

Knowledge Graph Extraction from Videos

Louis Mahon, Eleonora Giunchiglia, Bowen Li, Thomas Lukasiewicz

Department of Computer Science

University of Oxford, UK

email: louis.mahon@cs.ox.ac.uk

Abstract—Nearly all existing techniques for automated video annotation (or captioning) describe videos using natural language sentences. However, this has several shortcomings: (i) it is very hard to then further use the generated natural language annotations in automated data processing, (ii) generating natural language annotations requires solving the hard subtask of generating semantically precise and syntactically correct natural language sentences, which is actually unrelated to the task of video annotation, (iii) it is difficult to quantitatively measure performance, as standard metrics (e.g., accuracy and F1-score) are inapplicable, and (iv) annotations are language-specific. In this paper, we propose the new task of knowledge graph extraction from videos, i.e., producing a description in the form of a knowledge graph of the contents of a given video. Since no datasets exist for this task, we also include a method to automatically generate them, starting from datasets where videos are annotated with natural language. We then describe an initial deep-learning model for knowledge graph extraction from videos, and report results on MSVD* and MSR-VTT*, two datasets obtained from MSVD and MSR-VTT using our method.

I. INTRODUCTION

Recent progress in deep learning shows exciting potential for visual understanding. Promising results have been produced in the tasks of video annotation (or captioning) using natural language (see, e.g., [1]–[3]) and video question answering (see, e.g., [4]–[6]).

However, annotations written in natural language (NL) are not ideal, for several reasons. Firstly, it is difficult to use NL sentences in subsequent automated data processing tasks. One example of such data processing is database search. If we have a database of videos annotated with NL and want to search for all videos depicting a man, we cannot simply perform string matching on the annotations. This would, e.g., produce a false positive for the annotation “sailors man the ship”, and a false negative for “a firefighter puts on his coat”. Another example of automated data processing is logical inference. For example, if an NL annotation model correctly identifies that there is a man throwing a ball in an image, it cannot then conclude that there is a male throwing a ball in the image. That is, a logical inference rule of the form $\forall x \text{ man}(x) \Rightarrow \text{male}(x)$ cannot easily be applied. Automated inference in NL is known to be a difficult problem [8], [9]. Again, string matching would not provide a solution, as the above examples attest. A second drawback of NL annotations is that they require more work than just understanding the contents to be annotated. Traditionally, text generation has been divided into a semantic stage of determining what to say, and then a realization stage of determining how to say

it [10], [11]. The former corresponds to understanding the contents being annotated and is the task that we are interested in, but deep learning NL annotation models must, at the same time, learn to perform the latter, which means learning the language’s complex meaning and grammatical structure. This is irrelevant to the goal of annotation and so needlessly makes the task more difficult. Thirdly, it is not easy to interpret and quantitatively measure the performance of NL annotations. The agreement between a ground-truth and the predicted sentence cannot be measured by standard metrics, such as accuracy and F1-score. Instead, they require specially designed ones (e.g., BLEU [12], METEOR [13], and LEPOR [14]). Recognizing the imperfection of each of these, results typically report scores on multiple metrics, none of which have a simple and intuitive interpretation. This makes it difficult to evaluate how well a model is performing. A fourth problem is that NL annotations are specific to a single language. A model trained to produce annotations in English cannot produce annotations in German, and it is non-trivial to translate the English annotations into German ones. This is a particular drawback for low-resource languages.

Replacing NL annotations with knowledge graphs avoids all the above problems. A knowledge graph (KG) specifies the *individuals* (i.e., the objects) present in the video and the facts that hold true of these individuals. A *fact* expresses a relation between two individuals (e.g., *fold(person, paper)*) or an attribute of an individual (e.g., *white(paper)*). A visual representation of a KG is given in Fig. 1, where nodes representing individuals are depicted in red, nodes representing attributes are depicted in blue, and nodes representing relations are depicted in green. A KG can be represented visually or as a set of facts; see Fig. 1 (7) and (5), respectively. KGs are machine-readable, so enable automated data processing, such as searching for all videos depicting a man or applying the inference rule $\forall x \text{ man}(x) \Rightarrow \text{male}(x)$. They are determined only by the semantics of the video contents, and do not involve a semantically and syntactically complex natural language structure. They can be evaluated using standard machine-learning metrics, such as accuracy and F1-score, which makes it easy to interpret performance. For each data point, the ground truth and the prediction are sets of facts; so, e.g., the F1-score can be computed by counting the facts that appear in both as the true positives, those that appear in the former only as false negatives, etc. Finally, each fact, and hence each KG, can be trivially translated by translating one component at a time. For example, we can translate *white(paper)* to



(1) **MSR-VTT ground truth:**
a person folds a piece of white paper

(2) **Predicted caption from [7]:**
a person is folding a piece of paper

(3) **MSR-VTT* ground truth:**
fold(person, piece)

(4) **Predicted individuals:**
person, man, paper, origami

(5) **Predicted facts:**
demonstrate(person), fold(person, paper) white(paper), fold(man, origami)

(6) **Some inferred facts:**
change(person, paper), fold(male, origami)

(7) **Visual representation of the predicted knowledge graph:**




Fig. 1. The first frame from a video in MSR-VTT*, with (1) the ground-truth natural language annotation in MSR-VTT, (2) the natural language annotation produced by [7], (3) the ground-truth set of facts in MSR-VTT*, (4) the individuals predicted by our system, (5) the facts predicted by our system, and (6) some further inferences that can be made on top of these facts.

