.3	50.0 \pm 1.0	201.3 \pm 1.6	Scene	4.0 \pm 0.1	14.1 \pm 0.1	22.4 \pm 0.3	50.0 \pm 1.0	201.3 \pm 1.6
Scene+Speech	4.6 \pm 0.1	15.5 \pm 0.2	24.4 \pm 0.2	44.7 \pm 1.2	183.6 \pm 1.7	Prev.+Speech	4.6 \pm 0.1	15.5 \pm 0.2	24.4 \pm 0.2	44.7 \pm 1.2	183.6 \pm 1.7
Scene+Audio	5.6 \pm 0.0	18.7 \pm 0.1	28.2 \pm 0.1	33.7 \pm 0.6	140.8 \pm 0.3	Prev.+Audio	5.8 \pm 0.1	19.0 \pm 0.3	28.8 \pm 0.2	32.3 \pm 0.6	136.8 \pm 1.2
Scene+Action(KN)	5.3 \pm 0.3	17.6 \pm 0.8	27.1 \pm 0.9	36.0 \pm 1.7	158.7 \pm 1.6	Prev.+Action(KN)	6.7 \pm 0.2	21.8 \pm 0.4	32.5 \pm 0.5	25.3 \pm 0.6	115.9 \pm 1.0
Scene+Obj(IN)	5.0 \pm 0.2	16.6 \pm 0.7	25.5 \pm 1.0	40.7 \pm 2.1	173.1 \pm 3.3	Prev.+Obj(IN)	7.5 \pm 0.1	23.4 \pm 0.0	34.1 \pm 0.2	23.7 \pm 0.6	111.9 \pm 0.6
Scene+Obj(IG)	7.2 \pm 0.1	22.3 \pm 0.3	33.0 \pm 0.2	25.3 \pm 0.6	125.1 \pm 0.1	Prev.+Obj(IG)	9.5 \pm 0.2	27.7 \pm 0.1	39.4 \pm 0.1	18.0 \pm 0.0	92.6 \pm 0.4
Scene+Action(IG)	6.8 \pm 0.1	21.7 \pm 0.1	32.4 \pm 0.1	25.7 \pm 0.6	122.1 \pm 0.3	Prev.+Action(IG)	9.9 \pm 0.1	28.6 \pm 0.3	40.7 \pm 0.1	17.0 \pm 0.0	86.4 \pm 0.4
Scene+OCR	4.1 \pm 0.1	14.1 \pm 0.1	22.2 \pm 0.2	50.3 \pm 1.2	203.1 \pm 4.4	Prev.+OCR	10.0 \pm 0.1	28.8 \pm 0.2	40.9 \pm 0.2	16.7 \pm 0.6	87.3 \pm 0.8
Scene+Face	4.1 \pm 0.1	14.2 \pm 0.3	22.4 \pm 0.4	49.7 \pm 0.6	194.2 \pm 5.1	Prev.+Face	10.0 \pm 0.1	29.0 \pm 0.3	41.2 \pm 0.2	16.0 \pm 0.0	86.8 \pm 0.3

Table 5: **The importance of different experts** (Left): The value of different experts in combination with a baseline set for text-video retrieval and (right) their cumulative effect on MSR-VTT (here Prev. denotes the experts used in the previous row).

Number of Captions in training: An emerging idea in our community is that many machine perception tasks might be solved through the combination of simple models and large-scale training sets, reminiscent of the “big-data” hypothesis [19]. In this section, we perform an ablation study to assess the relative importance of access to pretrained experts and additional video description annotations. To do so, we measure the performance of the CE model as we vary (1) the number of descriptions available per-video during training and (2) the number of experts it has access to. The results are shown in Tab. 6 (right). We observe that increasing the number of training captions

Aggreg.	R@1	R@5	R@10	MdR	Params
Concat	0.0±0.0	0.0±0.0	0.0±0.0	1495.5±0.0	369.72k
CE - MW,PCG	8.5±0.1	25.9±0.3	37.6±0.2	19.0±0.0	246.22M
MoEE [35]	9.6±0.1	28.0±0.2	39.7±0.2	17.7±0.6	400.41 M
CE - CG	9.7±0.1	28.1±0.2	40.2±0.1	17.0±0.0	181.07 M
CE	10.0±0.1	29.0±0.3	41.2±0.2	16.0±0.0	183.45 M

Expert	Num. Captions	R@1	R@5	R@10	MdR
Obj(IN)	1	2.6±0.1	9.3±0.4	15.0±0.7	101.3±15.5
Obj(IN)	20	4.9±0.1	16.5±0.2	25.3±0.4	40.7±1.2
All	1	4.8±0.2	16.2±0.5	25.0±0.7	43.3±4.0
All	20	10.0±0.1	29.0±0.3	41.2±0.2	16.0±0.0

Table 6: (Left): Aggregation methods for text-video retrieval on MSR-VTT; (Right): The relative value of training with additional captions vs the value of experts.

per-video from 1 to 20 brings an improvement in performance, approximately comparable to adding the full collection of experts, suggesting that indeed, adding experts can help to compensate for a paucity of labelled data. When multiple captions and multiple experts are both available, they naturally lead to the most robust embedding. Some qualitative examples of videos retrieved by the multiple-expert, multiple-caption system are provided in Fig. 3.

