UvA-DARE (Digital Academic Repository)

Video2vec Embeddings Recognize Events when Examples are Scarce

Habibian, A.; Mensink, T.; Snoek, C.G.M.

DOI

[10.1109/TPAMI.2016.2627563](https://doi.org/10.1109/TPAMI.2016.2627563)

Publication date

2017

Document Version

Author accepted manuscript

Published in

IEEE Transactions on Pattern Analysis and Machine Intelligence

[Link to publication](#)

Citation for published version (APA):

Habibian, A., Mensink, T., & Snoek, C. G. M. (2017). Video2vec Embeddings Recognize Events when Examples are Scarce. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(10), 2089-2103. <https://doi.org/10.1109/TPAMI.2016.2627563>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Video2vec Embeddings Recognize Events when Examples are Scarce

Amirhossein Habibian, Thomas Mensink, and Cees G. M. Snoek

Abstract—This paper aims for event recognition when video examples are scarce or even completely absent. The key in such a challenging setting is a semantic video representation. Rather than building the representation from individual attribute detectors and their annotations, we propose to learn the entire representation from freely available web videos and their descriptions using an embedding between video features and term vectors. In our proposed embedding, which we call *Video2vec*, the correlations between the words are utilized to learn a more effective representation by optimizing a joint objective balancing descriptiveness and predictability. We show how learning the *Video2vec* using a multimodal predictability loss, including appearance, motion and audio features, results in a better predictable representation. We also propose an event specific variant of *Video2vec* to learn a more accurate representation for the words, which are indicative of the event, by introducing a term sensitive descriptiveness loss. Our experiments on three challenging collections of web videos from the NIST TRECVID Multimedia Event Detection and Columbia Consumer Videos datasets demonstrate: i) the advantages of *Video2vec* over representations using attributes or alternative embeddings, ii) the benefit of fusing video modalities by an embedding over common strategies, iii) the complementarity of term sensitive descriptiveness and multimodal predictability for event recognition. By its ability to improve predictability of present day audio-visual video features, while at the same time maximizing their semantic descriptiveness, *Video2vec* leads to state-of-the-art accuracy for both few- and zero-example recognition of events in video.

1 INTRODUCTION

This paper strives to recognize events such as *parking a vehicle*, *flash mob*, and *renovating a home* in web video content. A problem of increasing importance in a world that is swiftly adapting to video communication. The leading computer vision and multimedia retrieval solutions for this challenging problem, *e.g.*, [35], [53], [68], [74], [81], all learn to recognize events with the aid of many labeled video examples. However, as events become more and more specific, we consider it unrealistic to assume that ample examples to learn from will be commonly available. In practical use only a handful of video examples, an event name and an event definition may be present, such as the ones in Figure 1. We aim for event recognition when video examples are scarce or even completely absent.

Recognizing events from few or zero examples imposes constraints on the video representation. End-to-end deep learning of the video representation [30], [39], [71] demands too many video examples. In [39] for example, Karpathy *et al.* exploit more than 1 million YouTube videos and their

sport category labels to learn the video representation. Alternatively, Xu *et al.* [81] show how a very deep convolutional neural network (CNN) [63] intended for image classification can be leveraged as representation for event recognition in video. They use responses from intermediate layers of the network to represent frames, which are aggregated over the video using VLAD encoding [29]. When combined with a linear SVM, excellent results on the leading NIST TRECVID event detection benchmarks [57] are reported for scenarios where many and few examples are available. The CNN video representation outperforms more traditional video encodings such as improved dense trajectories [74], [75] and representations combining appearance, motion and audio features [52], [55], [84]. However, both the learned and engineered representations are incapable, nor intended, to recognize events when examples are completely absent. We propose a video representation that can leverage any underlying feature, be it a CNN, improved dense trajectories and/or audio features, while being capable of few- and zero-example event recognition with state-of-the-art accuracy.

The key to few- and zero-example event recognition is to add meaning to the video representation. Inspired by the success in image classification [17], [41], [58], many rely on the predictions made by a set of attribute classifiers [5], [16], [22], [26], [34], [44]–[46], [48], [50], [82]. In [22] Habibian *et al.* study the properties of 1,346 attribute classifiers trained from ImageNet [61] and TRECVID [57] for representing and recognizing events in web video. Rather than pre-specifying and manually labeling each individual attribute in advance, the attributes can be learned on top of imagery and (weak) labels harvested from the web [10], [12], [79], [83]. To assure visual predictability of the discovered attribute detectors a common tactic is to leverage part of the harvested data for validation [4]. A drawback is that many attribute labels rarely occur. For these infrequent labels only a limited number of positive examples are available, which leads to a biased estimation of their classification reliability. Hence, many of the discovered attributes overfit to their small training set and do not generalize well for new videos. We also discover our semantic representation from the web, but rather than selecting individual, and often unreliable, classifiers per attribute label, we prefer to combine labels automatically into more predictable attributes. By doing so, more training examples are available and a more robust representation is obtained, without losing descriptive ability.

In this paper we present a semantic representation learn-

Event Name: Winning a race without a vehicle

Definition: An individual (or more) succeeds in reaching a pre-determined destination before all other individuals, without vehicle assistance or assistance of a horse or other animal. Racing generally involves accomplishing a task in less time than other competitors. The only type of racing considered relevant for the purposes of this event is the type where the task is traveling to a destination, completed by a person(s) without assistance of a vehicle or animal. Different types of races involve different types of human ...

Fig. 1. Video exemplars (top) and the textual definition (bottom) of the event *winning a race without a vehicle* to illustrate one of the events studied in this paper. Following the NIST TRECVID evaluation guidelines [57], the textual definition is for zero-example recognition, and the ten provided video exemplars are for few-example event recognition.

ing algorithm for videos. Instead of relying on pre-specified attribute labels, we learn the representation from freely available web videos and their descriptions. We propose an embedding between the video features and their textual descriptions, which is learned by utilizing the correlations between the words in the descriptions. We learn the embedding by minimizing a joint objective function balancing the descriptiveness and predictability of the learned video representation. Our embedding is able to leverage the multiple modalities which coexist in video to learn a more reliable semantic representation. We call our embedding *Video2vec*.

A preliminary version of this article appeared as [21]. The current version adds *i)* zero-example event recognition *ii)* representation learning for multiple video modalities, *iii)* event specific representation learning, *iv)* an additional TRECVID video data set, revised experiments, and improved baselines, all using a recent video CNN feature [67], [81], and *v)* a new related work section, which will be discussed next.

2 RELATED WORK

We focus on three directions of related work we deem most closely connected to ours.

2.1 Representations for Event Recognition

Until recently most event recognition methods exploit video representations based on densely extracted low-level visual features, such as HOG/HOF [42], or MBH [56], [75], often combined with audio features like MFCC features [37], [53], [68]. Currently, most methods extract frame-based deep convolutional neural network (CNN) features [40], [63], [67], using the responses from intermediate layers of the CNN, which is pre-trained on ImageNet images [61], see e.g., [10], [34], [51], [81], [84]. To obtain per video descriptors the local/frame-based descriptors are aggregated by their mean, by using the Fisher vector [62] or VLAD encoding [29]. Despite the fact that these low-level based representations obtain state-of-the-art recognition performance,

they suffer from two drawbacks. First, because of their high dimensionality, training effective classifiers on these representations requires a sufficient number of training examples to prevent overfitting. Second, all these representations are incapable of providing a semantic interpretation of the video, which is crucial for zero-example recognition.

Semantic Representations To obtain a semantic representation for videos, inspiration is obtained from describing images with attributes [19], [41], and the video is represented by its attribute predictions. Creating the training data for a set of task-specific attribute classifiers manually involves lots of annotation effort, which is restrictive. Therefore, often public available datasets, such as ImageNet [61] and TRECVID [57], are used to train the attribute classifiers [22], [45], [48]. Although this overcomes the need for (additional) manual annotation, the attributes of these datasets are not necessarily descriptive for event recognition.

To tune the attributes for the task at hand, several works aim to automatically discover the attributes from web images/videos and their textual descriptions [10], [12], [79], [83]. For example Wu *et al.* [79], start by selecting the most frequent/relevant terms from a set of provided event descriptions as attributes. Then, inspired by Berg *et al.* [4], they use Internet search engines to gather positive examples. They assure visual predictability of the discovered attribute detectors by cross-validation. Similarly the work of Ye *et al.* [10] relies on WikiHow to obtain event descriptions, and the visual predictability of the selected terms is ensured by keeping only those terms which are present in existing image classification datasets. We refer to these methods as *term attributes* since they all discover the attributes from the terms in the descriptions.

Despite their effectiveness in obtaining training data, term attributes suffer from three drawbacks. First, descriptions have a long-tail distribution, therefore many terms rarely occur and these will not have enough positive examples to train reliable visual classifiers. Second, term attributes are selected mostly based on their visual prediction accuracy, while for effective event recognition the selected term attributes should also be descriptive for the target events. Third, contextual information is lost, since term attributes are learned independently by binary classifiers.

In this paper, we also learn the representation from videos and their descriptions. But, rather than selecting individual, and often unreliable, term attributes, we learn the entire representation by automatically combining the terms through embedding. In our embedding, the correlations between the terms are utilized to learn a more effective representation, which is predictable *and* descriptive.

2.2 Embedding Videos and their Descriptions

To learn correspondences between the visual domain and textual descriptions different embedding methods have been proposed for various purposes *i.e.*, image annotation [6], image classification [2], [78], image captioning [38], [73], video to text translation [15], [80] and cross-modal retrieval [11], [18], [54].

Canonical correlation analysis (CCA) [25] is the classical unsupervised manner to relate different modalities and can be used for cross-modal retrieval [11]. CCA is

the multimodal generalization of PCA, and computed as a generalized eigenvalue problem on the cross-covariance matrix between the visual and textual features. It finds a sequence of uncorrelated projections in which the cross-correlation between the modalities is maximized. This is not necessarily suited to obtain a discriminative representation, as is also validated by our experiments.