German by translating *white* to *weiss* and *paper* to *Papier*, giving *weiss(Papier)*. (As discussed in Sections II and V, the components of each fact denote particular senses of words, so word-sense disambiguation is not an issue.) For these reasons, we propose to produce annotations using KGs rather than NL.

These advantages have begun to be recognized in neighboring areas. The entire task of open information extraction is motivated by use cases that can only be met by a structured representation of information, and not by the same information expressed in natural language. In computer vision, researchers have started to argue for the importance of basic visual reasoning (see, e.g., [15]) and the ability to leverage external knowledge (see, e.g., [16]). A number of works aim to extract a structured description from an input image to reason about the identified objects (see, e.g., [17]). In the video understanding field, there have been only preliminary attempts [18], though, recently, [19] has sketched some planned future work to manually create a dataset of videos annotated with facts.

In this paper, we describe a method for automatically generating datasets of videos and corresponding KGs. The method uses a rule-based semantic parser, based on the Stanford parser [20]. The generated datasets are sufficient for the development of KG extraction models. Additionally, we introduce one such model, trained and tested on two generated datasets, which achieves a superior performance to [18], the only existing work to attempt the same task (F1-score of 14.0 vs. 6.1). We also compare to an artificial baseline, conduct ablation experiments that demonstrate the efficacy of each of the model’s components, and present qualitative results. This model enjoys all the advantages described above: it facilitates automated data processing, it does not require modelling complex NL syntax and semantics, its performance is easy to evaluate, and the resulting annotations are easy to translate.

Example 1: Fig. 1 shows an example where the input is a video of someone folding paper. Given this video, our model detects the contained individuals, namely, *person, man, paper*, and *origami*. Then, it predicts which facts are true of

these individuals: *demonstrate(person), fold(person, paper), white(paper)*, and *fold(man, origami)*. As an example of automated data processing, it may then exploit an ontology to predict more facts that are true for the video, e.g., the rule $\forall x, y (fold(x, y) \Rightarrow change(x, y))$ to infer the fact *change(person, paper)*. The corresponding KG is shown in (7). \triangleleft

The main contributions of this paper are briefly as follows.

- We propose the task of extracting KGs from videos. We provide arguments for the superiority of this approach to NL annotations, and support these arguments empirically.
- We describe and make available a method for automatically generating datasets of videos and corresponding KGs, and describe its application to generate two datasets. The corresponding code is available at https://github.com/LouisM/video_annotation.
- We introduce a model for KG extraction, trained and tested on the above two datasets. This outperforms the existing work to attempt the same task. We also compare to an artificial baseline and conduct ablation studies on the main components of the model.

The rest of this paper is organized as follows. In Section II, we present our method for automatically generating datasets. In Section III, we describe our KG extraction model, and then report experimental results in Section IV. Section V provides a further discussion, and Section VI summarizes our main contributions and gives an outlook on future work.

II. DATASET GENERATION

In the absence of appropriate datasets, we devise an automated method to generate them. There are several NL-annotated datasets, which we refer to as video-captioning datasets. Our method begins with these datasets and converts the NL annotations into sets of facts (which are equivalent to KGs). We now describe this method with reference to its application to two well-known video-captioning datasets: MSVD [21] and MSR-VTT [22]. We denote the generated datasets MSVD*

and MSR-VTT*, respectively, and Section IV reports the results of our proposed model on these datasets.

In both MSVD and MSR-VTT, the training examples are composed of a video and a NL sentence describing the contents of that video. We parse these sentences using a rule-based parser based on the Stanford NLP syntactic parser [20]. This produces a dependency parse of the sentence, which identifies the part of speech for each word, and the syntactic relations that hold between them. We then apply a sequence of rules to extract the semantic information from the sentence, and use it to form a KG, which is composed of a set of facts. Each fact contains a predicate and the corresponding arguments: $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ triples in the case of binary predicates, and $\langle \text{subject}, \text{predicate} \rangle$ pairs in the case of unary predicates (see Figs. 1 and 3). A full description of this parser can be found in the appendix.

The next step is to link all predicates, and individuals to entities in an ontology. The ontology that we use here is WordNet [23]. The linking method is context-sensitive, i.e., for a given word, it finds the most suitable WordNet synset using the surrounding words in the NL annotation. This is done by comparing the word vectors for the surrounding words to the word vectors for the synset definition. Again, full details are given in the appendix. Linking to an ontology formalizes the vocabulary and allows the application of inference rules to augment the information produced directly by our annotation system. It also means that the components of our KGs correspond to particular senses of words. For example, “bank” referring to the financial institution, and “bank” referring to the side of a river would be linked to different WordNet synsets, and so would appear as different items in a KG.

Words that appear only a small number of times across the dataset would be difficult to reach a good performance on. For this reason, we exclude all words that appear fewer than 50 times and also some common semantically weak verbs, as detailed in the appendix.

The method described so far produces only non-negated facts, from here on referred to as T . This means a model could learn to simply predict every potential fact to be true. To prevent this, we use the local closed-world assumption [24] to create a set of negated facts, from here on referred to as F , which the model must also learn to predict. For each fact in the description, we obtain a corrupted version by replacing the predicate with a different predicate from the vocabulary. For example, the fact $fold(\text{person}, \text{paper})$, which appears in Fig. 1, could be corrupted to $\neg throw(\text{person}, \text{paper})$. This negated fact is then added to F .

