

Zero-Shot Event Detection by Multimodal Distributional Semantic Embedding of Videos

Mohamed Elhoseiny^{§‡}, Jingen Liu[‡], Hui Cheng[‡], Harpreet Sawhney[‡], Ahmed Elgammal[§]
 m.elhoseiny@cs.rutgers.edu, {jingen.liu, hui.cheng, harpreet.sawhney}@sri.com, elgammal@cs.rutgers.edu
[§]Rutgers University, Computer Science Department [‡]SRI International, Vision and Learning Group


Abstract

We propose a new zero-shot Event Detection method by Multi-modal Distributional Semantic embedding of videos. Our model embeds object and action concepts as well as other available modalities from videos into a distributional semantic space. To our knowledge, this is the first Zero-Shot event detection model that is built on top of distributional semantics and extends it in the following directions: (a) semantic embedding of multimodal information in videos (with focus on the visual modalities), (b) automatically determining relevance of concepts/attributes to a free text query, which could be useful for other applications, and (c) retrieving videos by free text event query (e.g., "changing a vehicle tire") based on their content. We embed videos into a distributional semantic space and then measure the similarity between videos and the event query in a free text form. We validated our method on the large TRECVID MED (Multimedia Event Detection) challenge. Using only the event title as a query, our method outperformed the state-of-the-art that uses big descriptions from 12.6% to 13.5% with MAP metric and 0.73 to 0.83 with ROC-AUC metric. It is also an order of magnitude faster.

Introduction

Every minute, hundreds of hours of video are uploaded to video archival site such as YouTube (Google2014). Developing methods to automatically understand the events captured in this large volume of videos is necessary and meanwhile challenging. One of the important tasks in this direction is event detection in videos. The main objective of this task is to determine the relevance of a video to an event based on the video content (e.g., feeding an animal, birthday party; see Fig. 1). The cues of an event in a video could include visual objects, scene, actions, detected speech (by Automated Speech Recognition(ASR)), detected text (by Optical Character Recognition (OCR)), and audio concepts (e.g. music and water concepts).

Search and retrieval of videos for arbitrary events using only free-style text and unseen text in particular has been a dream in computational video and multi-media understanding. This is referred as "zero-shot event detection", because there is no positive exemplar videos to train a detector. Due to the proliferation of videos, especially consumer-generated



(a) Grooming an Animal

- (1) "brushing dog", weight = 0.67
- (2) "combing dog", weight = 0.66
- (3) "clipping nails", weight = 0.52



(b) Birthday Party

- (1) "cutting cake", weight = 0.72
- (2) "blowing candles", weight = 0.65
- (3) "opening presents", weight = 0.59

Figure 1: Top relevant Concepts from a pre-defined multimedia concept repository and their automatically-assigned weights as a part of our Event Detection method

videos (e.g., YouTube), zero-shot search and retrieval of videos has become an increasingly important problem.

Several research works have been proposed to facilitate performing the zero-shot learning task by establishing an intermediate semantic layer between events or generally categories (i.e., concepts or attributes) and the low-level representation of a multimedia content from the visual perspective. (Lampert, Nickisch, and Harmeling 2009) and (Farhadi et al. 2009) were the two first to use attribute learning representation for the zero-shot setting for object recognition in still images. Attributes were similarly adopted for recognizing human actions (Liu, Kuipers, and Savarese 2011); attributes are generalized and denoted by concepts in this context. Later, (Liu et al. 2013) proposed Concept Based Event Retrieval (CBER) for videos InTheWild. Even though these methods facilitate zero-shot event detection, they only capture the visual modality and more importantly they assume that the relevant concepts for a query event are manually defined. This manual definition of concepts, also known as semantic query editing, is a tedious task and may be biased due to the limitation of human knowledge. Instead, we aim at automatically generating relevant concepts by leveraging information from distributional semantics.

Recently, several systems were proposed for zero-shot event detection methods (Wu et al. 2014; Jiang et al. 2014b; 2014a; Chen et al. 2014; Habibian, Mensink, and Snoek

2014). These approaches rely on the whole text description of an event where relevant concepts are specified; see example event descriptions used in these approaches in the Supplementary Materials (SM)¹ (explicitly define the event explanation, scene, objects, activities, and audio). In practice, however, we think that typical use of event queries under this setting should be similar to text-search, which is based on few words instead that we model their connection to the multimodal content in videos.

