Learning Language-Visual Embedding for Movie Understanding with Natural-Language

Atousa Torabi[1], Niket Tandon[2], and Leonid Sigal[1]

[1] Disney Research Pittsburgh
[2] Max Planck Institute for Informatics
{atousa.torabi,lsigal}@disneyresearch.com
ntandon@mpi-inf.mpg.de

Abstract. Learning a joint language-visual embedding has a number of very appealing properties and can result in variety of practical application, including natural language image/video annotation and search. In this work, we study three different joint language-visual neural network model architectures. We evaluate our models on large scale LSMDC16 [17,18] movie dataset for two tasks: 1) Standard Ranking for video annotation and retrieval 2) Our proposed movie multiple-choice test. This test facilitate automatic evaluation of visual-language models for natural language video annotation based on human activities. In addition to original Audio Description (AD) captions, provided as part of LSMDC16, we collected and will make available a) manually generated re-phrasings of those captions obtained using Amazon MTurk b) automatically generated human activity elements in "Predicate + Object" (PO) phrases based on "Knowlywood", an activity knowledge mining model [22]. Our best model archives Recall@10 of 19.2% on annotation and 18.9% on video retrieval tasks for subset of 1000 samples. For multiple-choice test, our best model achieve accuracy 58.11% over whole LSMDC16 public test-set.

Keywords: Video Annotation, Video Retrieval, learning Joint Visual-Language Embedding, LSMT, Soft Attention Network

1   Introduction

Natural language-based video and image search has been a long standing topic of research among information retrieval, multimedia, and computer vision communities. Several existing on-line platforms (e.g. Youtube) rely on massive human curation efforts, manually assigned tags, click counts and surrounding text to match largely unstructured search phrases in order to retrieve ranked list of relevant videos from a stored library. However, as the amount of unlabeled video content grows, with advent of inexpensive mobile recording devices (e.g. smart phones), the focus is rapidly shifting to automated understand, tagging and search.

Over the last 1-2 years, there has been an increased interest in jointly modeling images/videos and natural language sentences. Models that jointly learn

from videos/images and natural language sentences have broad applicability to visual search, retrieval [11,24], captioning [4,30,26], or visual question answering [1,6,14] tasks. There are also abundant evidence [19,27] that jointly learning from language and visual modalities can be mutually beneficial. These recent trends of multi-modal learning are, at least in part, are enabled by recent advances in deep learning, with high capacity Convolutional Neural Networks (CNNs) driving up performance on image and video recognition [21,20,28] end, and Recurrent Neural networks (RNNs), particularly Long Short-Term Memory (LSTM), making advances in natural language translation and sentence generation. The learning of such high capacity models, with millions and tens of millions of parameters, is made possible by availability of large scale image- (e.g. COCO [13], Visual Genome[3]) and, more recently, video- (LSMDC [18], MSR-VTT [29]) description datasets, where an image or video is associated with one-or-more human generated natural language sentence descriptions.

While there are a number of recent works that look at joint image-language modeling, the progress on *video*-language models has been much more limited. The difficulty stems from additional challenges that come from the need to encode the temporal aspects of the video and the sheer volume of the video data required to be processed (e.g. Large Scale Movie Description Challenge (LSMDC) [18] dataset contains nearly 60-times the number of frames as images in COCO [13]). In the video-language domain, the few works that exist focus largely on video description generation [25,31]. These models typically use LSTM models to generate sentences given an encoding of the video. Evaluating the video description performance especially on LSMDC that contain audio description (AD) captions and are typically relatively verbose and very precise is not easy and usually have done based on human judgment. In this work we study two tasks: 1) Standard Ranking for video annotation and retrieval 2) Multiple-choice test, which enable us to automatically evaluate joint language-visual model based on precise matrics: recall and accuracy. We provide a baseline results for LSMDC16 video annotation and search and multiple-choice test data, that we built for LSMDC16.

The rest of the paper is organized as the following, in section 2, we describe exisiting joint visual-language learning related works, in section 3, we explain how we built LSMDC16 multiple-choice test. Section 4 describes our joint visual-language models structure and our experiments is presented in section 5.