As a final step, we merge all KGs for each video by taking the union of the facts of their parses. In MSVD and MSR-VTT, each video appears with multiple captions, and each produces a separate training example. In MSVD* and MSR-VTT*, each video appears in only one training example. After exclusion and merging, the numbers of distinct individuals, distinct predicates, and data points with at least one fact or negated fact are as reported in Table I.

III. OUR MODEL

In this section, we describe a deep learning model (illustrated in Fig. 2) for extracting KGs from videos. We first describe a general method for extracting a KG from an arbitrary input, and then detail how this method applies to the case where the inputs are videos.

A. Knowledge Graph Extraction Model

Given a set of possible inputs X (videos in our case) and a knowledge base KB (of ontological domain knowledge about X), let our vocabulary consist of the set P of all predicates that appear in KB and the set A of all individuals that appear in KB . Then, the neural architecture for KG extraction consists of the following four components:

- 1) an encoder $f: X \rightarrow Z$,
- 2) a multi-classifier $g: Z \rightarrow (0, 1)^{|A|}$,
- 3) a set of multilayer perceptrons (MLPs): $\{m_p \mid p \in P\}$,
- 4) a set of trainable individual vectors: $\{v_a \mid a \in A\}$,

where Z is the space of extracted feature vectors.

The data for training the above neural architecture consists of 4-tuples (x, c, T, F) , where $x \in X$ is the input to be annotated (a 4D video tensor in our case), $c \subseteq A$ is the set of individuals that are present in the input, and T (resp., F) is a set of facts (resp., negated facts) containing these individuals.

Example 2: Consider again the video in Fig. 1. The training tuple associated with the video is

$$(x, \{person, piece\}, \{fold(person, piece)\}, F),$$

where x represents the video itself, and F is a set of negated facts not containing $\neg fold(person, piece)$. An example of a negated fact that belongs to F is $\neg throw(person, piece)$. \triangleleft

During training, we compute $f(x) = e \in Z$ and feed this feature vector to the multi-classifier to predict the individuals present, $\hat{c} = g(e)$. Then, we compute a binary cross-entropy loss (BCE) for each individual in the vocabulary, giving

$$\mathcal{L}_c = \sum_{j=0}^{|A|} \mathcal{L}_{bin}(\hat{c}_j, c_j),$$

where $c_j \in \{0, 1\}$ and $\hat{c}_j \in (0, 1)$ are the ground truth and prediction as to whether the j th individual in the vocabulary is present, respectively, and \mathcal{L}_{bin} represents BCE. We also produce predictions for the facts and negated facts in T and F by selecting the individual vectors corresponding to the individuals in the ground truth, and feeding these vectors (along with the video encoding) to the predicate MLPs corresponding to the predicates in the ground truth. That is, for each fact $t \in T$ with $t = p(a, b)$, we compute a prediction $\hat{t} = m_p(e, v_a, v_b)$, and similarly for negated facts. A naive brute-force method would be to run every predicate MLP on every individual vector. This would require $|P| \times |A|$ forward passes for every video. Selecting only the present individuals reduces this to $|P| \times n$, where n is the number of individuals in the input video, which is typically less than three.

The vectors v_a and v_b are initialized to the corresponding word2vec word vectors [25]. For example, v_{man} is initialized

TABLE I
NUMBERS OF DISTINCT PREDICATES AND INDIVIDUALS THAT WERE INCLUDED IN OUR FINAL DATASETS. PREDICATES AND INDIVIDUALS WERE INCLUDED IF THEY APPEARED AT LEAST 50 TIMES. THE FINAL COLUMN SHOWS THE NUMBER OF DATA POINTS THAT REMAINED WITH NON-EMPTY ANNOTATIONS EVEN AFTER INFREQUENTLY OCCURRING WORDS WERE EXCLUDED.

	Num Training Examples	Num Individuals	Num Attributes	Num Relations	Num Facts	Num Non-empty Training Examples
MSVD*	1970	122	48	69	117	1800
MSR-VTT*	10000	372	235	113	348	9802

to the word2vec vector for the word “man”. These vectors are updated during training along with the network weights.

The training data contain many more negated facts than facts. To counteract this imbalance, we weigh the loss for facts and negated facts inversely to how many there are. The prediction loss is again calculated using BCE, this time applied to each fact (resp., negated fact). Formally,

$$\mathcal{L}_p = -\frac{1}{2|T|} \sum_{t \in T} \log(\hat{t}) - \frac{1}{2|F|} \sum_{f \in F} \log(1 - \hat{f}).$$

The total backpropagated loss is then a simple summation:

$$\mathcal{L} = \mathcal{L}_c + \mathcal{L}_p.$$

At inference time, given an input x , we first compute the feature vector (step (i)), then threshold the outputs of the multi-classifier to obtain the set I of predicted individuals (step (ii)):

$$\hat{I} = \{a \in A \mid g(f(x))_a > 0.5\}.$$

As we no longer have access to the ground truth set of individuals, we instead select the *predicted* individuals and pass them to the predicate-MLPs (step (iii)). This again avoids having to run all predicate-MLPs on all individuals. Specifically, to predict the unary relations that hold of these individuals, we run all predicate-MLPs in our vocabulary on all individual vectors corresponding to individuals in \hat{I} , and to predict the binary relations, we do the same on all pairs of individuals in \hat{I} . The total set \hat{T} of predicted facts is the union of the two:

$$\begin{aligned} U &= \{p(a) \mid p \in P, a \in \hat{I}, m_p(a) > 0.5\}, \\ B &= \{p(a, b) \mid p \in P, a, b \in \hat{I}, m_p(a, b) > 0.5\}, \\ \hat{T} &= U \cup B. \end{aligned}$$