Probabilistic topic models, such as the seminal correspondence latent Dirichlet allocation [6] and its extensions [14], [60], learn correspondences by extracting a set of correlated visual and textual topics from imagery and their captions. Despite their effectiveness for captioning images and videos [15], these methods are not designed to learn a representation for the purpose of recognition. Moreover, we note that by design these models are only applicable on discrete features, and therefore they cannot leverage the state-of-the-art video features used in our paper.

Embeddings for multi-class image classification, such as [2], [78] find a low-dimensional subspace in which multi-class classification is performed. The subspace is found jointly with the multi-class classifier by minimizing a classification loss [78], or a zero-shot classification loss [2]. In contrast to this multi-class image classification setting, we are interested in finding descriptions modeled as a multi-label video classification problem. Moreover, we argue that solely optimizing for classification is not sufficient to obtain a predictable and descriptive video embedding.

Recently deep neural network architectures have been proposed to learn multimodal correspondences for retrieval [18], [54] and captioning [38], [73]. Notably the multimodal recurrent neural networks capture the sequential ordering between the words in image captions and generate more accurate image captions, compared to the probabilistic topic models. However, training deep neural networks generally requires large amounts of training examples, while our purpose is to learn from few- and zero-examples.

2.3 Zero-Example Event Recognition

In zero-example event recognition the goal is to recognize an event based on a given textual *event definition*, without using any video examples. The event definition is usually provided in the form of a title and description, see Figure 1. This zero-example setting is beyond the conventional zero-shot image recognition of objects and scenes [3], [17], [41], [59], where often a training set of related classes is available as well as pre-specified class-to-attribute mappings. This challenging event recognition problem has attracted attention because of its high practical value and the corresponding TRECVID benchmark task initiated by NIST [57]. The common approach is to represent videos and the event queries using a semantic representation, and to rank all video representations based on the cosine similarity to the event query.

To represent the video mostly attributes and term-attributes [10], [12], [32], [33], [79], [83] are exploited. Extensions include combining attributes using logical operators [20], using video-segments to detect attributes [49], and adjustment of attribute scores using an ontology structure [34]. In addition to attributes, automatic speech recognition and optical character recognition have been considered to enrich the semantic video representations [13], [34], [79].

TABLE 1
Summary of the core notation used for Video2vec.

Notation	Description
N	Number of videos
M	Number of unique words in descriptions
D	Dimensionality of low-level feature
J	Number of low-level features to fuse
k	Dimensionality of Video2vec embedding
$\mathbf{X} \in \mathbb{R}^{D \times N}$	Matrix of low-level video features
$\mathbf{Y} \in \{0, 1\}^{M \times N}$	Matrix of binary term vectors
$\mathbf{W} \in \mathbb{R}^{D \times k}$	Video2vec visual projection
$\mathbf{A} \in \mathbb{R}^{M \times k}$	Video2vec textual projection
$\mathbf{S} \in \mathbb{R}^{k \times N}$	Video2vec embedding
$\mathbf{H} \in \mathbb{R}^{M \times M}$	Diagonal matrix with per-term weights
$\mathbf{x}_i, \mathbf{y}_i, \mathbf{s}_i$	The column representing the i -th video

To answer the event query, usually term matching between the semantic video representation and the event definition is performed. The retrieval can be enriched by using contextual information, such as word embeddings and term co-occurrences [8], [83]. For the final ranking, the cosine similarity can be extended by pseudo relevance feedback mechanisms, such as self-paced ranking [32]. This has been used to improve the zero-example event recognition by re-ranking [8], [33], [34]. Our representation learning is orthogonal to these efforts and can be joined with them to further improve the event recognition performance.

3 VIDEO2VEC EMBEDDINGS

Our goal is to learn a representation function $f : \mathcal{X} \rightarrow \mathcal{S}$, which maps each low-level video feature $\mathbf{x}_i \in \mathcal{X}$ into the semantic representation $\mathbf{s}_i \in \mathcal{S}$. The representation function is trained on a collection of videos and their semantic descriptions. We represent the descriptions in the form of binary term vectors $\mathbf{y}_i \in \mathcal{Y}$. For event recognition, the $\mathbf{s}_i$ should have two properties: First, it should be descriptive enough to provide a characteristic representation of each video. Second, each dimension of the representation, corresponding to a semantic concept, should be predictable from the low-level video features.

In the attribute approach to learning the representation, the semantic representations are defined to be in the same space as the term vectors ($\mathcal{S} = \mathcal{Y}$). However, in practice, the term vectors are noisy and sparse, which undermines their effectiveness as labels to train f . Therefore, we propose to learn the $\mathbf{s}_i$ on a lower dimensional projection of the term vectors, which are less sparse and less noisy. To remedy that the lower dimensional term vectors may be less descriptive for the video content, we formulate f as an embedding, coined as Video2vec, which is learned by balancing the descriptiveness and predictability as a joint optimization.

We first introduce the Video2vec embedding and its application for few- and zero-example event recognition (Section 3.1). Then we generalize our embedding to fuse multimedia features (Section 3.2). Finally, we extend the embedding to be learned specifically per event (Section 3.3). We summarize our notation conventions in Table 1.

3.1 Objective Function

We start from a dataset of videos, represented by video features $\mathbf{X}$, and their textual descriptions, represented by

binary term vectors $\mathbf{Y}$, indicating which words are present in each video description. Then, our Video2vec representation is learned by minimizing:

$$L_V(\mathbf{A}, \mathbf{W}) = \min_{\mathbf{S}} L_d(\mathbf{A}, \mathbf{S}) + L_p(\mathbf{S}, \mathbf{W}), \quad (1)$$

where $\mathbf{A}$ is the textual projection matrix, $\mathbf{W}$ is the visual projection matrix, and $\mathbf{S}$ is the Video2vec embedding. The loss function L_d corresponds to our first objective for learning a descriptive Video2vec, and the loss function L_p corresponds to our second objective for learning a predictable Video2vec. The Video2vec embedding $\mathbf{S}$ interconnects the two loss functions. **Descriptiveness.** For the L_d function, we use a variant of regularized Latent Semantic Indexing (LSI) [76]. This objective minimizes the quadratic error between the original video descriptions $\mathbf{Y}$, and the reconstructed translations obtained from $\mathbf{A}$ and $\mathbf{S}$:

$$L_d(\mathbf{A}, \mathbf{S}) = \frac{1}{2} \sum_{i=1}^N \|\mathbf{y}_i - \mathbf{A}\mathbf{s}_i\|_2^2 + \lambda_a \Omega(\mathbf{A}) + \lambda_s \Psi(\mathbf{S}), \quad (2)$$

where $\Psi(\cdot)$ and $\Omega(\cdot)$ denote regularization functions, and $\lambda_a \geq 0$ and $\lambda_s \geq 0$ are regularizer coefficients. We use the squared Frobenius norm for regularization, which is the matrix variant of the ℓ_2 regularizer, *i.e.*, $\Omega(\mathbf{A}) = \frac{1}{2} \|\mathbf{A}\|_F^2 = \frac{1}{2} \sum_i \|\mathbf{a}_i\|_2^2 = \frac{1}{2} \sum_{ij} a_{ij}^2$, the sum of the squared matrix elements. Similarly for the Video2vec matrix $\Psi(\mathbf{S}) = \frac{1}{2} \|\mathbf{S}\|_F^2$.

The main difference with regularized Latent Semantic Indexing [76] is their ℓ_1 regularizer, $\Omega(\mathbf{A}) = \sum_i \|\mathbf{a}_i\|_1$, which enforces sparsity in the textual projection $\mathbf{A}$. However, with our larger representation (typically we use a dimensionality of k between 512 and 2,048 in our experiments compared to only $k = 20$ in [76]) and fewer number of unique words (around 10K, compared to 100K), enforcing sparsity is not necessary for good performance.

Note that other textual embeddings can be formulated similar to Eq. (2), when appropriate regularization functions $\Omega(\cdot)$ and $\Psi(\cdot)$ are used. For example using $\Omega(\cdot) = \|\cdot\|_1$ enforces sparsity [24], [47], [76], or in the extreme case that $\mathbf{A}$ is constrained such that each column has a single non-zero value, the objective becomes very close to methods that select the best single term labels [4], and when $\mathbf{A}$ is enforced to preserve the taxonomical term relations, the objective resembles taxonomy embedding [72], [77].

Predictability. The L_p function measures the occurred loss between the Video2vec $\mathbf{S}$ and the embedding of video features using $\mathbf{W}$. Since the Video2vec $\mathbf{S}$ is real valued, as opposed to a binary or multi-class encoding, we can not rely on standard classification losses such as the hinge-loss used in SVMs. Therefore, we define L_p as a regularized regression, similar to ridge regression:

$$L_p(\mathbf{S}, \mathbf{W}) = \frac{1}{2} \sum_{i=1}^N \|\mathbf{s}_i - \mathbf{W}^\top \mathbf{x}_i\|_2^2 + \lambda_w \Theta(\mathbf{W}), \quad (3)$$

where we use (again) the Frobenius norm for regularization of the visual projection matrix $\mathbf{W}$, $\Theta(\mathbf{W}) = \frac{1}{2} \|\mathbf{W}\|_F^2$, and λ_w is the regularization coefficient.