2   Related Work

Visual Caption Generation: There is an increasing interest in jointly learning from images/videos and natural language sentences. In the last 1-2 years, several works have been developed for automatic image [3,4,5,10,11,30,26] and video [25,31] captioning. Most of these models use Recurrent Neural Networks (RNN), or LSTMs, in an encoder-decoder architecture (i.e. CNNs or CNN+LSTM encoder, for images or videos respectively, to generate a hidden semantic state and

---

[3] https://visualgenome.org/

RNN/LSTM decoder that decodes resulting hidden state using a trained language model to produce a final sentence description). Automatic captioning is a very challenging task, both from an algorithmic and evaluation point of view. The latter is particularly challenging with difficulties arising from evaluation of specificity, linguistic correctness and relevance of generated captions.

Visual Question Answering: Because of the aforementioned challenges, image [1,6,14] and, more recently, video [32] Visual Question Answering (VQA) has became a preferred alternative to caption generation. VQA datasets [1,6,23,32] consist of large number of structured (e.g. multiple choice or fill-in-the-blank) and unstructured (e.g. free form) image-specific questions with corresponding answers. As a result, evaluation in VQA setting tends to be considerably more objective, requiring algorithms to have, at a minimum, certain level of visual understanding to answer the poised questions.

Image-Caption Retrieval: Another alternative is image-caption retrieval that has been defined as a standard way to evaluate joint language-visual models [8,11,12,24]. The core idea is to rank a set of images according to their relevance to a caption query (*a.k.a*, image retrieval) or ranking captions according to their relevance to the given image query (*a.k.a*, caption retrieval). Image-caption retrieval approaches, typically, learn a joint embedding space that minimizes a pair-wise ranking objective function between images and captions. Particularly relevant to our paper, is the work of [24], where it is acknowledged that different forms of captions form a visual-semantic hierarchy and the order-preserving constraints are used as objective for learning.

Fig. 1. An example of LSMDC16 multiple-choice test.

3   LSMDC16 Multiple-Choice Test

In this section we explain our multiple-choice test collection for LSMDC16.

LSMDC16 multiple-choice test definition: Given a video with 5 possible description choices, the task is assigning correct(ground-truth) caption to the video among 5 possible choices. Figure 1 illustrate multiple-choice test for a video sample.

In order to define 5 possible captions for each video, first we tag each caption in the whole LSMDC16 with one or multiple activity phrase labels (i.e. described in the following paragraph). Second the correct answer is the ground-truth caption

and four other disctractor answers are randomly picked from the same subset (i.e. either training, test, or validation captions) with the condition that their activity phrase labels have no word intersection with correct answer activity phrase labels.

Original Audio Description (AD)
The nerd enters the Metropolitan City Police Academy wearing a blue uniform and takes his place in a registration line.

Activity elements (Predicate + Object)
enter metropolitan city police academy
enter police academy
wear blue uniform
wear uniform
take place

Fig. 2. An Example of Activity phrase mining from movie AD.

Activity knowledge mining: The activity elements are extracted from AD captions using "Knowlywood", a model for human activities mining from movie narratives such as AD captions [22] and it is presented in "Predicate + Object" (PO) format. Figure 2 illustrates an example of activity phrase labels from a AD caption (for more details about activity mining please see [22]).

Multiple-choice test evaluation: For evaluation, we compute test accuracy for public-test captions with 10053 videos. Test accuracy is the percentage of correctly answered questions (i.e. 10053 questions).
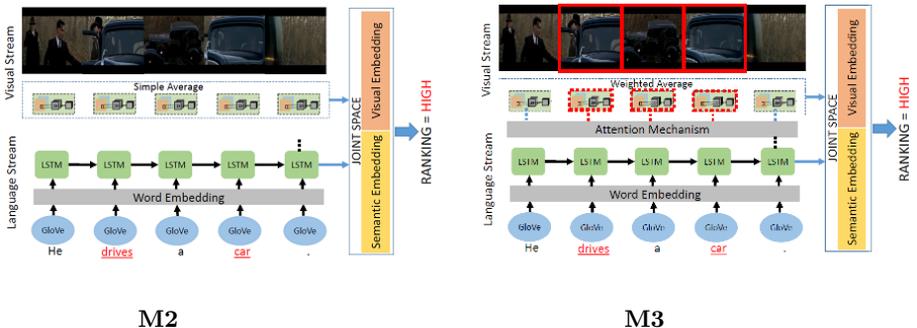

M2

M3

Fig. 3. The variations of neural networks models for learning joint visual-language embedding.