Now, as an example of automated information processing, we can apply the additional knowledge from KB to the predicted facts \hat{T} to obtain an additional set of inferred facts T' :

$$\begin{aligned} U' &= \{p(a) \mid p \in P, a \in \hat{I}, KB \cup \hat{T} \models p(a)\}, \\ B' &= \{p(a, b) \mid p \in P, a, b \in \hat{I}, KB \cup \hat{T} \models p(a, b)\}, \\ T' &= U' \cup B'. \end{aligned}$$

The annotation is then updated with these inferences (step (iv)):

$$\hat{T} \leftarrow \hat{T} \cup T'.$$

Example 3: Consider again the video in Fig. 1, and the schema of the model in Fig. 2. Following the steps in Fig. 2, the KG for the video in Fig. 1 is produced in four steps.

- (i) The input video x is mapped into a fixed-length feature vector e by the encoder f .

- (ii) The encoding e is fed into a multi-classifier MLP g , which returns the probability that each individual in the vocabulary is present in x . The probabilities are then thresholded at 0.5, and we obtain the set

$$I = \{person, man, paper, origami\}.$$

- (iii) Consider the individual *paper*. The encoding e , together with the vector v_{paper} is passed through all the MLPs, and a fact is returned only when the prediction of the MLP is greater than 0.5, such as *white(paper)*. This is then repeated for each predicted individual, and each ordered pair of predicted individuals, producing

$$\hat{T} = \{demonstrating(man), white(paper), fold(person, paper), fold(man, origami)\}.$$

- (iv) All facts that can be inferred from KB and the predicted facts are used to augment the KG extracted from the video. Examples of such inferred facts are *change(person, paper)* and *fold(male, origami)*. \triangleleft

B. KG Extraction from Videos

The above method can be used to extract a KG from any unstructured data X . For example, X may contain sequences of word embeddings representing text, spectrograms representing audio, or image tensors representing still images. In each case, we only have to choose an encoder f , so that $f(x)$ is a fixed-length feature vector for each $x \in X$. We now describe how this method is applied when the input domain is videos.

Our encoder, f , is composed of a convolutional neural network, followed by a recurrent neural network. Given a video $x = x_1, x_2, \dots, x_n$, we first process each video frame x_i using a pre-trained frozen copy of VGG19 [26] and take as i th frame encoding, ζ_i , the output of the 3rd last layer. (This is a standard approach taken by others in the field; e.g., [27]–[29].) The sequence of frame encodings ζ_1, \dots, ζ_n is then passed through a gated recurrent unit (GRU) [30] to produce a sequence $\bar{\zeta}_1, \dots, \bar{\zeta}_n$. As a second stream, we also compute the feature vector from a frozen copy of the I3D network [31]. The final output of the encoder is the concatenation of this I3D feature vector and a weighted sum of the $\bar{\zeta}_i$ ’s, weighted by a learnable n -dimensional vector w . Using a vector notation, with the VGG network applied along the first dimension, our encoder f is defined by:

$$f(x) = w \cdot GRU(VGG(x)) \cdot I3D(x). \quad (1)$$

The MLPs for each predicate have one hidden layer, as does the multi-classifier. The input size of the multi-classifier is the

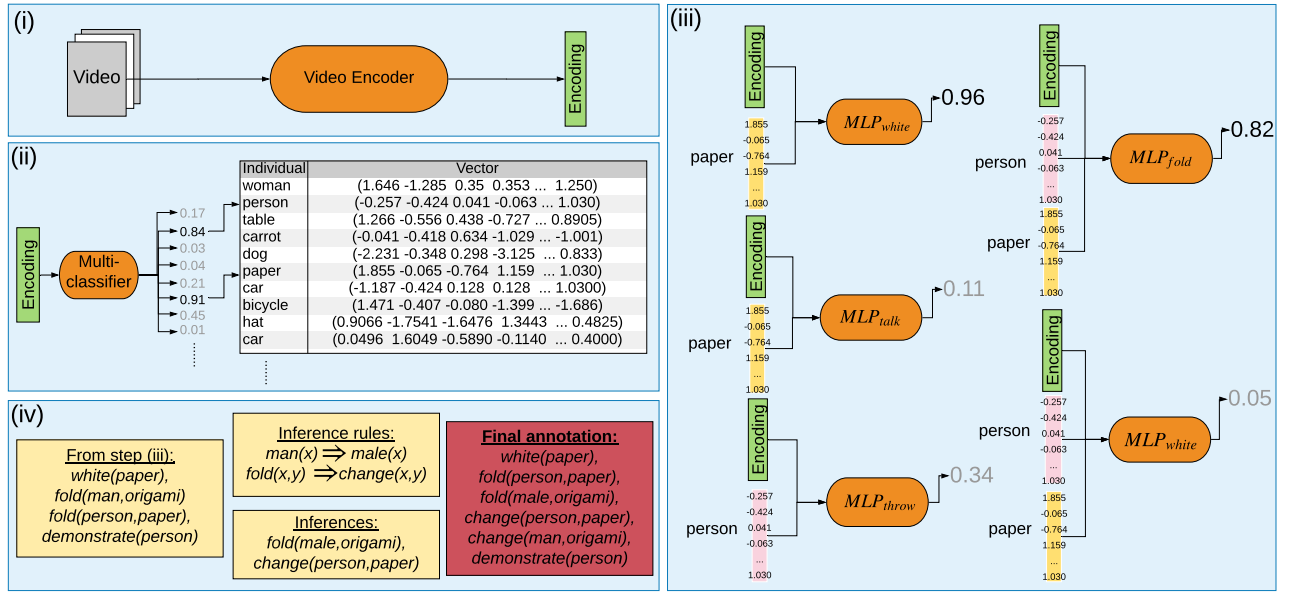


Fig. 2. Depiction of how the neural architecture processes a video of a man folding a piece of paper (as in Fig. 1). The four stages are: (i) encode the video, (ii) predict individuals, (iii) predict facts, and (iv) add facts inferred from *KB*.