The main question addressed in this paper is how to use an event text query (i.e. just the event title like “birthday party” or “feeding an animal”) to retrieve a ranked list of videos based on their content. In contrast to (Lampert, Nickisch, and Harmeling 2009; Liu et al. 2013), we do not manually assign relevant concepts for a given event query. Instead, we leverage information from a distributional semantic space (Mikolov et al. 2013b) trained on a large text corpus to embed event queries and videos to the same space, where similarity between both could be directly estimated. Furthermore, we only assume that query comes in the form of an “unstructured” few-keyword query (in contrast to (Wu et al. 2014; Jiang et al. 2014b; 2014a)). We abbreviate our method as EDiSE (Event-detection by Distributional Semantic Embedding of videos).

Contributions. The contributions of this paper can be listed as follows: (1) Studying how to use few-keyword unstructured-text query to detect/retrieve videos based on their multimedia content, which is novel in this setting. We show how relevant concepts to that event query could be automatically retrieved through a distributional semantic space and got assigned a weight associated with the relevance; see Fig. 1 “Birthday” and “Grooming an Animal” example events. (2) To the best of our knowledge, our work is the first attempt to model the connection between few keywords and multimodal information in videos by distributional semantics. We study and propose different similarity metrics in the distributional semantic space to enable event retrieval based on (a) concepts, (b) ASR, and (c) OCR in videos. Our unified framework is capable of embedding all of them into the same space; see Fig. 2. (3) Our method is also very fast, which makes it applicable to both large number of videos and concepts (*i.e.* 26.67 times faster than the state of the art (Jiang et al. 2014a)).

Related Work

Attribute methods for zero-shot learning are based on manually specifying the attributes for each category (e.g., (Lampert, Nickisch, and Harmeling 2009; Parikh and Grauman 2011)). Other methods focused on attribute discovery (Rohrbach, Stark, and Schiele 2011; Rohrbach, Ebert, and Schiele 2013) and then apply the same mechanism. Recently, several methods were proposed to perform zero shot recognition by representing unstructured text in document terms (*e.g.* (Elhoseiny, Saleh, and Elgammal 2013; Mensink, Gavves, and Snoek 2014)) One drawback of the TFIDF (Salton and Buckley 1988) in (Elho-

seiny, Saleh, and Elgammal 2013) and hardly matching tag terms in (Mensink, Gavves, and Snoek 2014; Rohrbach et al. 2010) is that they do not capture semantically related terms that our model can relate in noisy videos instead of still images. Also, WordNet (Miller 1995), adopted in (Rohrbach et al. 2010), does not connect objects with actions (e.g., person blowing candle), making it hard to apply in our setting and heavily depending on predefined information like WordNet.

There has been a recent interest especially in the computational linguistics’ community in word-vector representation (e.g., (Bengio et al. 2006)), which captures word semantics based on context. While word-vector representation is not new, recent algorithms (e.g. (Mikolov et al. 2013b; 2013a)) enabled learning these vectors from billions of words, which makes them much more semantically accurate. As a result, these models got recently adopted in several tasks including translation (Mikolov, Le, and Sutskever 2013) and web search (Shen et al. 2014). Several computer vision researchers explored using these word-vector representation to perform Zero-Shot learning in the object recognition (e.g. (Frome et al. 2013; Socher et al. 2013; Norouzi et al. 2014)). They embed the object class name into the word-vector semantic space learnt by models like (Mikolov et al. 2013b). It is worth mentioning that these zero-shot learning approaches (Frome et al. 2013; Socher et al. 2013) and also the aforementioned work (Elhoseiny, Saleh, and Elgammal 2013) assume that during training, there is a set of training classes and test classes. Hence, they learn a transformation to correlate the information between both domains (textual and visual). In contrast, zero-shot setting of event retrieval rely mainly on the event information without seeing any training events, as assumed in recent zero-shot event retrieval methods (e.g., (Dalton, Allan, and Mirajkar 2013; Jiang et al. 2014b; Wu et al. 2014; Liu et al. 2013)). Hence, there does not exist seen events to learn such transformation from. Differently, we also model multimodal connection from free text query to video information.