Query: Guy working on his engine with multiple parts
(GT rank: 29)



Similarity: 0.60

Similarity: 0.59

Similarity: 0.55

Query: shiny black sports car drives very slowly down road through orange and white safety cones (GT rank: 10)



Similarity: 0.48

Similarity: 0.46

Similarity: 0.46

Query: Awareness of mosquito bites by doctors
(GT rank: 2)



Similarity: 0.38

Similarity: 0.36

Similarity: 0.34

Query: Query: Awareness of mosquito bites by doctors
(without CE) - (GT rank: 7)



Similarity: 0.34

Similarity: 0.33

Similarity: 0.23

Figure 3: **Qualitative Results on MSR-VTT:** For each query, we show frames from the top three ranked videos (where present, the ground truth video is indicated by a green box around the similarity score). Top row: (left) Even for imperfect rankings, the model retrieves reasonable videos; Failure case (right) the embeddings can fail to differentiate between certain signals (in this case, ranking cars of the wrong colour above the ground truth video). Bottom row: (left) the videos retrieved by the proposed model (which assigns its second highest similarity to the correct video); (right) removing the proposed CE component produces a noisier ranking.

5 Conclusion

In this work, we introduced collaborative experts, a framework for learning a joint video-text embedding for efficient retrieval. We have shown that using a range of pretrained features and combining them through an appropriate gating mechanism can boost retrieval performance. In future work, we plan to explore the use of collaborative experts for other video understanding tasks such as clustering and summarisation.

Acknowledgements: Funding for this research is provided by the EPSRC Programme Grant Seebibyte EP/M013774/1 and EPSRC grant EP/R03298X/1. A.N. is supported by a Google PhD Fellowship. We would like to thank Antoine Miech, YoungJae Yu and Bowen Zhang for their assistance with experiment details. We would like to particularly thank Valentin Gabeur for identifying a bug in the software implementation that was responsible for the inaccurate results reported in the initial version of the paper. We would also like to thank Zak Stone and Susie Lim for their help with cloud computing.

References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5803–5812, 2017.
- [2] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017.
- [3] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5297–5307, 2016.
- [4] G. Bradski. The OpenCV Library. *Dr. Dobbs's Journal of Software Tools*, 2000.
- [5] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. VGGFace2: A dataset for recognising faces across pose and age. In *Proc. Int. Conf. Autom. Face and Gesture Recog.*, 2018.
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [7] Gal Chechik, Eugene Ie, Martin Rehn, Samy Bengio, and Dick Lyon. Large-scale content-based audio retrieval from text queries. In *Proceedings of the 1st ACM international conference on Multimedia information retrieval*, pages 105–112. ACM, 2008.
- [8] David L Chen and William B Dolan. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 190–200. Association for Computational Linguistics, 2011.
- [9] Dan Deng, Haifeng Liu, Xuelong Li, and Deng Cai. Pixellink: Detecting scene text via instance segmentation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [11] Jianfeng Dong, Xirong Li, and Cees GM Snoek. Word2visualvec: Image and video to sentence matching by visual feature prediction. *arXiv preprint arXiv:1604.06838*, 2016.
- [12] Jianfeng Dong, Xirong Li, and Cees GM Snoek. Predicting visual features from text for image and video caption retrieval. *IEEE Transactions on Multimedia*, 20(12):3377–3388, 2018.
- [13] Jianfeng Dong, Xirong Li, Chaoxi Xu, Shouling Ji, and Xun Wang. Dual dense encoding for zero-example video retrieval. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [14] Matthijs Douze, Arnau Ramisa, and Cordelia Schmid. Combining attributes and fisher vectors for efficient image retrieval. In *CVPR 2011*, pages 745–752. IEEE, 2011.
- [15] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improved visual-semantic embeddings. *arXiv preprint arXiv:1707.05612*, 2(7):8, 2017.
- [16] Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. Every picture tells a story: Generating sentences from images. In *European conference on computer vision*, pages 15–29. Springer, 2010.