size of the video encoding, denoted $\dim(f(x))$. The predicate-MLPs have input size $\dim(f(x)) + D$ in the case of unary relations and $\dim(f(x)) + 2D$ in the case of binary relations, where D is the size of the individuals' vectors (300 in our case). Using the encoder defined in (1), $\dim(f(x)) = 4096 + 1024 = 5120$. This framework could naturally be extended to predicates of arity greater than 2.

IV. EXPERIMENTAL RESULTS

In this section, we provide quantitative experimental results for the two datasets MSVD* and MSR-VTT* and compare these results. We also report on ablation studies for our neural architecture and give some further qualitative results.

A. Quantitative Results

Tables II and III display our performance on the two generated datasets MSVD* and MSR-VTT*, respectively, reporting both overall accuracy and F1-score. For each training example, these datasets contain many more negated facts than facts (see Section II for details). This is the reason for such a large difference between accuracy and F1-score: the predictions are dominated by true negatives, which increase the overall accuracy but have no effect on the F1-score. Accuracy and F1-score are harsh metrics, as they pertain to whole facts. An annotation would achieve 0 accuracy if it contained correct individuals and correct predicates, but related them in the wrong way, e.g., *fold(paper, person)* or *carry(person, paper)* in Fig. 1. All results are taken from a held-out test set, using the train/val/test splits defined in the original MSVD and MSR-VTT datasets.

There are limited existing baselines (only [18]) for the task of KG extraction from videos. To benchmark our performance, we also compare to an artificial baseline composed of the video-captioning network from [7], and the semantic parser

TABLE II
RESULTS ON THE MSVD* VIDEO AND DERIVED KGs, AS DESCRIBED IN SECTION II (BEST RESULTS IN BOLD).

	F1-score	Positive Accuracy	Negative Accuracy	Total Accuracy
Our Model	13.99	12.65	99.20	22.16
Baseline	13.5	7.53	99.96	25.55
V&L 2018	6.11	3.36	-	-

TABLE III
RESULTS ON THE MSR-VTT* VIDEO AND DERIVED KGs, AS DESCRIBED IN SECTION II (BEST RESULTS IN BOLD).

	F1-score	Positive Accuracy	Negative Accuracy	Total Accuracy
Our Model	8.90	7.33	99.48	59.19
Baseline	11.83	6.76	99.96	83.01

used to create our datasets (described briefly in Section II, and more in detail in the appendix).

Our model significantly outperforms the previous work by [18], and performs on par with the video-captioning baseline. Importantly, it shows the best positive accuracy, which is the most difficult metric to score highly on.

The metrics, both for the baseline and for our network, are lower on MSR-VTT*. This is consistent with other reported results in the literature. For example [7], [29], [32]–[34] all report lower video captioning performance on MSR-VTT than on MSVD. A likely reason is the greater vocabulary size of MSR-VTT: 29,316 vs. 13,010 for MSVD (figures taken from [22]). Although we exclude all but the most common individuals, a greater vocabulary size in the original NL captions would still make our task more difficult. It means that more facts will be excluded, and so there will be more annotations with missing information, which do not, therefore,

TABLE IV
ABLATION RESULTS FOR MSVD* (BEST RESULTS IN BOLD).

	F1	Positive Accuracy	Negative Accuracy	Total Accuracy
Our Model	13.99	12.65	99.20	22.16
Single MLP decoder	13.75	12.5	98.45	19.04
Single individual vector decoder	10.72	9.91	99.32	22.50
Without encoder	8.87	9.77	97.49	15.60

TABLE V
ABLATION RESULTS FOR MSR-VTT* (BEST RESULTS IN BOLD).

	F1	Positive Accuracy	Negative Accuracy	Total Accuracy
Our Model	8.90	7.33	99.48	59.19
Single MLP decoder	6.16	6.75	99.15	49.09
Single individual vector decoder	3.84	3.94	99.51	58.72
Without encoder	5.00	5.36	99.19	49.64

fully describe the input video.

B. Ablation Studies

To investigate the contribution of each part of our neural architecture, we perform ablation studies on the encoder f , the predicate MLPs, and the individual vector ($\{m_p \mid p \in P\}$ and $\{V_a \mid a \in A\}$ in the terminology of Section III).