4   Joint visual-language neural network models

Our overall goal is to learn, given a set of training video-caption pairs $\mathcal{D}_{vid} = \{(\mathbf{V}_i, \mathbf{S}_i)\}$, a ranking function such that the ground-truth pairs $(\mathbf{V}_i, \mathbf{S}_i)$ rank

higher than all other pairs (e.g. $(\mathbf{V}_i, \mathbf{S}_j)$ where $j \neq i$) by a margin. While this high level goal is easy to specify it requires a number of algorithmic choices and challenges to realize. we implement three different models:

M1. Simple Average Glove + Simple Average FC-7 (SA-G + SA-FC7): We use this simple model as the baseline. In this model the sentence is encoded using simple average word vector representations. Also video is encoded using simple average of video frames feature vectors as we will describe in Appendix B.1. Finally the video representation and sentence representation are separately linearly transformed into joint visual-language space. We constrain video and sentence embedding vectors to only positive values by computing their absolute values as suggested in [24]. The model is trained using a ranking objective function that will be described in Appendix B.2.

M2. Language LSTM + Simple Average FC-7 (LSTM + SA-FC7): Figure 3 (left) illustrates this model. In this model the sentence is encoded using LSTM [7] and represented in joint semantic space as LSTM last hidden layer output; video is encoded using simple average of video frames feature vectors as we will describe in Appendix B.1. Finally the video representation is linearly transformed into joint semantic space. In this model, we also constrain video and sentence embedding vectors to only positive values by computing their absolute values. The model is trained using a ranking objective function that will be described in Appendix B.2.

M3. Language LSTM + Weighted Average FC-7 (LSTM + WA-FC7): Figure 3 (right) illustrates this model. In this model, similar to above, the sentence is encoded using LSTM and its embedding is absolute values of last hidden state output of LSTM. A soft-attention aligns the output of last hidden state of language LSTM, $\mathbf{h}_N$, with all feature vectors of frames in $\mathbf{V}$. Video is then encoded using weighted average of frame feature vectors and the video representation is linearly transformed into joint embedding space obtained by minimizing a ranking loss that we will describe in Appnedix B.2.

5   Experiments

Datasets: We use two datasets in our experiments:

1. COCO image description dataset [13] with 123,287 images: 113,287 training images, 5,000 test and 5,000 validation images. Each image has 5 ground-truth captions.
2. LSMDC16 [18] movie description dataset with 101,079 training, 10,053 test and 7,408 validation videos. Each video has only one ground-truth caption.
3. LSMDC16 [18] rephrases, we have collected 17,000 captions generated by human using Amazon's Mechanical Turk (AMT) for LSMDC16 datasets. 11,000 captions are from training set and 6,000 captions are from test set. We asked people to re-phrase original captions with 3-10 word sentences using different wording compared to original captions as much as possible. As an example, "She walks over to the banquette and sits down." is original caption was re-phrased as "She walked to the feast and took a seat.".

5.1   Implementation details

We implemented our models in Theano-based framework named BLOCKS [2,15].
Our data pipeline has been implemented using FUEL framework [15].

Model configurations: We set LSTMs hidden-state dimension and the dimensions of the joint embedding space to 950 (we also tried 512 but it resulted in lower performance). The dimension of the word embedding is 300. The attention model matching space is also 300 (we also tried 900 which resulted in similar performance), and set the ranking margin $\alpha = 0.05$ (as proposed in [24]). Similar to other works, e.g. [11,24], to decrease overfitting issue, we constrain sentence and video embeddings to have a unit (L2) norm.