- [17] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, pages 2121–2129, 2013.
- [18] Deepti Ghadiyaram, Du Tran, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12046–12055, 2019.
- [19] Alon Halevy, Peter Norvig, and Fernando Pereira. The unreasonable effectiveness of data. 2009.
- [20] Alexander G Hauptmann, Rong Jin, and Tobun Dorbin Ng. Multi-modal information retrieval from broadcast video using ocr and speech recognition. In *Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries*, pages 160–161. ACM, 2002.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [22] Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron Weiss, and Kevin Wilson. Cnn architectures for large-scale audio classification. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2017. URL <https://arxiv.org/abs/1609.09430>.
- [23] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Enhua Wu. Squeeze-and-excitation networks. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [24] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. *arXiv preprint arXiv:1406.2227*, 2014.
- [25] Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [27] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014.
- [28] Benjamin Klein, Guy Lev, Gil Sadeh, and Lior Wolf. Associating neural word embeddings with deep image representations using fisher vectors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4437–4446, 2015.
- [29] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 706–715, 2017.
- [30] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
- [31] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [32] Yang Liu, Zhaowen Wang, Hailin Jin, and Ian Wassell. Synthetically supervised feature learning for scene text recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 435–451, 2018.

- [33] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 181–196, 2018.
- [34] Antoine Miech, Ivan Laptev, and Josef Sivic. Learnable pooling with context gating for video classification. *arXiv preprint arXiv:1706.06905*, 2017.
- [35] Antoine Miech, Ivan Laptev, and Josef Sivic. Learning a text-video embedding from incomplete and heterogeneous data. *arXiv preprint arXiv:1804.02516*, 2018.
- [36] 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. *arXiv preprint arXiv:1912.06430*, 2019.
- [37] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [38] Niluthpol Chowdhury Mithun, Juncheng Li, Florian Metze, and Amit K Roy-Chowdhury. Learning joint embedding with multimodal cues for cross-modal video-text retrieval. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, pages 19–27. ACM, 2018.
- [39] Arsha Nagrani, Samuel Albanie, and Andrew Zisserman. Learnable pins: Cross-modal embeddings for person identity. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 71–88, 2018.
- [40] Hyeonseob Nam, Jung-Woo Ha, and Jeonghee Kim. Dual attention networks for multimodal reasoning and matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 299–307, 2017.
- [41] Mayu Otani, Yuta Nakashima, Esa Rahtu, Janne Heikkilä, and Naokazu Yokoya. Learning joint representations of videos and sentences with web image search. In *European Conference on Computer Vision*, pages 651–667. Springer, 2016.
- [42] Yingwei Pan, Tao Mei, Ting Yao, Houqiang Li, and Yong Rui. Jointly modeling embedding and translation to bridge video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4594–4602, 2016.
- [43] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [44] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018. URL https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf.
- [45] N Radha. Video retrieval using speech and text in video. In *2016 International Conference on Inventive Computation Technologies (ICICT)*, volume 2, pages 1–6. IEEE, 2016.
- [46] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3202–3212, 2015.
- [47] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976, 2017.

- [48] Richard Socher, Andrej Karpathy, Quoc V Le, Christopher D Manning, and Andrew Y Ng. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics*, 2:207–218, 2014.
- [49] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. *arXiv preprint arXiv:1904.01766*, 2019.
- [50] Atousa Torabi, Niket Tandon, and Leonid Sigal. Learning language-visual embedding for movie understanding with natural-language. *arXiv preprint arXiv:1609.08124*, 2016.
- [51] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018.
- [52] Subhashini Venugopalan, Huijuan Xu, Jeff Donahue, Marcus Rohrbach, Raymond Mooney, and Kate Saenko. Translating videos to natural language using deep recurrent neural networks. *arXiv preprint arXiv:1412.4729*, 2014.
- [53] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. Sequence to sequence-video to text. In *Proceedings of the IEEE international conference on computer vision*, pages 4534–4542, 2015.
- [54] Less Wright. Project title. <https://github.com/lessw2020/Ranger-Deep-Learning-Optimizer>, 2019.
- [55] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [56] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5288–5296, 2016.
- [57] Ran Xu, Caiming Xiong, Wei Chen, and Jason J Corso. Jointly modeling deep video and compositional text to bridge vision and language in a unified framework. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [58] Natsuo Yamamoto, Jun Ogata, and Yasuo Ariki. Topic segmentation and retrieval system for lecture videos based on spontaneous speech recognition. In *Eighth European Conference on Speech Communication and Technology*, 2003.
- [59] Youngjae Yu, Hyungjin Ko, Jongwook Choi, and Gunhee Kim. Video captioning and retrieval models with semantic attention. *arXiv preprint arXiv:1610.02947*, 6(7), 2016.
- [60] Youngjae Yu, Jongseok Kim, and Gunhee Kim. A joint sequence fusion model for video question answering and retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 471–487, 2018.
- [61] Bowen Zhang, Hexiang Hu, and Fei Sha. Cross-modal and hierarchical modeling of video and text. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 374–390, 2018.
- [62] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [63] Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, and Josef Sivic. Cross-task weakly supervised learning from instructional videos. *arXiv preprint arXiv:1903.08225*, 2019.