Joint optimization. To handle large scale datasets and state-of-the-art high-dimensional visual features, *e.g.*, Fisher vectors [62] on video features [75] or deep learned representations [40], we employ a Stochastic Gradient Descent

Algorithm 1: Pseudocode for learning the Video2vec

```

input :  $\mathbf{X}, \mathbf{Y}, k, \eta$  (step-size),  $m$  (max-epochs)
output :  $\mathbf{W}$  and  $\mathbf{A}$ 

 $\mathbf{A}$ , and  $\mathbf{S} \leftarrow$  SVD decomposition of  $\mathbf{Y}$ 
 $\mathbf{W} \leftarrow$  random (zero-mean)

for  $e \leftarrow 1$  to  $m$  do
  for  $i \leftarrow 1$  to  $N$  do
    Pick a random video-description pair  $(\mathbf{x}_t, \mathbf{y}_t)$ 
    Compute gradients w.r.t.  $\mathbf{A}, \mathbf{W}$  and  $\mathbf{s}_t$ 
    Update parameters:
       $\mathbf{A} \leftarrow \mathbf{A} - \eta_t \nabla_{\mathbf{A}} L_V$       see Eq. (4)
       $\mathbf{W} \leftarrow \mathbf{W} - \eta_t \nabla_{\mathbf{W}} L_V$       see Eq. (5)
       $\mathbf{S} \leftarrow \mathbf{s}_t - \eta_t \nabla_{\mathbf{s}_t} L_V$       see Eq. (6)
  end
end
return  $\mathbf{W}$  and  $\mathbf{A}$ 

```

(SGD) [7] optimization, summarized in Algorithm 1. The number of passes over the datasets (*epochs*) and the step-size η are hyper-parameters of SGD.

The Video2vec objective function, as given in Eq. (1), is convex with respect to matrix $\mathbf{A}$ and $\mathbf{W}$ when the embedding $\mathbf{S}$ is fixed. In that case, the joint optimization is decoupled into Eq. (2) and Eq. (3), which are both reduced to a standard ridge regression for a fixed $\mathbf{S}$. Moreover, when both $\mathbf{A}$ and $\mathbf{W}$ are fixed, the objective in Eq. (1) is convex *w.r.t.* $\mathbf{S}$. Therefore we use standard SGD by computing the gradients of a sample *w.r.t.* the current value of the parameters, and we minimize $\mathbf{S}$ jointly with $\mathbf{A}$ and $\mathbf{W}$.

Lets denote a randomly sampled video and description pair at step t by $(\mathbf{x}_t, \mathbf{y}_t)$, and let $\mathbf{s}_t$ denote the current Video2vec embedding of sample t . The gradients of Eq. (1) for this sample *w.r.t.* $\mathbf{A}, \mathbf{W}$ and $\mathbf{s}_t$ are given by:

$$\nabla_{\mathbf{A}} L_V = -(\mathbf{y}_t - \mathbf{A}\mathbf{s}_t)\mathbf{s}_t^\top + \lambda_a \mathbf{A}, \quad (4)$$

$$\nabla_{\mathbf{W}} L_V = -\mathbf{x}_t(\mathbf{s}_t - \mathbf{W}^\top \mathbf{x}_t)^\top + \lambda_w \mathbf{W}, \text{ and} \quad (5)$$

$$\nabla_{\mathbf{s}_t} L_V = -\mathbf{A}^\top(\mathbf{y}_t - \mathbf{A}\mathbf{s}_t) + (\mathbf{s}_t - \mathbf{W}^\top \mathbf{x}_t) + \lambda_s \mathbf{s}_t. \quad (6)$$

The effect of joint learning descriptiveness and predictability, becomes clear in Eq. (6), where both the textual projection matrix $\mathbf{A}$ and visual projection matrix $\mathbf{W}$ contribute to learning the Video2vec embedding $\mathbf{S}$. This embedding $\mathbf{S}$ is subsequently used to obtain the textual projection $\mathbf{A}$ matrix, in Eq. (4), and the visual projection $\mathbf{W}$ matrix, in Eq. (5). This leads to the Video2vec embedding, which is both descriptive, by preserving the textual information, and predictable, by minimizing the visual prediction loss.

The parameters $\mathbf{A}$, $\mathbf{S}$, and $\mathbf{W}$ can be initialized randomly. However, to speed up convergence, we propose to initialize $\mathbf{A}$ and $\mathbf{S}$ by a rank-reduced singular value decomposition (SVD) of term vectors $\mathbf{Y}$. This corresponds to the solution of the unregularized LSI objective of Eq. (2). For this purpose, $\mathbf{A}$ and $\mathbf{S}$ are initialized by the k first eigenvectors of $\mathbf{Y}\mathbf{Y}^\top$ and $\mathbf{Y}^\top \mathbf{Y}$, respectively.

After training the visual and textual projection matrices, they are used to predict the Video2vec representation and the term vector of each video. In the case that both a video $\mathbf{x}_i$ and description $\mathbf{y}_i$ are given, we could obtain the Video2vec representation by returning $\mathbf{s}_i$ from Eq. (1), while keeping both $\mathbf{A}$ and $\mathbf{W}$ fixed. However, in practice most videos are

not provided with a description. Therefore, we use

$$\mathbf{s}_i = \mathbf{W}^\top \mathbf{x}_i, \quad (7)$$

to predict our Video2vec representation from the low-level video features $\mathbf{x}_i$. Moreover, using the predicted representation $\mathbf{s}_i$, the term vectors for each unseen video are predicted as follows:

$$\hat{\mathbf{y}}_i = \mathbf{A} \mathbf{s}_i = \mathbf{A} \mathbf{W}^\top \mathbf{x}_i, \quad (8)$$

where the words with the highest values are most relevant for this video, see the illustration in Figure 2

Zero-Example Event Recognition. To enable zero-example recognition using the Video2vec embedding, we employ the following steps: First, each test video is represented by predicting its term vector $\hat{\mathbf{y}}_i$ using Eq. (8), based on the pre-trained embeddings. Second, we translate the textual event definition into the event query, denoted as $\mathbf{y}^e \in \mathbb{R}^M$, by matching the words in the event definition with the M unique words in the Video2vec dictionary. Finally, the zero-example ranking is obtained by measuring the similarity between the video representations and the event query based on the cosine similarity:

$$s_e(\mathbf{x}_i) = \frac{\mathbf{y}^{e\top} \hat{\mathbf{y}}_i}{\|\mathbf{y}^e\| \|\hat{\mathbf{y}}_i\|}. \quad (9)$$

3.2 Video2vec Fusion

Videos are inherently multimodal. In general any video contains appearance, motion, and audio cues and sometimes even textual information in the form of subtitles or speech recognition scripts. Fusing the different modalities is typically achieved by early-fusion, *i.e.*, fusion at the level of the representations, and late-fusion, *i.e.*, fusion at the level of prediction scores [64]. Both fusion strategies have been shown to be effective for understanding complex events as well, *e.g.*, [37], [53], [70], [79]. We propose *Video2vec^F*, which extends the Video2vec embedding by learning the semantic representation from multiple modalities.

A straightforward approach to learn the multimodal semantic representation is by fusing multiple Video2vec embeddings, which are independently trained per modality. Our intuition is that the semantic representation is more effective if it is predictable from *all* the modalities rather than from each individual modality. For *Video2vec^F* we adjust the predictability loss, of Eq. (3), to incorporate a weighted combination of the predictability from all J modalities as follows:

$$L_p^F(\mathbf{S}, \mathcal{W}) = \sum_{j=1}^J L_p(\mathbf{S}, \mathbf{W}^j) \quad (10)$$

where $\mathbf{S}$ is the multimodal Video2vec embedding, and $\mathcal{W} = \{\mathbf{W}^j, j = 1 \dots J\}$ is a set of projection matrices from the J modalities. Each feature projection matrix $\mathbf{W}^j \in \mathbb{R}^{D_j \times k}$ projects the low-level feature $\mathbf{x}_i^j \in \mathbb{R}^{1 \times D_j}$ extracted from the video into the Video2vec representation $\mathbf{s}_i$. To balance the impact of feature dimensionality in the multimodal predictability loss, all low-level features from $\mathbf{x}_i^j$ are ℓ_2 normalized. Moreover, instead of weighting each modality equally, as in Eq. (10), a term $\gamma_j \geq 0$ could be added to

Fig. 2. Video2vec prediction: From the low-level video features the Video2vec representation and the term vector are predicted using the visual projection matrix $\mathbf{W}$ and the textual projection matrix $\mathbf{A}$.

weight the importance of each modality, if sufficient training examples are available for their cross-validation.

The objective function Eq. (1) is still convex with respect to the parameters $\mathbf{S}$, $\mathbf{A}$, and $\mathbf{W}^j$ when the other parameters are fixed. However, the gradient with respect to $\mathbf{s}_t$, Eq. (6) becomes:

$$\nabla_{\mathbf{s}_t} L_V = -\mathbf{A}^\top (\mathbf{y}_t - \mathbf{A} \mathbf{s}_t) + (\mathbf{s}_t - \sum_j \mathbf{W}^{j\top} \mathbf{x}_t^j) + \lambda_s \mathbf{s}_t. \quad (11)$$

It can be seen that all the modalities are jointly contributing to learn the multimodal Video2vec representations $\mathbf{S}$.

After training the textual and feature projection matrices, they are used to extract the multimodal Video2vec representation. Each feature projection matrix $\mathbf{W}^j$ predicts the Video2vec representation based on its underlying modality as follows:

$$\mathbf{s}_i^j = \mathbf{W}^{j\top} \mathbf{x}_i^j. \quad (12)$$

The final Video2vec representation is obtained by aggregating the per-modality representations. We experimentally observe that concatenation performs slightly better than sum or max pooling. The final representation we use is a concatenation of the per-modality embeddings: $\mathbf{s}_i = [\mathbf{s}_i^1, \dots, \mathbf{s}_i^J]$.

3.3 Event Specific Video2vec

Not all words are equally important for recognizing an event. For each event there are a few terms which are highly informative, while the others are not. For the event “attempting a bike trick” words such as *bicycle*, *jumping*, and *ramp* are highly distinctive while *beach*, *cake*, and *children* are less informative.