Training: We found it useful to train with both video-sentence (LSMDC) and image-sentence (COCO) datasets; to accommodate this for standard ranking task we treat images as one-frame videos. For both standard ranking and multi-choice test training phase is the same. However for multi-choice test we tried both pairwise ranking and annotation ranking as we described in Appendix B.2. In training phase, we randomly sample minibatch of size 200 image/video-caption pairs. The contrastive terms for the minibatch are computed by using 199 contrastive videos for each caption and 199 contrastive captions for each video, similar to [24]. Our models were trained with Adam optimizer with learning rate 0.001 and we applied clipping gradient with threshold 2.0. We early-stop and save the best model based on monitored loss function of 1000 randomly selected validation samples. For all datasets, including LSMDC16 and COCO, we use original train, valid, and test splits.

Testing: For evaluation standard video annotation and retrieval, we use Recall@K and Median Rank (medR) metrics. "r@k" means the percentage of ground-truth captions/videos in the first K retrieved captions/videos and "medR" means the median ranks of ground-truth videos/captions. We report results on COCO test images, LSMDC16 test set, and LSMDC16 test rephrasing captions that we will make public. For multi-choice test evaluation, we compute the accuracy over test-data. In order to do that for each test query, first we compute order similarity (Equation 3 in Appendix B.2) for all 5 answer choices and then video is assigned to description with highest score, then accuracy is computed for all test data as we described in Section 3.

5.2   Qualitative Results for temporal attention network

Figure 4 shows results of top 5 (at each row, from left to right) phrase-based video search. A phrase query have been shown on top of each row. Each row shows the video frame with highest attention weight from the retrieved video computed for corresponding phrase query. In order to compute attention weights for a new phrase query, we forward pass top 5 videos, retrieved by model **M1** or **M2**, and new phrase query through attention network in **M3** model, trained on combination of COCO and LSMDC. We observe that our temporal attention network is able to highlight salient video frames in LSMDC video clips by associating a higher attention weights to salient video frames.
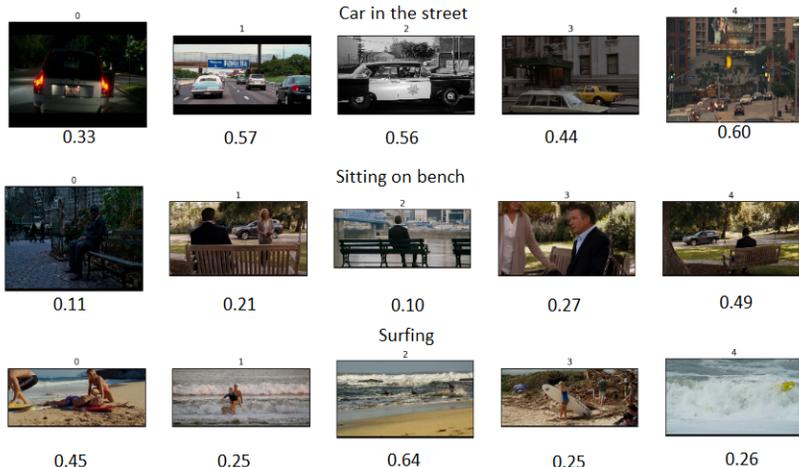
Fig. 4. Top 5 (from left to right) phrase-based video search results. Each row shows video frmaes with maximum attention weight computed with model **M3** (*C+L'16*). The phrase query is on top of each row. The value below each image is the attention weight for corresponding query.

Table 1. Quantitative comparison of models for standard ranking. **M1** is *SA-G + SA-FC7* model and **M2** is *LSTM + SA-FC7* (details are in Section 4), and **OE** stands for order-embeddings model [24]. The **C** stands for *COCO* image dataset, **L'16** for *LSMDC'16* datasets, and **RP** for *LSMDC'16* rephrases data. The **1000 AD** stands for 1000 randomly picked audio description captions from LSMDC16 public test-set; **1000 re-phrase** stands for 1000 captions re-phrasings from LSMDC16 public test-set.