In the context of videos, (Wu et al. 2014) proposed a method for zero-shot event detection by using the salient words in the whole structured event description, where relevant concept are already defined in the event structured text description; also see Eq. 1 in (Wu et al. 2014). Similarly, (Dalton, Allan, and Mirajkar 2013) adopted a Markov-Random-Field language model proposed by (Metzler and Croft 2005). One drawback of this model is that it performs an intensive processing for each new concept. This is since it determines the relevance of the concept to a query event by creating a text document to represent each concept. This document is created by web-querying the concept name and some of its keywords and merging the top retrieved pages. In contrast, our model does not require this step to determine relevance of an event to a query. Once the language model is trained, any concept can be instantly added and captured in our multimodal semantic embedding of videos.

In contrast to both (Wu et al. 2014) and (Dalton, Allan, and Mirajkar 2013), we focus on retrieving videos only with the event title (*i.e.*, few-words query) and without semantic

¹Supplementary Materials (SM) could be found here https://sites.google.com/site/mhelseiny/EDiSE_supp.zip

editing. The key difference is in modeling and embedding concepts to allow zero-shot event retrieval. In (Wu et al. 2014) and (Dalton, Allan, and Mirajkar 2013), the semantic space is a vector whose dimensionality is the number of the concepts. Our idea is to embed concepts, video information, and the event query into a distributional semantic space whose dimensionality is independent of the number of concepts. This property, together with the semantic properties captured by distributional semantics, feature our approach with two advantages (a) scalability to any concept size. Having new concepts does not affect the representation dimensionality (i.e., in all our experiments concepts, videos, event queries are embedded to M dimensional space; M is few hundreds in our experiments). (b) facilitating automatic determination of relevant concepts given an unstructured short event query: For example, being able to automatically determine that “blowing a candle” concept is a relevant concept to “birthday party” event. (Wu et al. 2014) and (Dalton, Allan, and Mirajkar 2013) used the complete text description of an event for retrieval that explicitly specifies relevant concepts.

There is a class of models that improve zero-shot Event Detection performance by reranking. Jiang et al. proposed multimodal pseudo relevance feedback (Jiang et al. 2014b) and self-paced reranking (Jiang et al. 2014a) algorithms. The main assumption behind these models is that all unlabeled test examples are available and the top few examples by a given initial ranking have high top K precision ($K \sim 10$). This means that reranking algorithms can not update confidence of a video for an event without knowing the confidences of the remaining videos to perform reranking. In contrast, our goal is different which is to directly model the probability of a few-keyword event-query given an arbitrary video. Hence, our work does not require an initial ranking and can compute the conditional probability of a video without any information about other videos. Our method is also 26.67 times faster, as detailed in our experiments.

Method

Problem Definition

Given an arbitrary event query e and a video v (e.g. just “birthday party”), our objective is to model $p(e|v)$. We start by defining the representation of event query e , the concept set \mathbf{c} , the video v in our setting.

Event-Query representation e : We use the unstructured event title to represent an event query for concept based retrieval. Our framework also allows additional terms specifically for ASR or OCR based retrieval. While we show retrieval on different modalities, concept based retrieval is our main focus in this work. The few-keyword event query for concept based retrieval is denoted by e_c , while query keywords for OCR and ASR are denoted by e_o and e_a , respectively. Hence, under our setting $e = \{e_c, e_o, e_a\}$.

Concept Set \mathbf{c} : We denote the whole concept set in our setting as \mathbf{c} , which include visual concepts \mathbf{c}_v and audio concepts \mathbf{c}_d , i.e., $\mathbf{c} = \{\mathbf{c}_v, \mathbf{c}_d\}$. The visual concepts include object, scene and action concepts. The audio concepts include acoustic related concepts like water sound. We performed an experiment on a set of audio concepts trained

on MFCC audio features (Davis and Mermelstein 1980; Logan and others 2000). However, we found their performance $\approx 1\%$ MAP, and hence we excluded them from our final experiments. Accordingly, our final performance mainly relies on the visual concepts for concept based retrieval; i.e., $\mathbf{c}_d = \emptyset$. We denote each member $c_i \in \mathbf{c}$ as the definition of the i^{th} concept in \mathbf{c} . c_i is defined by the i^{th} concept’s name and optionally some related keywords; see examples in SM. Hence, $\mathbf{c} = \{c_1, \dots, c_N\}$ is the the set of concept definitions, where N is the number concepts.