In Section 3.1, the descriptiveness loss L_d is defined as the overall error in reconstructing *all* the words from the Video2vec representations, see Eq. (2). With this definition, the loss is biased towards the more frequent words, as minimizing their reconstruction error leads to a higher decrease in the overall error. Consequently, the words which

are infrequent in the descriptions might be discarded, which degrades their prediction accuracy from video features. This undermines the effectiveness of our representation learning for event recognition.

We extend the Video2vec embedding to learn a video representation, able to predict the informative words of events. Our extension, which we coin as *Video2vec_E*, is learned per event. It minimizes the reconstruction error of the terms with respect to their importance for event description, rather than their frequency in the Video2vec train data. We introduce a *term sensitive* descriptiveness loss:

$$L_d^{ts}(\mathbf{A}, \mathbf{S}) = \frac{1}{2} \sum_{i=1}^N \|\mathbf{H}_e^{\frac{1}{2}} (\mathbf{y}_i - \mathbf{A}\mathbf{s}_i)\|_2^2 + \lambda_a \Omega(\mathbf{A}) + \lambda_s \Psi(\mathbf{S}), \quad (13)$$

where $\mathbf{H}_e \in \mathbb{R}^{M \times M}$ is a diagonal matrix, denoting the importance of each word for describing an event. By setting a relatively high value for h_{jj} for term j , its reconstruction error is more penalized compared to the other terms. Hence the word is expected to be more precisely reconstructed.

We determine the term importance matrix $\mathbf{H}_e$ by relying on the presence/absence of terms in the textual event definitions. Our assumption is that the words, which are present in event definitions are more important than the absent words. We set each element of the importance matrix h_{jj} to α , if the word j is present, and $1 - \alpha$ if the word j is absent in the event definition. α is a balancing parameter between 0 and 1, which should be higher than 0.5 to assign more importance to the present words. We empirically set this parameter to 0.75 in all our zero-example event recognition experiments.

Multimodal fusion. Finally, we can leverage the multimodal features for learning event specific embeddings by combining the multimodal predictability loss and the term sensitive descriptiveness in a joint objective:

$$L_V(\mathbf{A}, \mathbf{W}) = \min_{\mathbf{S}} L_d^{ts}(\mathbf{A}, \mathbf{S}) + L_p^{\mathcal{F}}(\mathbf{S}, \mathbf{W}). \quad (14)$$

We coin the learned video representation *Video2vec_E^F*, since it is learned event specific on multiple video modalities.

4 EXPERIMENTAL SETUP

4.1 Datasets

We first introduce the dataset used for learning the Video2vec embeddings. Then, we detail the video datasets by which we evaluate the event recognition experiments.

4.1.1 Video2vec Learning

In all experiments, we learn the Video2vec embeddings on the *VideoStory46K* [21] dataset. This collection encompasses 45,826 videos harvested from YouTube, with a total length of 743 hours. Every video comes with a short title caption provided by the user who has uploaded the video, as shown in Figure 3. There are 19,159 unique terms in the captions, most of them occurring infrequently. We filter out the terms occurring only once as they generally are misspelled terms, numbers, or noisy terms. It provides us with 9,828 unique terms, which are used in our experiments.

Fig. 3. Example videos and title captions from the VideoStory46K dataset [21], which we use for Video2vec representation learning.

4.1.2 Event Recognition Evaluation

We perform our event recognition on the challenging TRECVID Multimedia Event Detection corpus [65] and the Columbia Consumer Video collection [36]. These contain more than 42K videos in total, including user generated web videos with a large variation in quality, length and content.

TRECVID Multimedia Event Detection (MED) [65]. This dataset is introduced by NIST as a benchmark for event recognition. We perform our experiments on the two latest releases of the dataset, referred to as *MED 2013* and *MED 2014*. Each dataset includes videos from 20 complex events, with 10 events overlapping as listed in Table 5. Each dataset includes three labeled video partitions: Event Kit training, Background training, and test set MED including 200, 5K, and 27K videos, respectively¹. Apart from the videos, a textual definition is provided per event, which explicates the event as unformatted plain text, such as the one shown in Figure 1.

We perform our few- and zero-example event recognition experiments by exactly following the *10Ex* and *0Ex* evaluation procedure outlined by the NIST TRECVID event recognition task [57]. In the few-example event recognition experiments, training data for each event is composed of 10 positive videos from the Event Kit training data along with about 5K negative videos from the Background training data. The event recognition results for each event classifier are reported on the 27K videos from test set MED. In our zero-example event recognition experiments, we rely on the provided textual event definitions to create an event query vector. Then the event recognition performance is reported on the test set MED.

Columbia Consumer Video (CCV) [36]. This dataset contains 9,317 user-generated videos from YouTube including over 210 hours in total. The dataset contains ground truth annotations at video-level for 20 semantic categories, *i.e.*, *wedding reception* and *music performance*. We use the standard partitioning of the dataset, but we use only 10 positive examples per event in the training data. These 10 are selected based on alphabetical order of the respective video names, we ignore the remaining positive examples in the train set. We report event recognition results on the standard test partition. We denote our redefinition of the CCV dataset for few-example event recognition as *CCV₁₀*.

4.2 Video Features

To cancel out accidental effects of the choice for the underlying features, we consider the same set of appearance,

1. There is also a PROGRESS set with 98K videos, but this partition is for blind testing by NIST only.

Fig. 4. Effect of embedding. The Video2vec embedding outperforms the term attribute and term attribute-f baselines, which are directly learned from the terms without embedding.

motion and audio features for all our experiments, the various baselines, and our Video2vec variants. The appearance features are used for all the experiments. The motion and audio features are added for the experiments involving the multimodal fusion.

Appearance. We adopt the *video CNN* representation [81] as appearance features for event recognition, but found the very deep network of [67] to perform slightly better than [63]. For each video the frames are extracted by uniformly sampling every two seconds. Then, the CNN descriptors are extracted per frame as the 1K dimensional responses from the last fully connected layer (pool₅) of the Google Inception network [67]. We train the network on the 15,293 ImageNet categories with more than 200 examples, using the Caffe toolbox [31]. The final video CNN features are obtained by aggregating the frame descriptors over each video by VLAD encoding [29] with a codebook size of 20, resulting in a 20,480 dimensional vector.

Motion. We use the MBH descriptors along the motion trajectories [75] as motion features. The extracted 288-dimensional descriptors are reduced to 128 dimensions using PCA and are then aggregated per video using a Fisher vector [62], with 128 Gaussians resulting in a 32,768 dimensional vector. Each Fisher vector is power normalized, with $\alpha = 0.2$, as in [27].

Audio. We extract MFCCs descriptors [1] over a 10ms window. The descriptors consist of 13 values, 30 coefficients and the log-energy, along with their derivatives and the second derivatives. The MFCC descriptors are aggregated by Fisher vectors using a Gaussian Mixture Model with 256 components, resulting in a 46,080 dimensional vector.

4.3 Implementation Details

We learn the Video2vec embeddings by using 75% of the VideoStory46K dataset for training and 25% for validation to set the hyper-parameters of our model ($\lambda_w, \lambda_a, \lambda_s$) and of SGD (number of epochs, η). As the validation criterion we rely on the objective function value, when using $S = W^T X$. The step size η_t is fixed during training.

For few-example event recognition, the event classifiers are trained as binary SVM with RBF kernels, as suggested in [50]. Following [81], [84] we set the SVM regularization and the RBF kernel parameters by a default value of 1, as the train set is not big enough for cross-validation.

As evaluation criteria we follow the standard convention in the literature [36], [57] by relying on the average precision

TABLE 2

Five selected dimensions of the Video2vec embedding trained on VideoStory46K. Each dimension is visualized by reporting its ten most relevant words, as determined by A^{-1} . By design, correlated terms are usually combined into one dimension of Video2vec.

car	diamond	cute	cell	room
wheel	style	cat	mobile	living
man	pear	play	home	designer
truck	set	dog	call	furniture
test	gold	meow	solavei	picture
woman	engagement	baby	plan	fedisa
drive	body	kitty	business	car
front	heart	cutest	girl	truck
driver	dress	black	woman	traffic
rack	card	bath	man	fire

(AP) per event, and we report the mean average precision (mAP) for overall accuracy.

5 EXPERIMENTS

5.1 Video2vec for Few-Example Event Recognition

5.1.1 Effect of Embedding

We first compare Video2vec with term attribute baselines. The baselines learn the representations directly from terms without any embedding. We evaluate all video representations for few-example event recognition using a dimensionality varying from 32 to 8,192.

1. Term attributes. This representation is extracted by following the tradition of predicting relevant individual words from the video descriptions *e.g.*, [10], [12], [79], [83]. A linear SVM classifier is trained per term. The classifiers which have the highest prediction accuracy, based on a 2-fold cross-validation [4], are selected as term attributes.

2. Term attributes-f. Rather than using cross-validation to select the term attributes it simply selects the words with the highest *frequency* in the descriptions.

Results. The results in Figure 4 demonstrate that Video2vec embeddings outperform the term attribute and term attribute-f baselines on all the three test sets.

Term attributes, which relies on the estimated reliability of individual word classifiers suffers from two drawbacks. First, many of the visual terms are very specific and therefore incapable of characterizing the events of interest *i.e.*, necklace, suitcase, and earring. Although

Fig. 5. Video2vec vs Other Embeddings. Video2vec outperforms the CCA and the Description embedding on all three test sets. The Description embedding is the closest competitor, but it suffers from embedding correlated terms which are visually dissimilar. CCA, uses the same objective function to learn the visual and textual embeddings, which is suboptimal due to intrinsic differences between the visual and textual features.

these words can be accurately predicted from videos, they are incapable of providing a characteristic representation for event recognition. Second, many of the words rarely occur in video descriptions. Hence, only a limited number of positive examples are available to learn the word classifiers, which leads to a biased estimation of their reliability. Consequently, many of the discovered visual terms overfit to their small training set and do not generalize well for new videos.