The results for MSVD* and MSR-VTT* are shown in Tables IV and V, respectively. In the “without encoder” setting, the video encoding feature vector is replaced with a randomly generated vector. This tests whether the system is actually using information from the video or is just predicting common individuals and predicates, e.g., *talk(man)*. That the results in this setting are worse shows that information from the video is indeed being used. In the “single MLP decoder” setting, one MLP is used for all predicates, in contrast to the main system in which each predicate has its own MLP. This is to test whether predicate-specific information is being used, e.g., whether the system makes meaningfully different predictions for *throws(woman, ball)* and *sees(woman, ball)*. Similarly, in the single-individual decoder setting, one vector is used for all individuals, in contrast to the main system in which each individual has its own vector. This tests if individual-specific information has been learned by the individual vectors. Both decoder ablation settings produce worse results, which shows that the predicate-specific MLP and individual-specific vectors have learned some meaningful semantics of their respective predicates and individuals. Note that the results suffer more in these ablation settings on MSR-VTT* than on MSVD*. This is to be expected, as MSR-VTT* contains a larger vocabulary, which allows for more semantic differentiation between different predicates and between different individuals.

C. Qualitative Results

To further evaluate the quality of our KG extractions, we present some manual inspections of images and the predicted

facts. Fig. 3 shows the first frame from an MSVD* video with our predicted facts, and the same for a video from MSR-VTT*. (Fig. 1 shows another from MSR-VTT*.)

We can now infer additional facts from KB and the facts generated by the video annotation network, which is one of the advantages of our approach of using KGs. The class *man* is a subclass of *person* and *male* in WordNet, and we can apply this inference to all facts mentioning a man, producing *stand(male)* and *stand(person)* in Fig. 3. If we wanted to determine how many videos in a database depicted at least one person, or how many depict males, or how many depict males sitting vs. standing, this would not be possible by merely annotating the videos with NL sentences. It is, on the other hand, possible with our system, because of the inferences that can be made on top of a KG. Finally, the qualitative examples show the limitations imposed by the smaller vocabulary size in MSVD*. As discussed in Section II, we exclude individuals and relations that appear fewer than 50 times across the dataset. This means that there is sometimes insufficient material to fully describe a video. For example, the man in the video shown in Fig. 3 is eating a banana, and there was one MSVD caption that expressed this. However, *banana* appears fewer than 50 times, so it is excluded from our training data, and we cannot predict *eat(man, banana)*. (Recall from Section II that there are multiple annotations for each video, and we merge all of them in our dataset. The annotation shown in Fig. 3 is one that was not excluded.)

V. FURTHER DISCUSSION

In this section, we further discuss the advantages of generating KGs compared to NL annotations with reference to the output of our model. We also examine our model as compared to the video-captioning baseline.

A. Value of KGs

One advantage of KGs is that they enable the application of automated data processing. An example is the application of inference rules to infer additional facts, which is illustrated by the results shown in Fig. 3. The facts predicted by our model (3) allow the inference of further information (6), whereas the NL annotation produced by [7] does not, even when it is exactly correct, as in the case of the MSR-VTT video. As well as producing additional information about the video contents, these inferences are especially useful for producing abstract classes and relations, such as *male* in Fig. 3 or *change* Fig. 1. Other examples of automated data processing are query answering and search. If we are interested in videos that depict a male changing some object in some way, as happens in the video shown in Fig. 1, then using the extracted KGs (including further inferences), we can simply return the videos where, for some individual x , the fact *change(male, x)* has been predicted. In contrast, NL annotations do not enable us to perform such a search. We cannot simply replace all occurrences of *man*, *boy*, *guy* etc. with *male*, and all occurrences of *fold*, *open*, *close* etc. with *change* and then count the annotations that contain both *male* and *change*. This would

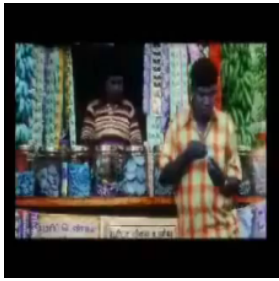

 <p>(1) MSVD ground truth: a man is standing at the bus stop</p> <p>(2) Predicted caption from [7]: a man is talking on the phone</p> <hr/> <p>(3) MSVD* ground truth: <i>stand(man)</i></p> <p>(4) Predicted individuals: <i>man</i></p> <p>(5) Predicted facts: <i>stand(man), sit(man)</i></p> <p>(6) Some inferred facts: <i>stand(person), stand(male)</i></p>	 <p>(1) MSR-VTT ground truth: a man is driving a car</p> <p>(2) Predicted caption from [7]: a man is driving a car</p> <hr/> <p>(3) MSR-VTT* ground truth: <i>drive(man, car)</i></p> <p>(4) Predicted individuals: <i>man, car, person</i></p> <p>(5) Predicted facts: <i>drive(man, car), drive(person, car)</i></p> <p>(6) Some inferred facts: <i>drive(man, vehicle)</i></p>
--	---

Fig. 3. Left: the first frame from a video in the MSVD* test set, with the ground-truth caption, the caption produced by our system, and some further inferences made using *KB* and these facts. Compared with a video captioning system, our output can be used to infer additional information about the video, such as the fact that it contains a male. Right: the second frame (first is very unclear) from a video in the MSR-VTT* test set, with the ground-truth caption, the caption produced by [7], the facts predicted by our system, and some further inferred facts. Although the video captioning model produces a sentence similar to the ground truth, it cannot be used for inference and so misses the additional inferred information.