Model	Annotation (1000 coco test)				Retrieval (1000 coco test)			
	R@1	R@5	R@10	medR	R@1	R@5	R@10	medR
M2 (C)	44.6	78.4	**89.5**	2	37.2	72.3	**85.9**	2
OE(C)[24]	**46.7**	-	88.9	2	**37.9**	-	**85.9**	2

Model	Annotation (1000 AD)				Retrieval (1000 AD)			
M1 (C+L'16)	2.4	9.0	14.0	113	3.0	8.8	13.2	114
M2 (L'16)	3.5	10.8	17.1	98	3.3	10.2	15.6	**88**
M2 (C+L'16)	**4.0**	12.2	**19.2**	91	**4.3**	12.6	18.9	98
M2(C+L'16 + RP)	3.9	**12.6**	18.1	**90**	4.2	**13.0**	**19.5**	90

Model	Annotation (1000 re-phrase)				Retrieval (1000 re-phrase)			
M1 (C+L'16)	1.8	6.7	11.0	149	2.0	6.8	10.8	155
M2(L'16)	2.4	6.9	11.8	139	2.0	6.9	11.8	131
M2 (C+L'16)	**2.6**	**8.5**	**12.2**	137	**2.5**	7.9	**12.7**	138
M2(C+L'16+RP)	2.2	7.0	12.1	**124**	1.8	**8.1**	12.5	**124**

Table 2. Quantitative comparison of models for multiple-choice test. **M1** is *SA-G + SA-FC7* model and **M2** is *LSTM + SA-FC7* (details are in Section 4). The **AR** stands for annotation ranking (details in Appendix B.2). The **C** stands for *COCO* image dataset, **L'16** for *LSMDC'16* datasets, and **RP** for *LSMDC'16* rephrases data.

Model	multiple-choice (10053 LSMDC16 test)
	accuracy
M1 (C+L'16)	55.1%
M2 (L'16)	56.3%
M2 (C+L'16)	56.6%
M2 (C+L'16+AR)	**58.1%**

5.3   Quantitative Results

Standard ranking: Table 1 summarizes all our evaluations for **M1** and **M2** model architectures described in Section 4. The best results are shown in bold. Following are some observations. First, using model **M2**, we are able to reproduced COCO image retrieval SOTA results in [24] (see row 1 and 2). Second, **M2** (*C+L'16*) trained on combination of COCO (C) and LSMDC16 (L'16), has consistently a better performance, for both 1000 original captions and re-phrases, compared to the same model only trained on LSMDC16 (L'16) or baseline model **M1** (*C+L'16*). Finally model **M2** (*C+L'16+RP*) trained on combination of COCO (C), LSMDC16 (L'16), and LSDMC16 re-phrases has better medR performance compared to **M2** (*C+L'16*).