The drawbacks of term attributes are relaxed by simply relying on the most frequent words. We observe the most frequent terms usually refer to characteristic attributes of events which are frequently used by humans when describing a video, *i.e.*, *car*, *girl*, and *kid*. Because of their large number of positive examples, the trained visual classifiers from term attributes-*f* are in general more reliable.

The Video2vec embedding represents the words in a reduced-dimensional space, where correlated terms are usually combined together, as visualized in Table 2. Combining correlated words leads to less correlation between dimensions of the learned representation. Moreover, as the positive examples for all correlated terms are combined, it provides more positive video examples to train visual classifiers, often leading to better accuracy.

Besides its effectiveness for few-example event recognition, the Video2vec embedding also improves the representation learning efficiency by training visual embeddings for combination of words rather than individual words.

5.1.2 Video2vec vs Other Embeddings

We compare the effectiveness of our proposed Video2vec embedding with two alternative embeddings:

1. **CCA embedding.** This baseline learns the textual and visual projections by CCA [11], which maximizes the cross-correlation between the video features and descriptions. We experimentally observed that the embedding is even more effective when we PCA-reduce the video features to a dimensionality of 1,024 before learning the CCA embedding.

2. **Description embedding.** Similar to the Video2vec, this embedding is learned by minimizing the descriptiveness and predictability losses, but in two disjoint steps: The textual projection is first learned based on the regularized Latent Semantic Indexing [76], as in Eq. (2). Then the visual projection is learned separately, by minimizing the error for predicting the embedded descriptions from the video features based on ridge regression, as in Eq. (3).

TABLE 3
Video2vec vs Other Representations for few-example event recognition. Video2vec outperforms the alternatives on all three test sets.

	Low-Level	Attributes	Video2vec
MED 2013	28.1	22.5	32.4
MED 2014	21.8	17.2	24.2
CCV₁₀	50.0	48.8	54.8

Results. The results in Figure 5 demonstrate that the Video2vec embedding outperforms the CCA and the Description embedding baselines on all the three test sets for a dimensionality larger than 256.

We explain the gain over CCA by the fact that CCA is a symmetric embedding, it learns both the textual and visual projections in the same way. However, the textual and visual features have different distributions and properties, which may require different objective functions to learn their projections. This is achieved by Video2vec as it relies on two separate LSI and ridge regression loss functions to learn the textual and visual projections.

We explain the improvement over the Description embedding by the fact that combining the words based on textual correlation only does not necessarily imply that the corresponding video is visually correlated as well. As it happens, the terms *puppy* and *kid* have a high correlation in the descriptions but are visually dissimilar. Combining these two words together, as is done by Description embedding, undermines the accuracy of the classifiers predicting them from videos. In contrast, in a Video2vec the correlated words are combined only if their combination improves their classifier prediction. It prevents the combination of correlated terms which are visually dissimilar.

From now on we fix the Video2vec representation to 2,048 dimensions, which is optimal based on Figure 4.

5.1.3 Video2vec vs Other Representations

Next, we compare Video2vec with state-of-the-art video representations for few-example event recognition:

1. **Low-Level.** In this baseline, the event classifiers are trained directly on the low-level video representations, without extracting a semantic representation. We rely on the video CNN features of [81], based on our implementation detailed in Section 4.2.

Fig. 6. From 1 to 100 examples. When the number of event exemplars is limited Video2vec outperforms the low-level representation. By increasing training examples the difference becomes more subtle.

2. Attributes. This representation is obtained by adopting the public ImageNet dataset as the source for training attribute classifiers as proposed in [22]. However, instead of training SVM classifiers on bag-of-words encoding of color SIFT descriptors [22], we upgrade the attributes by training a deep CNN with Google Inception architecture on the 15,293 ImageNet categories as detailed in Section 4.2.

Results. Table 3 shows that Video2vec outperforms the state-of-the-art attributes and low-level video representations on all three test sets.

By comparing the Video2vec and the attribute representation we observe a higher event recognition accuracy of 32.4 vs 22.5 for the MED 2013, 17.2 vs 21.8 for the MED 2014, and 54.8 vs 48.8 for the CCV₁₀ test set. We explain it by

the reliance of the attribute baseline on ImageNet categories. Many of these pre-specified categories are not semantically relevant for the events of interest. For example, many of the categories are devoted to specific animal species. In contrast, the Video2vec embedding is automatically derived from the VideoStory46K dataset, which includes many descriptions relevant to events. Different from previous work, our low-level representation outperforms the attribute representation, indicating the strong low-level features used in this work.

The results further demonstrate that the Video2vec outperforms the low-level representation. As an explanation, we speculate that the low-level representation is prone to overfitting due to its high dimensionality. More specifically, in the low-level baseline, the event classifiers are trained on the 20,480 dimensional low-level features from only 10 positive exemplars. This may lead to overfitting as a result of the curse of dimensionality. In contrast, the Video2vec representation transfers the semantics from descriptions into the video representation to alleviate the overfitting as a sort of regularization.

From 1 to 100 examples. For further investigation, we gradually increase the number of positive examples from 1 to 100, and evaluate the accuracy of the event classifiers trained on both the low-level and Video2vec representation. The positive examples are selected randomly² and the results are reported by averaging over 10 repetitions to compensate for the random effect. As Figure 6 shows,

2. The additional event exemplars are selected from the 100Ex evaluation procedure provided in the MED dataset

TABLE 4

Video2vec vs other semantic representations for zero-example event recognition. Our proposed Video2vec outperforms the alternatives.

	Attributes	Term-attributes	Term-attributes-f	Video2vec
MED 2013	3.2	10.7	14.2	15.9
MED 2014	1.4	3.8	5.0	5.2

when the number of event train examples is less than 10, the Video2vec representation outperforms the low-level representation. By increasing the training examples to 100, the difference becomes more subtle, which confirms our hypothesis.

5.1.4 Video2vec for Zero-Example Event Recognition

We also evaluate the effectiveness of Video2vec for zero-example event recognition. We compare Video2vec with the three semantic representation baselines used before: *Attributes*, *Term attributes*, and *Term attributes-f*.

Results. The results in Table 4 demonstrate that Video2vec outperforms the baselines on both test sets. We again explain the modest performance of the attribute baseline by the fact that this representation relies on ImageNet categories as attributes. Many of which are not semantically relevant to the events of interest. This drawback is relaxed by the term attributes and term-attributes-f baselines, as they are trained on more relevant video descriptions from the VideoStory46K dataset. However, both term attribute baselines rely on individual terms for learning the representation. In contrast, Video2vec learns the representation from term combinations, which leads to a better zero-example event recognition accuracy.

5.2 Video2vec Fusion

We evaluate the effectiveness of the Video2vec fusion (Video2vec^F) proposed in Section 3.2 for few-example event recognition. We perform the experiments using the appearance, motion, and audio modalities, as detailed in Section 4.2.

5.2.1 Effect of fusion

We study the impact of fusing multiple modalities for learning the video representation. We start from using only the appearance features (*Appearance*) and gradually add the motion (+ *Motion*) and audio features (+ *Audio*) to learn the video representation by Video2vec^F.

Results. Table 5 confirms that incorporating more modalities for estimating the predictability loss leads to a more effective representation learning.

We observe that some events are benefiting more from adding the auditory modality, i.e., *birthday party* and *music performance*. For these events, there are some distinctive terms which are more effectively predicted from the audio features, i.e., *singing*, *clapping*, and *piano*. Other events improve by including motion features, i.e., *parkour* and *parade*, for which some distinctive terms such as *jumping*, *rolling*, and *running* are well predictable from the motion features. It demonstrates that different modalities are complementary for predicting the terms, so they are all required to effectively predict the video descriptions.

TABLE 5
Effect of fusion. Fusing more modalities to learn Video2vec^F leads to a more effective semantic representation.

MED 2013				MED 2014				CCV ₁₀			
Event	Appearance	+ Motion	+ Audio	Event	Appearance	+ Motion	+ Audio	Event	Appearance	+ Motion	+ Audio
Birthday party	37.1	38.8	43.4	Attempting a bike trick	9.0	9.4	8.8	Basketball	65.6	67.9	68.6
Changing a vehicle tire	64.7	65.5	67.2	Cleaning an appliance	11.1	13.7	16.0	Baseball	58.6	59.6	59.5
Flash mob gathering	55.3	63.8	65.0	Dog show	83.7	88.0	89.7	Soccer	58.7	64.5	64.7
Getting a vehicle unstuck	59.3	65.2	65.3	Giving directions to a location	0.5	1.4	1.0	Ice Skating	70.2	74.5	75.1
Grooming an animal	24.3	29.4	28.2	Marriage proposal	0.3	0.5	0.6	Skiing	79.4	81.8	81.9
Making a sandwich	16.7	18.9	21.3	Renovating a home	11.5	11.9	12.9	Swimming	78.9	81.8	81.7
Parade	33.9	44.7	44.0	Rock climbing	13.8	14.2	13.7	Biking	67.3	69.0	70.2
Parkour	61.3	72.8	72.2	Town hall meeting	40.1	36.3	42.4	Cat	66.2	68.3	70.8
Repairing an appliance	44.5	49.7	57.4	Winning a race without a vehicle	22.1	26.8	26.7	Dog	65.4	68.9	68.6
Working on a sewing project	47.2	48.0	46.0	Working on a metal crafts	15.0	17.9	20.6	Bird	33.9	35.9	36.4
Attempting a bike trick	8.8	9.3	8.9	Beekkeeping	54.1	53.5	48.1	Graduation	25.3	27.6	25.1
Cleaning an appliance	10.5	13.2	15.5	Wedding shower	19.7	31.5	40.3	Birthday	54.2	56.2	61.9
Dog show	81.6	84.9	85.9	Non-motorized vehicle repair	65.3	66.0	66.5	Wedding Reception	17.3	18.3	19.5
Giving directions to a location	0.6	1.0	0.9	Fixing musical instrument	25.4	31.0	44.0	Wedding Ceremony	45.6	53.5	58.4
Marriage proposal	0.3	0.4	0.5	Horse riding competition	40.5	37.9	35.6	Wedding Dance	51.3	58.6	61.7
Renovating a home	11.4	12.1	13.8	Felling a tree	12.4	16.2	22.3	Music Performance	41.0	41.6	51.1
Rock climbing	13.7	14.9	14.4	Parking a vehicle	17.3	19.5	21.9	NonMusic Performance	30.6	35.6	36.4
Town hall meeting	40.2	38.9	44.8	Playing fetch	1.3	1.3	1.4	Parade	49.4	64.0	64.0
Winning a race without a vehicle	21.9	28.8	27.8	Tailgating	32.9	34.5	37.7	Beach	74.5	73.9	72.9
Working on a metal crafts	15.2	18.0	20.4	Tuning musical instrument	8.7	8.3	15.1	Playground	63.4	61.9	58.1
<i>mAP</i>	32.4	35.9	37.1	<i>mAP</i>	24.2	26.0	28.3	<i>mAP</i>	54.8	58.2	59.3