produce false positives with annotations like “a man is handed a folded sheet of paper”. In such a sentence, although *male* and *change* would both be present, they are not connected in the way that we are interested in. The sentence does not describe a scenario in which a male is changing something, even though it contains the words *man* and *fold*. It is non-trivial to represent the information encoded in a NL sentence as a set of individuals and relations that hold between members of this set. To apply such a translation, we would need to employ a semantic parser, i.e., we would be using the method of the video-captioning baseline, and the demerits of this approach are discussed below. String matching in NL is not even able to detect the individuals that are present. The sentence “a person folds a boy’s school uniform” could describe a video in which no boy is present. Our annotations represent information in the semantically structured format of a KG, from which information such as the above can be easily read off. Finally, we can see that the set of facts (equivalent to a KG) shown in Fig. 3 could be easily represented in a language other than English. The individuals and predicates in the vocabulary correspond to particular senses of words, so disambiguating polysemous words is not required. To translate, e.g., the fact *drive(man, car)* to German, we could simply look up the translation for the senses of each of these three words: *drive* → *fahren*, *man* → *Mann*, and *car* → *Auto*. Translating one item at a time, we then get *fahren(Mann, Auto)*. In contrast, translating the NL sentence would require both the disambiguation of polysemous words, and the extraction of English syntactic features, and their expression in German.

B. Comparison with Video-Captioning Baseline

One possible method of producing KGs, and hence the advantages discussed above, is to first produce an NL annotation and then semantically parse this NL sentence to convert it to a set of facts. This is the method employed in the video-captioning baseline, to which we compare our performance in Section IV. The quantitative comparison, across the multiple metrics and datasets, shows that it performs about equally well to our model. However, there are three reasons to believe that

the latter is more promising for future progress. Firstly, the video-captioning baseline must solve two problems to produce an annotation: the problem of annotating videos with NL, and the problem of semantically parsing NL. This means it has two potential sources of error. However, errors in the second stage (that of semantic parsing) are not properly penalized in the results in Tables II and III, because the baseline uses the same semantic parser as was used to create the dataset. Therefore, if the video-captioning network correctly predicts the original NL caption, the result, after parsing, will automatically be correct according to our dataset. This will happen even if there was a mistake in the semantic parse. For example, in Fig. 1, the NL annotation predicted by [7] is “a person is folding a piece of paper”, which the parser then, incorrectly, parses as *fold(person, piece)* (instead of *fold(person, paper)*). However, this same incorrect parse was made when creating the dataset, so the ground truth in MSR-VTT* is also *fold(person, piece)*, and the baseline is marked correct in this prediction. If there existed a dataset of videos with human-made KGs, the baseline would still require a semantic parser, and its accuracy would be limited by the parsing accuracy, whereas our model would not.

A second advantage of our model is that we can restrict the vocabulary of our annotations when desired. We may, on occasion, only be interested in a certain set of individuals, and we could easily modify our network to only output annotations concerning these individuals. If, for example, we were only interested in determining whether a video contained any animals and what those animals were doing, we could remove all facts from our dataset other than those pertaining to animals and then train our model on the modified dataset. This would be a simpler task and be expected to give a higher accuracy. The video-captioning baseline, however, would still output an NL sentence as an intermediary step, and it would not be simple to apply the same restriction to this space of NL sentences. One may think that the vocabulary of the output language model in the video-captioning network could be restricted to only a certain, animal-related vocabulary. However, many

words that have one meaning denoting an animal have also other meanings, e.g., “fly” or “bear”. Again, this highlights the complexities of working with NL and so the undesirability of using it as an intermediary step to produce a KG.

Thirdly, the captioning and parsing approach is only possible where there already exist networks mapping the input domain to NL. This is true for videos, but need not be true in general. Our method, on the other hand, can be applied to any input domain, as long as a suitable encoder is used.

VI. CONCLUSION

This paper has proposed the task of KG extraction from videos, where a KG is composed of a set of facts that describes relations holding between individuals. We have provided arguments and empirical support for the advantages of KGs over NL annotations. No datasets currently exist of videos and corresponding KGs, so we have proposed a method to generate them from existing datasets of videos and NL annotations. Further, we have introduced a deep learning model for KG extraction, and evaluated its performance both qualitatively and quantitatively on two generated datasets.

Future work includes exploring KG extraction from other input domains. Of particular interest is the application to text, where our model would perform a task akin to open information extraction. Another extension is to manually construct a dataset designed specifically for the purpose of KG extraction, rather than using those generated by our automated method.

ACKNOWLEDGMENTS

This work was supported by the Alan Turing Institute under the UK EPSRC grant EP/N510129/1 and by the AXA Research Fund. Eleonora Giunchiglia is supported by an Oxford-DeepMind Graduate Scholarship and by the EPSRC under grant EP/N509711/1. We also acknowledge the use of the EPSRC-funded Tier 2 facility JADE (EP/P020275/1) and GPU computing support by Scan Computers International Ltd.