Even the Standard ranking performance for original AD and rephrased captions is relatively low due to the caption complexity, still with simple phrase queries, we can obtain quiet impressive video retrieval results using trained joint language-visual models for LSMDC16 (see more qualitative results in Appendix A).
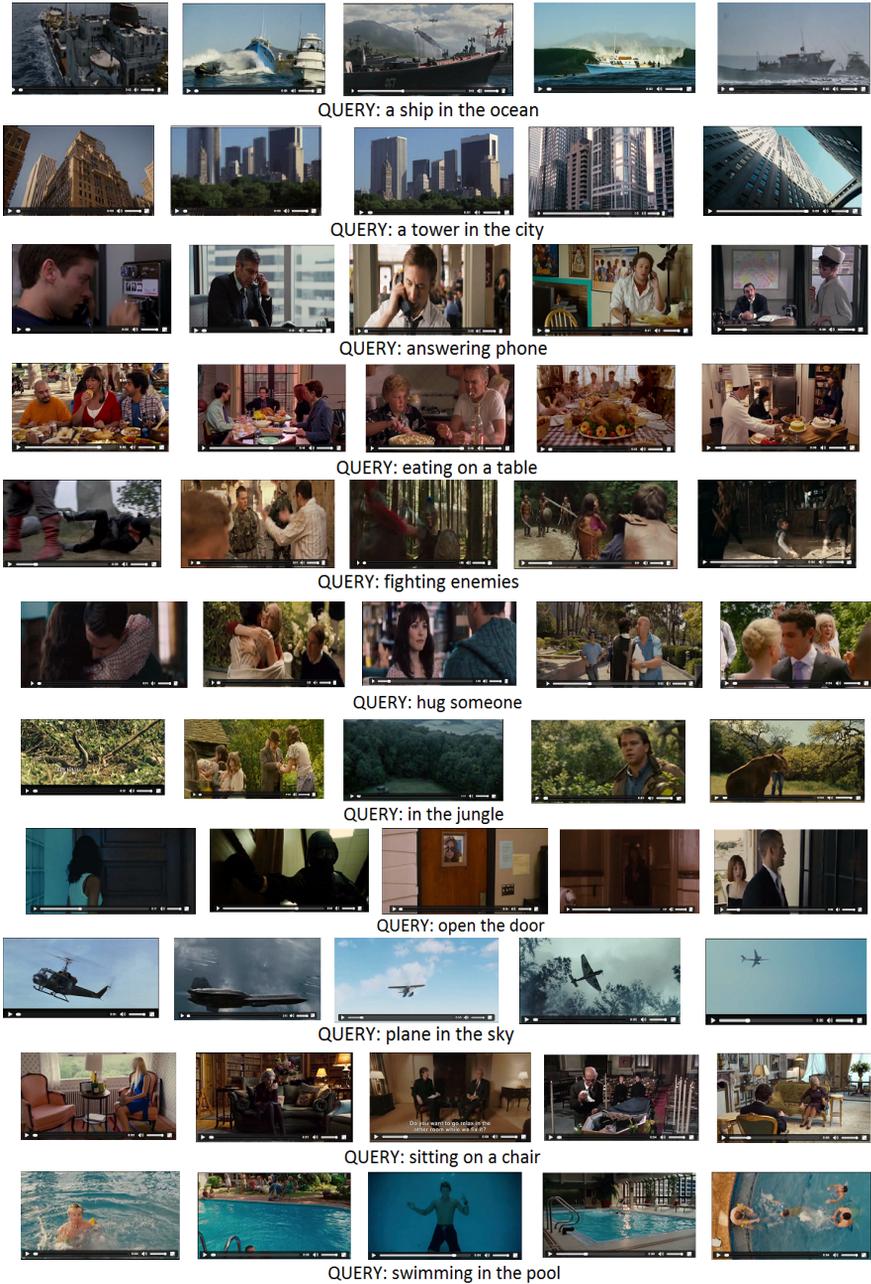
Multiple-choice test: Table 2 summarizes all our evaluations for **M1** and **M2** model architectures described in Section 4. The best result are shown in bold. First, **M2** (*C+L'16+AR*) which is trained on combination of COCO and LSMDC has the best performance. Second, training a model using **AR** (Annotation Ranking) loss has 2% improvement in accuracy compared to pairwise ranking for multiple-choice test (see ranking loss details in Appendix B.2).

References

1. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: Vqa: Visual question answering. In: ICCV (2015)
2. Bastien, F., Lamblin, P., Pascanu, R., Bergstra, J., Goodfellow, I.J., Bergeron, A., Bouchard, N., Bengio, Y.: Theano: new features and speed improvements. In: NIPS (2012)
3. Chen, X., Zitnick, C.L.: Mind's eye: A recurrent visual representation for image caption generation. In: CVPR (2015)
4. Donahue, J., Hendricks, L.A., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T.: Long-term recurrent convolutional networks for visual recognition and description. In: CVPR (2015)
5. Fang, H., Gupta, S., Iandola, F.N., Srivastava, R., Deng, L., Dollár, P., Gao, J., He, X., Mitchell, M., Platt, J.C., Zitnick, C.L., Zweig, G.: From captions to visual concepts and back. In: CVPR (2015)
6. Gao, H., Mao, J., Zhou, J., Huang, Z., Wang, L., Xu, W.: Are you talking to a machine? dataset and methods for multilingual image question answering. In: NIPS (2015)
7. Graves, A.: Generating sequences with recurrent neural networks. CoRR abs/1308.0850 (2013)
8. Hodosh, M., Young, P., Hockenmaier, J.: Framing image description as a ranking task: Data, models, and evaluation metrics. In: JAIR (2013)
9. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. arXiv preprint arXiv:1408.5093 (2014)
10. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for gen. image descriptions. In: CVPR (2015)
11. Kiros, R., Salakhutdinov, R., Zemel, R.S.: Unifying visual-semantic embeddings with multimodal neural language models. TACL (2015)
12. Lin, D., Fidler, S., Kong, C., Urtasun, R.: Visual semantic search: Retrieving videos via complex textual queries. In: CVPR. pp. 2657–2664 (2014)
13. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
14. Malinowski, M., Rohrbach, M., Fritz, M.: Ask your neurons: A neural-based approach to answering questions about images. In: ICCV (2015)
15. van Merriënboer, B., Bahdanau, D., Dumoulin, V., Serdyuk, D., Warde-Farley, D., Chorowski, J., Bengio, Y.: Blocks and fuel: Frameworks for deep learning. CoRR abs/1506.00619 (2015)
16. Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: EMNLP (2014)
17. Rohrbach, A., Torabi, A., Rohrbach, M., Larochelle, H., Pal, C., Courville, A., Schiele, B.: Describing and understanding video and the large scale movie description challenge. In: ICCV (2015), https://sites.google.com/site/describingmovies
18. Rohrbach, A., Torabi, a., Rohrbach, M., Tandon, N., Pal, C., Larochelle, H., Courville, A., Schiele, B.: Movie description. arXiv:1605.03705 (2016)
19. Sadeghi, F., Divvala, S.K., Farhadi, A.: Viske: Visual knowledge extraction and question answering by visual verification of relation phrases. In: CVPR (2015)
20. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)

21. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS (2014)
22. Tandon, N., de Melo, G., De, A., Weikum, G.: Knowlywood: Mining activity knowledge from hollywood narratives. In: Proc. CIKM (2015)
23. Tapaswi, M., Zhu, Y., Stiefelhagen, R., Torralba, A., Urtasun, R., Fidler, S.: Movieqa: Understanding stories in movies through question-answering. In: CVPR (2016)
24. Vendrov, I., Kiros, R., Fidler, S., Urtasun, R.: Order-embeddings of images and language. arXiv preprint arXiv:1511.06361v6 (2015)
25. Venugopalan, S., Rohrbach, M., Donahue, J., Mooney, R.J., Darrell, T., Saenko, K.: Sequence to sequence - video to text. In: ICCV (2015)
26. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. In: CVPR (2015)
27. Wu, Q., Wang, P., Shen, C., Dick, A., van den Hengel, A.: Ask me anything: Free-form visual question answering based on knowledge from external sources. In: CVPR (2016)
28. Wu, Z., Wang, X., Jiang, Y., Ye, H., Xue, X.: Modeling spatial-temporal clues in a hybrid deep learning framework for video classification. In: ACM Multimedia (2015)
29. Xu, J., Mei, T., Yao, T., Rui, Y.: Msr-vtt: A large video description dataset for bridging video and language. CVPR (2016)
30. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A.C., Salakhutdinov, R., Zemel, R.S., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: ICML (2015)
31. Yao, L., Torabi, A., Cho, K., Ballas, N., Pal, C., Larochelle, H., Courville, A.: Describing videos by exploiting temporal structure. In: ICCV (2015)
32. Zhu, L., Xu, Z., Yang, Y., Hauptmann, A.G.: Uncovering temporal context for video question and answering. arXiv preprint arXiv:1511.04670v1 (2015)

A   Phrase-based movie shots search

QUERY: a ship in the ocean

QUERY: a tower in the city

QUERY: answering phone

QUERY: eating on a table

QUERY: fighting enemies

QUERY: hug someone

QUERY: in the jungle

QUERY: open the door

QUERY: plane in the sky

QUERY: sitting on a chair

QUERY: swimming in the pool

Fig. 5. Top 5 (from left to right) phrase-based video search results with natural language using model **M2** (C+L'16).

B   Model details

Detail about sentence and video encoding and ranking loss that we used in our models.