5.2.2 Comparison with other fusions

We compare Video2vec^F with the following fusion baselines:

1. *Early fusion*. This baseline fuses the modalities by simply concatenating the low-level features from all the three modalities into a longer feature vector. Then the event classifiers are trained and applied on the concatenated low-level features. This simple baseline is shown to be competitive to more complex multiple kernel learning [70].

2. *Late fusion*. This baseline fuses the modalities at the level of event classification scores. A separate event classifier is learned by training an SVM on low-level video features per modality. For each test video, the final event detection score is obtained by averaging the detection scores predicted by each classifier, as evaluated in [53].

3. *Video2vec early fusion*. This baseline learns one Video2vec embedding on the concatenation of various low-level video features, using the standard Video2vec objective of Eq. (1). For each video, a semantic representation is extracted by applying the learned feature projection on the concatenation of the low-level video features. Then, the event classifiers are trained and applied on the semantic video representations.

4. *Video2vec late fusion*. This baseline learns three separate Video2vec embeddings, one per modality, using the standard Video2vec objective of Eq. (1). For each video, a semantic representation is obtained by concatenating the three Video2vec representations, which are predicted from each modality. Then, the event classifiers are trained and applied on the semantic video representations.

Results. The results are reported in Table 6. The Video2vec^F outperforms the alternative fusions on all three test sets. We explain the better performance of Video2vec^F over early fusion and late fusion by the fact that both baselines train the event classifiers directly from low-level video features. However, in Video2vec^F the event classifiers are trained on the semantic Video2vec representations, which are more effective than the low-level features in general, as shown

TABLE 6
Comparison with other fusions. The Video2vec^F outperforms the alternative fusion tactics.

	Early fusion	Late fusion	Video2vec early fusion	Video2vec late fusion	Video2vec ^F
MED 2013	32.6	33.8	33.8	33.6	37.1
MED 2014	27.0	27.0	27.0	26.1	28.3
CCV₁₀	56.3	55.8	56.9	55.4	59.3

in Section 5.1.3 for the appearance features.

In the Video2vec late fusion the embeddings are learned *separately* per modality. In contrast, Video2vec^F relies on all the video modalities jointly to estimate the predictability loss, which in general is more reliable than the per modality predictability estimations. For further investigation, we visualize the learned textual projection matrices $\mathbf{A}$ in Figure 7. The Video2vec^F prevents some undesirable combination of terms that happen when the predictability losses are estimated separately per modality. For example, in the Video2vec which is learned only on audio features, the terms *laugh*, *cheer*, *bark*, *dog*, and *woof* are all combined as they have similar auditory features (see the right plot). However, the Video2vec^F does not combine the *laugh* and *cheer* with the *bark*, *dog*, and *woof* terms, as these terms are different in the appearance and motion features. Video2vec^F learns a more reliable combination of terms, leading to a more effective video representation.

Video2vec early fusion learns a single feature projection matrix on the concatenated video features. However, the low-level features from different modalities have a different intrinsic dimensionality, distribution, and meaning, which aggravates the learning from their concatenation. In contrast, the Video2vec^F learns separate feature projection matrices per modality, where each feature projection is optimized based on the features from one modality. It alleviates learning the feature projections, which leads to a more effective video representation.

To conclude, the Video2vec^F effectively fuses the features from various video modalities by embedding them

Fig. 7. Effect of learning the embeddings *jointly* over all the modalities by comparing Video2vec^F with the Video2vec late fusion baseline. Each textual projection matrix is visualized by plotting $\mathbf{A} \times \mathbf{A}^T$, which reveals the learned term combinations. The Video2vec^F learn a more reasonable combination of terms compared to the Video2vec late fusion, where the embeddings are learned disjointedly per modality.

TABLE 7

Effect of learning event specific Video2vec for zero-example event recognition. The event specific Video2vec_ε and Video2vec_ε^F embeddings improve their generic counterparts, Video2vec and Video2vec^F. The results demonstrate that the term sensitive descriptiveness loss (L_d^s) and the multimodal predictability loss (L_p^F) are both effective and complementary.

MED 2013					MED 2014				
Event	Video2vec	Video2vec ^F	Video2vec _ε	Video2vec _ε ^F	Event	Video2vec	Video2vec ^F	Video2vec _ε	Video2vec _ε ^F
Birthday party	24.6	32.5	30.3	37.4	Attempting a bike trick	8.8	5.3	7.1	5.9
Changing a vehicle tire	43.9	36.0	42.1	38.2	Cleaning an appliance	8.2	10.5	9.5	12.3
Flash mob gathering	14.5	30.1	22.4	33.8	Dog show	4.0	5.2	5.8	5.4
Getting a vehicle unstuck	40.2	29.9	36.0	32.2	Giving directions to a location	0.6	0.6	0.6	1.1
Grooming an animal	18.7	21.4	28.0	23.4	Marriage proposal	0.3	0.7	0.4	0.7
Making a sandwich	19.4	15.5	21.1	17.1	Renovating a home	5.2	5.1	6.8	6.3
Parade	17.6	32.5	26.4	38.2	Rock climbing	1.6	1.0	1.4	1.0
Parkour	26.1	34.6	28.0	40.3	Town hall meeting	1.9	2.3	9.2	9.2
Repairing an appliance	39.8	42.4	41.4	46.3	Winning a race without a vehicle	9.4	7.8	8.4	7.6
Working on a sewing project	30.8	34.3	36.4	36.0	Working on a metal crafts	1.5	4.8	4.7	4.8
Attempting a bike trick	8.8	5.1	7.1	5.9	Beekeeping	0.3	1.2	0.6	0.9
Cleaning an appliance	8.2	12.6	9.1	12.5	Wedding shower	2.3	2.8	3.3	3.3
Dog show	4.0	5.8	5.8	5.9	Non-motorized vehicle repair	33.2	39.6	44.2	46.9
Giving directions to a location	0.6	0.7	0.7	0.8	Fixing musical instrument	4.7	18.3	7.9	25.5
Marriage proposal	0.3	0.8	0.4	0.7	Horse riding competition	7.7	12.4	12.8	13.7
Renovating a home	5.2	5.3	6.8	6.4	Felling a tree	7.7	5.7	6.5	6.8
Rock climbing	1.6	1.1	1.4	1.0	Parking a vehicle	2.8	2.3	2.3	2.9
Town hall meeting	1.9	2.5	9.2	9.6	Playing fetch	2.2	3.7	2.9	3.3
Winning a race without a vehicle	9.4	8.1	8.4	7.9	Tailgating	0.2	0.3	0.2	0.4
Working on a metal crafts	1.6	4.1	4.9	7.0	Tuning musical instrument	0.5	2.4	0.9	2.0
mAP	15.9	17.8	18.3	20.0	mAP	5.2	6.6	6.8	8.0

into a mutual semantic representation learned jointly over all the modalities.

5.3 Event Specific Video2vec

We evaluate the impact of learning event specific Video2vec embeddings, as proposed in Section 3.3, by comparing the Video2vec_ε and its multimodal fusion Video2vec_ε^F with Video2vec and Video2vec^F. Different from Video2vec_ε and Video2vec_ε^F, which are learned based on the term sensitive descriptiveness loss, in Eq. (13), the others are learned based on LSI, in Eq. (2).

Results. The results are reported in Table 7 and Table 8. Both event specific embeddings outperform their generic counterparts on both the MED 2013 and MED 2014 test sets. For zero-example event recognition reported in Table 7, the Video2vec_ε improves the Video2vec from 15.9 to 18.3 on MED 2013, and from 5.2 to 6.8 on MED 2014. Similarly, the Video2vec_ε^F improves the Video2vec^F from 17.8 to 20.0 on MED 2013 and from 6.6 to 8.0 on MED 2014. A consistent improvement is observed for few-example event recognition also albeit with a smaller margin, as detailed in Table 8. The results confirm the effectiveness of learning event specific Video2vec embeddings based on the term sensitive descriptiveness loss.

We explain the lower performance of the Video2vec and Video2vec^F baselines by the fact that the LSI loss treats all the terms equally when measuring the reconstruction error. Hence this loss is biased toward minimizing the reconstruction error for the frequent terms. As a consequence, the Video2vec and Video2vec^F are less accurate for predicting infrequent terms. This drawback is addressed by the term sensitive descriptiveness loss, which minimizes the reconstruction error for the indicative terms, even if those terms are infrequent in the train data. As a result, the Video2vec_ε and Video2vec_ε^F can predict the distinctive terms of an event more accurately, which leads to an improved event recognition accuracy. In Figure 8, we compare the term vectors, which are predicted by Video2vec and Video2vec_ε for two video examples.