Video Representation: For our zero-shot purpose, a video v is defined by three pieces of information, which are video OCR denoted by v_o , video ASR denoted by v_a , and video concept representation denoted by v_c . v_o and v_a are the detected text in OCR and ASR, respectively. We used (Myers et al. 2005) to extract v_o and (van Hout et al. 2013) to extract v_a . In this paper, we mainly focus on the visual video content, which is the most challenging. The video concept based representation v_c is defined as

$$v_c = [p(c_1|v), p(c_2|v), \dots, p(c_N|v)] \quad (1)$$

where $p(c_i|v)$ is a conditional probability of concept c_i given video v , detailed later. We denote $p(c_i|v)$ by v_c^i .

In zero-shot event detection setting, we aim at recognizing events in videos without training examples based on its multimedia content including still-image concepts like objects and scenes, action concepts, OCR, and ASR². Given a video $v = \{v_c, v_o, v_a\}$, our goal is to compute $p(e|v)$ by embedding both the event query e and information of video v of different modalities (v_c , v_o , and v_a) into a distributional semantic space, where relevance of v to e could be directly computed; see Fig. 2. Specifically, our approach is to model $p(e|v)$ as a function \mathcal{F} of $\theta(e)$, $\psi(v_c)$, $\theta(v_o)$, and $\theta(v_a)$, which are the distributional semantic embedding of e , v_c , v_o , and v_a , respectively

$$p(e|v) \propto \mathcal{F}(\theta(e), \psi(v_c), \theta(v_o), \theta(v_a)) \quad (2)$$

We remove the stop words from e , v_o , v_a before applying the embedding $\theta(\cdot)$. The rest of this section is organized as follows. First, we present the distributional semantic manifold and the embedding function $\theta(\cdot)$ which is applied to e , v_a , v_o , and the concept definitions \mathbf{c} in our framework. Then, we show how to determine automatically relevant concepts to an event title query and assign a relevance weight to them, as illustrated in Fig. 1. We present this concept relevance weighting in a separate section since it might be generally useful for other applications. Finally, we present the details of $p(e|v)$ where we derive v_c embedding (i.e. $\psi(v_c)$), which is based on the proposed concept relevance weighting.

Distributional Semantic Model & $\theta(\cdot)$ Embedding

We start by the distributional semantic model by (Mikolov et al. 2013b; 2013a) to train our semantic manifold. We denote the trained semantic manifold by \mathcal{M}_s , and the vectorization function that maps a word to \mathcal{M}_s space as $vec(\cdot)$. We denote the dimensionality of the real vector returned from

²Note that OCR and ASR are not concepts. They are rather detected text in video frames and speech