B.1   Video and sentence encoding

Each sentence is a sequence of words. The length of word sequences are variable from one sentence to the next. Each word is initially encoded using GloVe [16] distributed word vector representation that has been pre-trained on a 6 billion word corpus that includes Wikipedia and Gigaword 5th edition datasets. Each sentence is therefore represented by a matrix:

$$\mathbf{S} = [\mathbf{w}_1, ..., \mathbf{w}_N] \in \mathbb{R}^{N \times d_{\mathbf{w}}} \tag{1}$$

of N word feature vectors, and each vector has $d_w = 300$ dimension in our GloVe representation. In model $M1$, fixed length representation of the sentence can then be obtained as simple average of the word representations, $\frac{1}{N}\sum_{i=1}^{N}\mathbf{w}_i$. The sentence representation is then fed into a linear transformation and the absolute values of transformation output is taken as semantic representation of the full sentence. In model $M2$ and $M3$, the sentence representation $\mathbf{S}$ is fed into an LSTM which results in a sequence of hidden states $[\mathbf{h}_1, ..., \mathbf{h}_N] \in \mathbb{R}^{N \times d_h}$, where d_h is the dimension of the LSTM hidden state. The absolute values of the output of the last hidden state $\mathbf{h}_N$ is taken as semantic representation of the full sentence.

A video is represented as sequence of frames which are sampled every 10 video frames (images are treated as single frame videos). We extract frame features using pre-trained VGG-19 [20] convolutional neural network (CNN) in Caffe [9]. VGG-19 network [20] is trained for object classification on ImageNet, to classify among $1,000$ object categories. Video sequence length varies from one sequence to the next, which results in a matrix representation:

$$\mathbf{V} = [\mathbf{v}_1, ..., \mathbf{v}_M] \in \mathbb{R}^{M \times d_{\mathbf{v}}} \tag{2}$$

of M video frame feature vectors, each of dimensionality $d_{\mathbf{v}}$. VGG-19, proposed in [20], has 16 convolutinal layers and 3 fully connected layers, followed by a soft-max output layer. For our experiments we extract second Fully-Connected layer (FC7), therefore $d_{\mathbf{v}} = 4096$. Fixed length representation of the video can then be obtained as simple average of the frame-based representations, $\frac{1}{M}\sum_{i=1}^{M}\mathbf{v}_i$, a weighted average obtained using soft attention or using an LSTM encoding.

B.2   Learning to rank

Two-level partial order similarity: We use the negative order-violation penalty, proposed in [24], for similarity metric, which is defined as

$$S(c, v) = -\|max\,(0, c - v)\|^2 \tag{3}$$

using order violating penatly, where c is either an embedding (for absolute values of simple average: $|\mathbf{W_{word}} \cdot \frac{1}{M} \sum_{i=1}^{M} \mathbf{v}_i|$, with embedding matrix $\mathbf{W_{word}}$)) or $c = |\mathbf{h}_N|$ denotes the absolute values of last hidden state of the language LSTM and v is an embedding (for absolute values of simple/weighted average: $|\mathbf{W_{video}} \cdot \frac{1}{M} \sum_{i=1}^{M} \mathbf{v}_i|$, with embedding matrix $\mathbf{W_{video}}$). $\mathbf{W_{word}}$, $\mathbf{W_{video}}$, and $\mathbf{h}_N$ are netowrk parameters that are learned during training phase.

The advantage of this similarity distance is that it is asymmetric, compared to cosine similarity which is symmetric, so it can capture the relatedness of captions with very different lengths that describe the same visual content but in different levels of detail (detailed discussion is given in [24]).

Pairwise ranking loss: As baseline, we use the standard loss that has been used in multiple works for image retrieval task in [11,24] and is defined as

$$\sum_{(c,v)} \left(\sum_{c'} max\left\{0, \alpha - S(c,v) + S(c',v)\right\} + \sum_{v'} max\left\{0, \alpha - S(c,v) + S(c,v')\right\} \right)$$
(4)

where (c,v) is an encoding for the ground-truth video/sentence pair, c' (contrastive captions) are captions that do not belong to v and v' (contrastive videos) are video/images that are not captioned by c; α denotes a margin hyperparameter and S is the similarity function in Eq. 3.

Annotation ranking loss: The first term in the above pairwise ranking 4, is a hinge loss that promote a pair of video and ground-truth caption is required to score higher than all other pairings of same video with contrastive captions by a margin of α. We name this annotation ranking loss and we show for multi-choice test, training models with this loss has better performance compared to pairwise ranking loss. Annotation ranking loss is defined as

$$\sum_{(c,v)} \left(\sum_{c'} max\left\{0, \alpha - S(c,v) + S(c',v)\right\} \right)$$
(5)