We explain the bigger gain for zero-example event recognition by two reasons: First, in few-example event recognition, an SVM event classifier is trained on the provided video exemplars, which makes the recognition somewhat robust to the noise in the video representations. For zero-example recognition the event is recognized by directly matching the event definition with the video representation. Hence, accurate prediction of the terms from videos become crucial. As the second explanation, we highlight that in our

Fig. 8. Unseen video examples and their term vectors predicted by the Video2vec (middle) and the Video2vec_ε embeddings (bottom). The size of each term indicates its prediction confidence score. The Video2vec_ε more effectively predicts the indicative terms of the event.

TABLE 8

Effect of learning event specific Video2vec for few-example event recognition. Confirming the conclusion of Table 7.

	Video2vec	Video2vec ^F	Video2vec _ε	Video2vec _ε ^F
MED 2013	32.4	37.1	32.6	37.8
MED 2014	24.2	28.3	24.3	29.1

experiments the term importance matrix H_e is determined based on the event definitions. This is reasonable for zero-example event recognition, where the events are recognized using their definition only. However, for few-example event recognition it might be more effective to learn the term importance matrix H_e from the video examples rather than pre-define them from the event definitions.

Finally, the most effective representation is learned by the Video2vec_ε^F. It indicates that the term sensitive descriptiveness loss (L_d^b) and the multimodal predictability loss (L_p^F) are both effective and complementary when learning semantic video representations for event recognition.

5.4 Comparison with the state-of-the-art

We evaluate the merit of Video2vec by comparing it with several other recent works on few- and zero-example event recognition. Since in most papers the results are reported only on the MED 2013 test set, we limit our comparisons to this test set. Each approach may vary the data sources used for representation pre-training, yet they all abide to the standard data partitioning for training and testing of event classifiers as prescribed by the NIST evaluation protocol.

The results are reported in Table 9. It shows that our proposed representation learning sets a new state-of-the-art for the both few- and zero-example event recognition. It should be noted that these works rely on modeling the temporal aspects of the events [9], [43], [66], [69], using larger train set for representation learning [8], [32], [34], and query expansion [8], [79] to improve the event recognition. These improvements can also be applied together with our representation learning for a more effective event recognition. Very recently Jiang *et al.* [34] improved their results for zero-example recognition from 18.3 to 20.8 after adding re-

TABLE 9

Comparison with the state-of-the-art. Video2vec embeddings set a new state-of-the-art for the both few- and zero-example event recognition.

Few-example on MED 2013	Zero-example on MED 2013	
Habibian <i>et al.</i> [21]	19.6	Ye <i>et al.</i> [83] 9.0
Nagel <i>et al.</i> [51]	21.8	Chang <i>et al.</i> [8] 9.6
Li <i>et al.</i> [43] [†]	23.7	Mazloom <i>et al.</i> [49] 11.9
Tang <i>et al.</i> [69] [†]	26.8	Wu <i>et al.</i> [79] 12.7
Sun <i>et al.</i> [66] [†]	28.7	Jiang <i>et al.</i> [32] 12.9
Chang <i>et al.</i> [9]	30.9	Jiang <i>et al.</i> [34] 18.3
This paper: Video2vec _ε ^F	37.8	This paper: Video2vec _ε ^F 20.0

[†] Based on implementation by [9]

ranking by pseudo relevance feedback, we expect a similar gain for Video2vec_ε^F.

We also perform zero-example event recognition using only the event name as query, rather than the full definition, following the setting of [23], [28]. The results are reported in Table 10. Surprisingly, by using just event names as query, the Video2vec_ε^F performance is increased from 20.0 to 22.5 mAP. The improvement is most significant for the events with informative names *i.e.*, *flash mob gathering* and *parkour*, whose AP increases from 33.8 and 40.3 to 46.8 and 64.2 respectively. For some events such as *repairing an appliance* and *cleaning an appliance*, the event definitions contain specific informative words, *i.e.*, *microwave*, *stove* and *oven*, which are crucial for recognizing the event. Hence, the event definitions are more effective for recognizing these events. Finally, Video2vec_ε^F is more effective than [28] and [23] for zero-example recognition using event names.

6 CONCLUSION

In this paper, we attack the problem of event recognition in video when examples are scarce. We propose the Video2vec embedding that learns a semantic video representation from a set of videos and their textual descriptions by minimizing a joint objective function balancing term descriptiveness and video predictability losses. As a result, the words which are correlated in the descriptions are combined together to improve their video predictability.

In addition, we propose the Video2vec_ε^F embedding with a multimodal predictability loss learned jointly over

TABLE 10
Zero-example event recognition on MED 2013 by querying with event names instead of event definitions.

	Event Definitions	Event Names
Jain <i>et al.</i> [28]	NA	4.2
Han <i>et al.</i> [23]	NA	16.4
This paper: $Video2vec_{\mathcal{E}}^{\mathcal{F}}$	20.0	22.5

video appearance, motion, and audio features. The different modalities are complementary for predicting the words, so they are all required for a richer video description. Moreover, embedding the heterogeneous video features into a mutual semantic space leads to few-example event recognition that is more effective than traditional fusion tactics.

We also propose the $Video2vec_{\mathcal{E}}$ embedding to learn event specific video representations. This embedding relies on a term sensitive descriptiveness loss to learn a more accurate representation for the words, which are indicative of the event.

Finally, the $Video2vec_{\mathcal{E}}^{\mathcal{F}}$ embedding demonstrates that the term sensitive descriptiveness loss and the multimodal predictability loss are both effective and complementary to learn semantic video representations for few- and zero-example event recognition with state-of-the-art accuracy.

We consider $Video2vec$'s ability to generate human interpretable representations for previously unseen videos most appealing, as it opens up new connections with natural language processing and computational linguistics for describing and querying videos.

REFERENCES

- [1] CMU sphinx open source toolkit for speech recognition. <http://cmusphinx.sourceforge.net/>. Accessed: 2015-08-03.
- [2] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid. Label-embedding for attribute-based classification. In *CVPR*, 2013.
- [3] Y. Aytar, M. Shah, and J. Luo. Utilizing semantic word similarity measures for video retrieval. In *CVPR*, 2008.
- [4] T. Berg, A. Berg, and J. Shih. Automatic attribute discovery and characterization from noisy web data. In *ECCV*, 2010.
- [5] S. Bhattacharya, M. M. Kalayeh, R. Sukthankar, and M. Shah. Recognition of complex events: Exploiting temporal dynamics between underlying concepts. In *CVPR*, 2014.
- [6] D. M. Blei and M. I. Jordan. Modeling annotated data. In *SIGIR*, 2003.
- [7] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *ICCS*, 2010.
- [8] X. Chang, Y. Yang, A. G. Hauptmann, E. P. Xing, and Y.-L. Yu. Semantic concept discovery for large-scale zero-shot event detection. In *IJCAI*, 2015.
- [9] X. Chang, Y.-L. Yu, Y. Yang, and A. G. Hauptmann. Searching persuasively: Joint event detection and evidence recounting with limited supervision. In *MM*, 2015.
- [10] J. Chen, Y. Cui, G. Ye, D. Liu, and S.-F. Chang. Event-driven semantic concept discovery by exploiting weakly tagged internet images. In *ICMR*, 2014.
- [11] J. Costa Pereira, E. Coviello, G. Doyle, N. Rasiwasia, G. R. Lanckriet, R. Levy, and N. Vasconcelos. On the role of correlation and abstraction in cross-modal multimedia retrieval. *IEEE TPAMI*, 2014.
- [12] Y. Cui, D. Liu, J. Chen, and S.-F. Chang. Building a large concept bank for representing events in video. *arXiv:1403.7591*, 2014.
- [13] J. Dalton, J. Allan, and P. Mirajkar. Zero-shot video retrieval using content and concepts. In *CIKM*, 2013.
- [14] P. Das, R. Srihari, and J. Corso. Translating related words to videos and back through latent topics. In *WSDM*, 2013.
- [15] P. Das, C. Xu, R. Doell, and J. Corso. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In *CVPR*, 2013.
- [16] A. Dehghan, H. Idrees, and M. Shah. Improving semantic concept detection through the dictionary of visually-distinct elements. In *CVPR*, 2014.
- [17] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. In *CVPR*, 2009.
- [18] F. Feng, X. Wang, and R. Li. Cross-modal retrieval with correspondence autoencoder. In *MM*, 2014.
- [19] V. Ferrari and A. Zisserman. Learning visual attributes. In *NIPS*, 2007.
- [20] A. Habibian, T. Mensink, and C. G. M. Snoek. Composite concept discovery for zero-shot video event detection. In *ICMR*, 2014.
- [21] A. Habibian, T. Mensink, and C. G. M. Snoek. Videostory: A new multimedia embedding for few-example recognition and translation of events. In *MM*, 2014.
- [22] A. Habibian and C. G. M. Snoek. Recommendations for recognizing video events by concept vocabularies. *CVIU*, 124, 2014.
- [23] X. Han, B. Singh, V. I. Morariu, and L. S. Davis. VRFP: On-the-fly video retrieval using web images and fast fisher vector products. *arXiv:1512.03384*, 2016.
- [24] T. Hofmann. Probabilistic latent semantic indexing. In *SIGIR*, 1999.
- [25] H. Hotelling. Relation between two sets of variates. *Biometrika*, 28:322–377, 1936.
- [26] H. Izadinia and M. Shah. Recognizing complex events using large margin joint low-level event model. In *ECCV*, 2012.
- [27] M. Jain, H. Jégou, and P. Bouthemy. Better exploiting motion for better action recognition. In *CVPR*, 2013.
- [28] M. Jain, J. C. van Gemert, T. Mensink, and C. G. M. Snoek. Objects2action: Classifying and localizing actions without any video example. In *ICCV*, 2015.
- [29] H. Jégou, F. Perronnin, M. Douze, J. Sanchez, P. Perez, and C. Schmid. Aggregating local image descriptors into compact codes. *IEEE TPAMI*, 2012.
- [30] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *IEEE TPAMI*, 35(1):221–231, 2013.
- [31] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *MM*, 2014.
- [32] L. Jiang, D. Meng, Q. Zhao, S. Shan, and A. G. Hauptmann. Self-paced curriculum learning. In *AAAI*, 2015.
- [33] L. Jiang, T. Mitamura, S.-I. Yu, and A. G. Hauptmann. Zero-example event search using multimodal pseudo relevance feedback. In *ICMR*, 2014.
- [34] L. Jiang, S.-I. Yu, D. Meng, Y. Yang, T. Mitamura, and A. G. Hauptmann. Fast and accurate content-based semantic search in 100m internet videos. In *MM*, 2015.
- [35] Y.-G. Jiang, S. Bhattacharya, S.-F. Chang, and M. Shah. High-level event recognition in unconstrained videos. *IJMIR*, 2(2):73–101, 2013.
- [36] Y.-G. Jiang, G. Ye, S.-F. Chang, D. Ellis, and A. Loui. Consumer video understanding: A benchmark database and an evaluation of human and machine performance. In *ICMR*, 2011.
- [37] Y.-G. Jiang, X. Zeng, G. Ye, S. Bhattacharya, D. Ellis, M. Shah, and S.-F. Chang. Columbia-UCF TRECCVID 2010 multimedia event detection: Combining multiple modalities, contextual concepts, and temporal matching. In *TRECCVID*, 2010.
- [38] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.
- [39] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.
- [40] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [41] C. H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 2009.
- [42] I. Laptev, M. Marszałek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *CVPR*, 2008.
- [43] W. Li, Q. Yu, A. Divakaran, and N. Vasconcelos. Dynamic pooling for complex event recognition. In *ICCV*, 2013.
- [44] Z. Ma. *From concepts to events: A progressive process for multimedia content analysis*. PhD thesis, University of Trento, 2013.
- [45] Z. Ma, Y. Yang, N. Sebe, and A. G. Hauptmann. Knowledge adaptation with partially shared features for event detection using few exemplars. *IEEE TPAMI*, 2014.