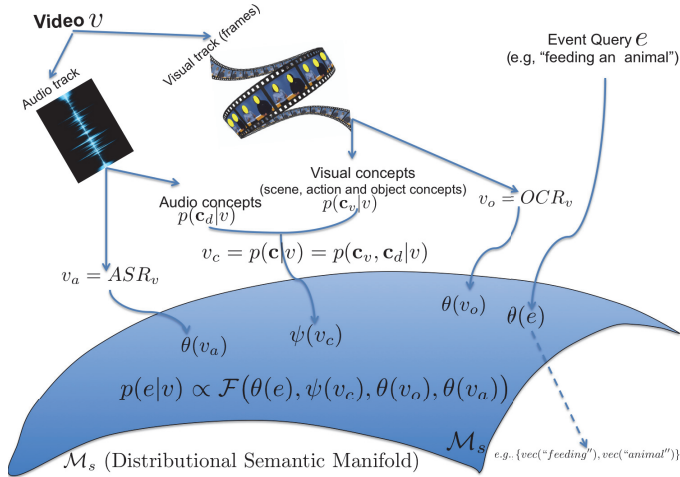


Figure 2: EDiSE Approach

$vec(\cdot)$ by M . These models learn a vector for each word w_n , such that $p(w_n | (w_{i-L}, w_{i-L+1}, \dots, w_{i+L-1}, w_{i+L}))$ is maximized over the training corpus; $2 \times L$ is the context window size. Hence similarity between $vec(w_i)$ and $vec(w_j)$ is high if they co-occurred a lot in context of size $2 \times L$ in the training text-corpus (i.e., semantically similar words share similar context). Based on the trained \mathcal{M}_s space, we define how to embed the event query e , and c . Each of e_c , e_a , and e_o is set of one or more words. Each of these words can be directly embedded into \mathcal{M}_s manifold by $vec(\cdot)$ function. Accordingly, we represent these sets of word vectors for each of e_c , e_a , and e_o as $\theta(e_c)$, $\theta(e_a)$, and $\theta(e_o)$. We denote $\{\theta(e_c), \theta(e_a), \theta(e_o)\}$ by $\theta(e)$. Regarding embedding of c , each concept $c^* \in c$ is defined by its name and optionally some related keywords. Hence, the corresponding word vectors are then used to define $\theta(c^*)$ in \mathcal{M}_s space.

Relevance of Concepts to Event Query

Let us define a similarity function between $\theta(c^*)$ and $\theta(e_c)$ as $s(\theta(e_c), \theta(c^*))$. We propose two functions to measure the similarity between $\theta(e_c)$ and $\theta(c^*)$. The first one is inspired by an example by (Mikolov et al. 2013b) to show the quality of their language model, where they indicated that $vec(\text{"king"}) + vec(\text{"woman"}) - vec(\text{"man"})$ is closest to $vec(\text{"queen"})$. Accordingly, we define a version of $s(X, Y)$, where the sets X and Y are firstly pooled by the sum operation; we denote the sum pooling operation on a set by an overline. For instance, $\overline{X} = \sum_i x_i$ and $\overline{Y} = \sum_i y_j$, where x_i and y_j are the word vectors of the i^{th} element in X and the j^{th} element in Y , respectively. Then, cosine similarity between \overline{X} and \overline{Y} is computed. We denote this version as $s_p(\cdot, \cdot)$; see Eq. 3. Fig. 3 shows how $s_p(\cdot, \cdot)$ could be used to retrieve the top 20 concepts relevant to $\theta(\text{"Grooming An Animal"})$ in \mathcal{M}_s space. The figure also shows embedding of the query and the relevant concept sets in 3D PCA visualization. $\theta(e_c = \text{"Grooming An Animal"})$ and each of $\theta(c_i)$ for the most relevant 20 concepts are represented by their



Figure 3: PCA visualization in 3D of the "Grooming an Animal" event (in green) and its most 20 relevant concepts in \mathcal{M}_s space using $s_p(\cdot, \cdot)$. The exact $s_p(\theta(\text{"Grooming An Animal"}), \theta(c_i))$ is shown between parenthesis

corresponding pooled vectors $(\overline{\theta(e_c)} \text{ and } \overline{\theta(c_i)}) \forall i$, normalized to unit length under L2 norm. Another idea is to define $s(X, Y)$ as a similarity function between the X and Y sets. For robustness (Torki and Elgammal 2010), we used percentile-based Hausdorff point set metric, where similarity between each pair of points is computed by the cosine similarity. We denote this version by $s_t(X, Y)$; see Eq. 3. We used $l = 50\%$ (i.e., median).

$$s_p(X, Y) = \frac{(\sum_i x_i)^T (\sum_j y_j)}{\|\sum_i x_i\| \|\sum_j y_j\|} = \frac{\overline{X}^T \overline{Y}}{\|\overline{X}\| \|\overline{Y}\|} \quad (3)$$

$$s_t(X, Y) = \min_j \{ \max_i \frac{x_i^T y_j}{\|x_i\| \|y_j\|}, \min_i \{ \max_j \frac{x_i^T y_j}{\|x_i\| \|y_j\|} \}$$

Event Detection $p(e|v)$

In practice, we decomposed $p(e|v)$ into $p(e_c|v)$, $p(e_a|v)$, $p(e_o|v)$, which makes the problem reduces to deriving $p(e_c|v)$ (concept based retrieval), $p(e_a|v)$ (OCR based retrieval), and $p(e_o|v)$ (ASR based retrieval) under \mathcal{M}_s . We start by $p(e_c|v)$ then we will how later in this section how $p(e_a|v)$, and $p(e_o|v)$ could be estimated.