REFERENCES

- [1] J. Zhang and Y. Peng, “Hierarchical vision-language alignment for video captioning,” in *MultiMedia Modeling*. Springer, 2019, pp. 42–54.
- [2] L. Gao, Z. Guo, H. Zhang, X. Xu, and H. T. Shen, “Video captioning with attention-based LSTM and semantic consistency,” *IEEE T. Multimedia*, vol. 19, no. 9, pp. 2045–2055, 2017.
- [3] L. Zhou, Y. Zhou, J. J. Corso, R. Socher, and C. Xiong, “End-to-end dense video captioning with masked transformer,” in *Proc. CVPR*, 2018, pp. 8739–8748.
- [4] Z. Zhao, Q. Yang, D. Cai, X. He, and Y. Zhuang, “Video question answering via hierarchical spatio-temporal attention networks,” in *Proc. IJCAI*, 2017, pp. 3518–3524.
- [5] K. Zeng, T. Chen, C. Chuang, Y. Liao, J. C. Niebles, and M. Sun, “Leveraging video descriptions to learn video question answering,” in *Proc. AAAI*, 2017, pp. 4334–4340.
- [6] J. Lei, L. Yu, M. Bansal, and T. L. Berg, “TVQA: localized, compositional video question answering,” in *Proc. EMNLP*, 2018, pp. 1369–1379.
- [7] S. Olivastrì, G. Singh, and F. Cuzzolin, “An end-to-end baseline for video captioning,” *arXiv preprint arXiv:1904.02628*, 2019.
- [8] S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning, “A large annotated corpus for learning natural language inference,” *arXiv preprint arXiv:1508.05326*, 2015.
- [9] Y. Belinkov, A. Poliak, S. M. Shieber, B. V. Durme, and A. M. Rush, “Don’t take the premise for granted: Mitigating artifacts in natural language inference,” in *Proc. ACL*, 2019, pp. 877–891.
- [10] K. R. McKeown, “Discourse strategies for generating natural-language text,” *Artificial intelligence*, vol. 27, no. 1, pp. 1–41, 1985.
- [11] W. J. Levelt, *Speaking: From intention to articulation*. MIT press, 1993, vol. 1.
- [12] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “BLEU: A method for automatic evaluation of machine translation,” in *Proc. ACL*, 2002, pp. 311–318.
- [13] S. Banerjee and A. Lavie, “METEOR: An automatic metric for MT evaluation with improved correlation with human judgments,” in *Proc. of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, 2005, pp. 65–72.
- [14] A. L. Han, D. F. Wong, and L. S. Chao, “LEPOR: A robust evaluation metric for machine translation with augmented factors,” in *Proc. COLING, Posters*, 2012, pp. 441–450.
- [15] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. B. Girshick, “CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning,” in *Proc. CVPR*, 2016, pp. 1988–1997.
- [16] P. Wang, Q. Wu, C. Shen, A. R. Dick, and A. van den Hengel, “FVQA: Fact-based visual question answering,” *IEEE TPAMI*, vol. 40, pp. 2413–2427, 2016.
- [17] P. Wang, Q. Wu, C. Shen, A. Dick, and A. Van Den Henge, “Explicit knowledge-based reasoning for visual question answering,” in *Proc. IJCAI*, 2017, pp. 1290–1296.
- [18] D. Vasilis and T. Lukasiewicz, “Learning structured video descriptions: Automated video knowledge extraction for video understanding tasks,” in *On the Move to Meaningful Internet Systems. OTM 2018 Conferences*, 2018, pp. 315–332.
- [19] K. Curtis, G. Awad, S. Rajput, and I. Soboroff, “HLVU: A new challenge to test deep understanding of movies the way humans do,” in *Proc. ICMR*, 2020, pp. 355–361.
- [20] P. Qi, T. Dozat, Y. Zhang, and C. D. Manning, “Universal dependency parsing from scratch,” in *Proc. CoNLL*, 2018, pp. 160–170.
- [21] D. L. Chen and W. B. Dolan, “Collecting highly parallel data for paraphrase evaluation,” in *Proc. HMM*, 2011, pp. 190–200.
- [22] J. Xu, T. Mei, T. Yao, and Y. Rui, “MSR-VTT: A large video description dataset for bridging video and language,” in *Proc. CVPR*, 2016, pp. 5288–5296.
- [23] G. A. Miller, R. Beckwith, C. Fellbaum, D. Gross, and K. J. Miller, “Introduction to WordNet: An on-line lexical database,” *International Journal of Lexicography*, vol. 3, no. 4, pp. 235–244, 1990.
- [24] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmann, S. Sun, and W. Zhang, “Knowledge Vault: A web-scale approach to probabilistic knowledge fusion,” in *Proc. ACM SIGKDD*, 2014, pp. 601–610.
- [25] Q. Le and T. Mikolov, “Distributed representations of sentences and documents,” in *Proc. ICML*, 2014, pp. 1188–1196.
- [26] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [27] Y. Pan, T. Yao, H. Li, and T. Mei, “Video captioning with transferred semantic attributes,” in *Proc. CVPR*, 2017, pp. 6504–6512.
- [28] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui, “Jointly modeling embedding and translation to bridge video and language,” in *Proc. CVPR*, 2016, pp. 4594–4602.
- [29] X. Zhang, K. Gao, Y. Zhang, D. Zhang, J. Li, and Q. Tian, “Task-driven dynamic fusion: Reducing ambiguity in video description,” in *Proc. CVPR*, 2017, pp. 3713–3721.
- [30] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the properties of neural machine translation: Encoder-decoder approaches,” in *Proc. SSST@EMNLP*, 2014, pp. 103–111.
- [31] J. Carreira and A. Zisserman, “Quo vadis, action recognition? A new model and the kinetics dataset,” in *Proc. CVPR*, 2017, pp. 6299–6308.
- [32] J. Xu, T. Yao, Y. Zhang, and T. Mei, “Learning multimodal attention LSTM networks for video captioning,” in *Proc. ACM MM*, 2017, pp. 537–545.
- [33] Y. Chen, S. Wang, W. Zhang, and Q. Huang, “Less is more: Picking informative frames for video captioning,” in *Proc. ECCV*, 2018, pp. 358–373.
- [34] B. Wang, L. Ma, W. Zhang, and W. Liu, “Reconstruction network for video captioning,” in *Proc. CVPR*, 2018, pp. 7622–7631.