- [46] Z. Ma, Y. Yang, Z. Xu, S. Yan, N. Sebe, and A. Hauptmann. Complex event detection via multi-source video attributes. In *CVPR*, 2013.
- [47] J. Mairal, F. Bach, and J. Ponce. Sparse modeling for image and vision processing. *FTCV*, 8(2-3):85–283, 2012.
- [48] M. Mazloom, E. Gavves, and C. G. M. Snoek. Conceptlets: Selective semantics for classifying video events. *TMM*, 16, 2014.
- [49] M. Mazloom, A. Habibian, D. Liu, C. G. M. Snoek, and S.-F. Chang. Encoding concept prototypes for video event detection and summarization. In *ICMR*, 2015.
- [50] M. Merler, B. Huang, L. Xie, G. Hua, and A. Natsev. Semantic model vectors for complex video event recognition. *TMM*, 2012.
- [51] M. Nagel, T. Mensink, and C. G. M. Snoek. Event fisher vectors: Robust encoding visual diversity of visual streams. In *BMVC*, 2015.
- [52] P. Natarajan et al. BBN VISER TRECVID 2013 multimedia event detection and multimedia event recounting systems. In *TRECVID*, 2013.
- [53] P. Natarajan, S. Wu, S. Vitaladevuni, X. Zhuang, S. Tsakalidis, U. Park, R. Prasad, and P. Natarajan. Multimodal feature fusion for robust event detection in web videos. In *CVPR*, 2012.
- [54] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng. Multimodal deep learning. In *ICML*, 2011.
- [55] S. Oh, S. McCloskey, I. Kim, A. Vahdat, K. J. Cannons, H. Hamjirsadeghi, G. Mori, A. A. Perera, M. Pandey, and J. J. Corso. Multimedia event detection with multimodal feature fusion and temporal concept localization. *MVA*, 2014.
- [56] D. Oneata, J. Verbeek, and C. Schmid. Action and event recognition with fisher vectors on a compact feature set. In *ICCV*, 2013.
- [57] P. Over, J. Fiscus, G. Sanders, et al. TRECVID 2013—an overview of the goals, tasks, data, evaluation mechanisms and metrics. In *TRECVID*, 2013.
- [58] M. Palatucci, D. Pomerleau, G. E. Hinton, and T. M. Mitchell. Zero-shot learning with semantic output codes. In *NIPS*, 2009.
- [59] D. Parikh and K. Grauman. Relative attributes. In *ICCV*, 2011.
- [60] D. Putthividhy, H. T. Attias, and S. S. Nagarajan. Topic regression multi-modal latent dirichlet allocation for image annotation. In *CVPR*, 2010.
- [61] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 2015.
- [62] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek. Image classification with the fisher vector: Theory and practice. *IJCV*, 105, 2013.
- [63] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [64] C. G. M. Snoek, M. Worring, and A. W. M. Smeulders. Early versus late fusion in semantic video analysis. In *MM*, 2005.
- [65] S. Strassel et al. Creating havic: Heterogeneous audio visual internet collection. In *LREC*, 2012.
- [66] C. Sun and R. Nevatia. Discover: Discovering important segments for classification of video events and recounting. In *CVPR*, 2014.
- [67] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [68] A. Tamrakar, S. Ali, Q. Yu, J. Liu, O. Javed, A. Divakaran, H. Cheng, and H. Sawhney. Evaluation of low-level features and their combinations for complex event detection in open source videos. In *CVPR*, 2012.
- [69] K. Tang, L. Fei-Fei, and D. Koller. Learning latent temporal structure for complex event detection. In *CVPR*, 2012.
- [70] K. Tang, B. Yao, L. Fei-Fei, and D. Koller. Combining the right features for complex event recognition. In *ICCV*, 2013.
- [71] G. W. Taylor, R. Fergus, Y. LeCun, and C. Bregler. Convolutional learning of spatio-temporal features. In *ECCV*, 2010.
- [72] I. Tsochanaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. In *JMLR*, pages 1453–1484, 2005.
- [73] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko. Translating videos to natural language using deep recurrent neural networks. In *NAACL*, 2015.
- [74] H. Wang, D. Oneata, J. Verbeek, and C. Schmid. A robust and efficient video representation for action recognition. *IJCV*, 2015.
- [75] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013.
- [76] Q. Wang, J. Xu, H. Li, and N. Craswell. Regularized latent semantic indexing. In *SIGIR*, 2011.
- [77] K. Q. Weinberger and O. Chapelle. Large margin taxonomy embedding for document categorization. In *NIPS*, 2009.
- [78] J. Weston, S. Bengio, and N. Usunier. Wsabi: Scaling up to large vocabulary image annotation. In *IJCAI*, 2011.
- [79] S. Wu, S. Bondugula, F. Luisier, X. Zhuang, and P. Natarajan. Zero-shot event detection using multi-modal fusion of weakly supervised concepts. In *CVPR*, 2014.
- [80] R. Xu, C. Xiong, W. Chen, and J. J. Corso. Jointly modeling deep video and compositional text to bridge vision and language in a unified framework. In *AAAI*, 2015.
- [81] Z. Xu, Y. Yang, and A. G. Hauptmann. A discriminative cnn video representation for event detection. In *CVPR*, 2015.
- [82] Y. Yan, Y. Yang, H. Shen, D. Meng, G. Liu, A. Hauptmann, and N. Sebe. Complex event detection via event oriented dictionary learning. In *AAAI*, 2015.
- [83] G. Ye, Y. Li, H. Xu, D. Liu, and S.-F. Chang. Eventnet: A large scale structured concept library for complex event detection in video. In *MM*, 2015.
- [84] S.-I. Yu, L. Jiang, Z. Mao, X. Chang, X. Du, C. Gan, Z. Lan, Z. Xu, X. Li, Y. Cai, et al. Informedia@ TRECVID 2014 MED and MER. In *TRECVID*, 2014.

Amirhossein Habibian received a B.Sc. in computer engineering (2008) and a M.Sc. in artificial intelligence (2011), both from University of Tehran, Iran. Currently, he is a Ph.D. candidate at the University of Amsterdam, The Netherlands. He is also a senior engineer at Qualcomm Research in Amsterdam working on deep learning and computer vision. He received the Best Paper Award of ACM Multimedia 2014. His research interests include machine learning, computer vision, and multimedia retrieval.

Thomas Mensink received a PhD degree in computer science in 2012 from the University of Grenoble, France, while working both at the LEAR team of INRIA Grenoble and at Xerox Research Centre Europe. He has received a cum laude MSc degree in Artificial Intelligence from the University of Amsterdam, The Netherlands, in 2007. Since 2012 he is a PostDoc researcher at the University of Amsterdam, and in 2015 he received a prestigious NWO VENI grant. His research interests include machine learning and computer vision, with special interest in learning semantic representations for zero-example prediction.

Cees G.M. Snoek received the M.Sc. degree in business information systems (2000) and the Ph.D. degree in computer science (2005) both from the University of Amsterdam, The Netherlands. He is currently a director of the QUVA Lab, the joint research lab of Qualcomm and the University of Amsterdam on deep learning and computer vision. He is also a principal engineer at Qualcomm and an associate professor at the University of Amsterdam. His research interests focus on video and image recognition. Cees is recipient of an NWO Veni award (2008), a Fulbright Junior Scholarship (2010), an NWO Vidi award (2012), and the Netherlands Prize for ICT Research (2012). Several of his Ph.D. students and Post-docs have won awards, including the IEEE Transactions on Multimedia Prize Paper Award (2012) the ACM SIGMM Best Ph.D. Thesis Award (2013), the Best Paper Award of ACM Multimedia (2014) and an NWO Veni award (2015). Three of his former mentees serve as assistant professor.