Estimating $p(e_c|v)$: In our work, concepts are linguistic meanings that have corresponding detection functions given the video v . From Fig. 3, \mathcal{M}_s space could be viewed as a space of meanings captured by a training text-corpus, where only sparse points in that space has a corresponding visual detection functions given v , which are the concepts c (e.g., "blowing a candle"). For zero shot event detection, we aim at exploiting these sparse points by the information captured by $s(\theta(e_c), \theta(c^* \in c))$ in \mathcal{M}_s space. We derive $p(e_c|v)$ from probabilistic perspective starting from marginalizing $p(e_c|v)$ over the concept set c

$$p(e_c|v) \propto \sum_{c_i} p(e_c|c_i) p(c_i|v) \propto \sum_{c_i} s(\theta(e_c), \theta(c_i)) v_{c_i}^i \quad (4)$$

where $p(e_c|c_i) \forall i$ are assumed to be proportional to $s(\theta(e_c), \theta(c_i))$ in our framework. From semantic embedding perspective, each video v is embedded into \mathcal{M}_s by the

set $\psi(v_c) = \{\theta_v(c_i) = v_c^i \theta(c_i), \forall c_i \in \mathbf{c}\}$, where $v_c^i \theta(c_i)$ is a set of the same points in $\theta(c_i)$ scaled with v_c^i ; $\psi(v_c)$ could be then directly compared with $\theta(e_c)$; see Eq. 5

$$p(e_c|v) \propto \sum_{c_i} s(\theta(e_c), \theta(c_i)) v_c^i \quad (5)$$

$$\propto s'(\theta(e_c), \psi(v_c) = \{\theta_v(c_i), \forall c_i \in \mathbf{c}\})$$

where $s'(\theta(e_c), \psi(v_c)) = \sum_i s(\theta(e_c), \theta_v(c_i))$ and $s(\cdot, \cdot)$ could be replaced by $s_p(\cdot, \cdot)$, $s_t(\cdot, \cdot)$, or any other measure in \mathcal{M}_s space. An interesting observation is that when $s_p(\cdot, \cdot)$ is chosen, $p(e_c|v) \propto \frac{\theta(e_c)^T}{\|\theta(e_c)\|} \left(\sum_i \frac{\theta(c_i)}{\|\theta(c_i)\|} v_c^i \right)$ which is a direct similarity between $\frac{\theta(e_c)}{\|\theta(e_c)\|}$ representing the query and the embedding of $\psi(v_c)$ as $\sum_i \frac{\theta(c_i)}{\|\theta(c_i)\|} v_c^i$; see proof in Appendix A. $s_p(\cdot, \cdot)$ performs consistently better than $s_t(\cdot, \cdot)$ in our experiments. In practice, we only include $\theta_v(c_i)$ in $\psi(v_c)$ such that c_i is among the top R concepts with highest $p(e_c|c_i)$. This is assuming that the remaining concepts are assigned $p(e_c|c_i) = 0$ which makes those items vanish; we used R=5. Hence, only a few concept detectors needs to be computed for on v which is a computational advantage.

Estimating $p(e_o|v)$ and $p(e_a|v)$: Both v_o and v_a can be directly embedded into \mathcal{M}_s since they are sets of words. Hence, we can model $p(e_o|v)$ and $p(e_a|v)$ as follows

$$p(e_o|v) \propto s_d(\theta(e_o), \theta(v_o)), p(e_a|v) \propto s_d(\theta(e_a), \theta(v_a)) \quad (6)$$

where $s_d(X, Y) = \sum_{ij} x_i^T y_j$. We found this similarity function more appropriate for ASR/OCR text since they normally contains more text compared to concept definition. We also exploited an interesting property in \mathcal{M}_s that nearest words to an arbitrary point can be retrieved. Hence, we automatically augment e_a and e_o with the nearest words to the event title in \mathcal{M}_s using cosine similarity before retrieval. We found this trick effective in practice since it automatically retrieve relevant words that might appear in v_o or v_a .

Fusion: We fuse $p(e_c|v)$, $p(e_o|v)$, and $p(e_a|v)$ by weighted geometric mean with focus on visual concepts, i.e. $p(e|v) = \sqrt[w+1]{p(e_c|v)^w \sqrt{p(e_o|v)p(e_a|v)}}$; $w = 6$. $p(e_c|v)$, $p(e_o|v)$, and $p(e_a|v)$ involves the similarity between $\theta(e)$ and each of $\psi(v_c)$, $\theta(v_o)$, and $\theta(v_a)$, leading to Eq. 2 view.

Visual Concept Detection functions ($p(\mathbf{c}|v)$)

We leverage the information from three types of visual concepts in \mathbf{c}_v : object concepts \mathbf{c}_o , action concepts \mathbf{c}_a , and scene concepts \mathbf{c}_s . Hence, $\mathbf{c} = \mathbf{c}_v = \{\mathbf{c}_o \cup \mathbf{c}_a \cup \mathbf{c}_s\}$; the list of concepts are attached in SM. We define object and scene concept probabilities per video frame, and action concepts per video chunks. The rest of this section summarizes the concept detection for objects and scenes per frame f , and action concepts per video chunk u . Then, we show how they can be reduced to video level probabilities. Fig. 4 shows example high confidence concepts in a ‘‘Birthday Party’’ video. **Object Concepts $p(o_i|f)$, $o_i \in \mathbf{c}_o$:** We involve 1000 Overfeat (Sermanet et al. 2014) object concept detectors which maps to 1000-ImageNet categories. We also adopt the concept detectors of face, car and person from a publicity available detector (i.e., (Felzenszwalb, McAllester, and Ramanan 2008))

Scene Concepts $p(s_i|f)$, $s_i \in \mathbf{c}_s$: We represented scene concepts ($p(s_i|f)$) as bag of word representation on static features (i.e., SIFT (Lowe 2004) and HOG (Dalal and Triggs 2005)) with 10000 codebooks. We used TRECVID 500 SIN concepts concepts, including scene categories like ‘‘city’’ and ‘‘hall’’ way; these concepts are provided by provided by TRECVID2011 SIN track.

Action Concepts $p(a_i|u)$, $a_i \in \mathbf{c}_a$: We use both manually annotated (i.e. strongly supervised) and automatically annotated (i.e. weakly supervised) concepts; detailed in SM. We have ~500 action concepts; please refer to (Liu et al. 2013) for the action concept detection method that we adopt.

Video level concept probabilities $p(\mathbf{c}|v)$ We represent probabilities of the \mathbf{c}_v set given a video v by a pooling operation over the the chunks or the frames of the videos similar to (Liu et al. 2013). In our experiments, we evaluated both max and average pooling. Specifically, $p(o_i|v) = \rho(\{p(o_i|f_k), f_k \in v\})$, $p(s_i|v) = \rho(\{p(s_i|f_k), f_k \in v\})$, $p(a_k|v) = \rho(\{p(a_i|u_k), u_k \in v\})$, where $p(o_i|v)$ and $p(s_i|v)$ are the video level probabilities of for the i^{th} object and the l^{th} scene concepts respectively, pooled over frames $f_k \in v$. $\{f_k \in v\}$ are selected every M frames in v (M= 250). $p(a_k|v)$ is the video level probability of the k^{th} action concept, pooled over a set of video chunks $\{u_k \in v\}$. The chunk size is set to the mean chunk length of all concept training chunks. Finally, ρ is the pooling function. We denote average and max pooling as $\rho_a(\cdot)$ and $\rho_m(\cdot)$, respectively.

EDiSE Computational Performance Benefits

Here we discuss the computational complexity of concept based EDiSE, and ASR/OCR based EDiSE. The fusion part is negligible since it is constant time.

Concept based EDiSE

The computational complexity for computing $p(e_c|v)$ is mainly linear in the number of videos, denoted by $|V|$. We here detail why computational complexity of $p(e_c|v)$ is almost constant and hence video retrieval is almost $O(|V|)$.

From Eq. 5, $p(e_c|v)$ has a computational complexity of $O(N \cdot Q)$ for on e video, where Q is the computational complexity of computing $s(\cdot, \cdot)$ and N is the number of concepts. We detail next the computational complexity of $s_p(\cdot, \cdot)$ and $s_t(\cdot, \cdot)$ for the whole set of videos $|V|$.

Complexity of $p(e_c|v)$ for $s_p(\cdot, \cdot)$ Let’s assume that there $\theta(e_c)$ set has $|e_c|$ terms and $\theta(c_i)$ has $|c_i|$ terms. Then, the computational complexity of $s_p(\theta(e_c), \theta(c_i))$ is




Figure 4: Concept probabilities from videos ($p(\mathbf{c}|v)$)

$O(M(|e_c| + |c_i|))$. $|c_i|$ and $|e_c|$ are usually few terms in our case (< 10). Hence the computational complexity of $s_p(\theta(e_c), \theta(c_i))$ is $O(M)$, where M is the dimensionality of the word vectors. In our experiments $M = 300$. Given the complexity of $s_p(\theta(e_c), \theta(c_i))$, the computational complexity of $p(e_c|v)$ will be $O(N \cdot M)$, where N is the number of concepts. Hence, the computational complexity for computing $p(e_c|v)$ for $|V|$ videos is $O(|V| \cdot N \cdot M)$. However, for a given event, only few concepts are relevant, which are computed based on $s_p(\theta(e_c), \theta(c_i))$ and only few concepts 5 in our case are sufficient for event zero shot retrieval, retrieved by Nearest Neighbor search of $c_i \in \mathbf{c}$ that is close the e_c . Hence the computational complexity reduced to $O(|V| \cdot M)$, $M = 300$ for the GoogleNews word2vec model that we used. Hence, the complexity for $|V|$ videos is basically linear $O(|V|)$, given M is a constant and $M \ll |V|$.

Complexity of $p(e_c|v)$ for $s_t(\cdot, \cdot)$ The previous argument applies here in all elements except the complexity of the similarity function $s_t(\theta(e_c), \theta(c_i))$, which is $O(M(|e_c| \cdot |c_i|))$. Assuming that $|e_c| \cdot |c_i|$ is bounded by a constant, then the complexity of $|V|$ videos is also $O(|V| \cdot M)$, but with a bigger constant compared to $s_p(\cdot, \cdot)$ (linear in $|V|$ for constant $M \ll |V|$).

ASR/OCR based EDiSE

The computational complexity of $s_d(\theta(e_o), \theta(v_o))$ and $s_d(\theta(e_a), \theta(v_a))$ are $O(|e_o| \cdot |v_o| \cdot M)$ and $O(|e_a| \cdot |v_a| \cdot M)$, respectively. There is no concepts for ASR/OCR based retrieval. Hence, the computational complexity of $p(e_o, v)$ and $p(e_a|v)$ are $O(|V| \cdot |e_o| \cdot |v_o| \cdot M)$ and $O(|V| \cdot |e_a| \cdot |v_a| \cdot M)$, respectively. Since $|e_o| \ll |V|$, $|v_o| \ll |V|$, $|e_a| \ll |V|$, $|v_a| \ll |V|$, and $M \ll |V|$, the dominating factor in the complexity for both $p(e_o, v)$ and $p(e_a|v)$ will be $|V|$.

Experiments

We evaluated our method on the large TRECVID MED (Felzenszwalb, McAllester, and Ramanan 2013). We show the MAP (Mean Average Precision) and ROC AUC performance of the designated MEDTest set (Felzenszwalb, McAllester, and Ramanan 2013), containing more than 25,000 videos. Unless otherwise mentioned, our results are on TRECVID MED2013. There are two distributional semantic models in our experiments, trained on Wikipedia and GoogleNews using (Mikolov et al. 2013b). The Wikipedia model got trained on 1 billion words resulting in a vocabulary of size of $\approx 120,000$ words and word vectors of 250 dimensions. The GoogleNews model got trained on 100 billion words resulting in a vocabulary of size 3 million words

and word vectors of 300 dimensions. The objective of having two models is to compare how well our EDiSE method performs depending on the size of the training corpus, used to train the language model. In the rest of this section, we present Concepts, OCR, ASR, and fusion results.

Concept based Retrieval

All the results in this section were generated by automatically retrieved concepts using only the event title. We start by comparing different settings of our method against (Dalton, Allan, and Mirajkar 2013). We used the language model in (Dalton, Allan, and Mirajkar 2013) for concept based retrieval to rank the concepts. This indicates that $p(e|c_i)$ in Eq. 4 is computed by the language model in (Metzler and Croft 2005) as adopted in (Dalton, Allan, and Mirajkar 2013), that we compare with under exactly the same setting. For our model, we evaluated the two pooling operations $\rho_m(\cdot)$ and $\rho_p(\cdot)$ and also the two different similarity measures on \mathcal{M}_s space $s_p(\cdot, \cdot)$ and $s_t(\cdot, \cdot)$. Furthermore, we evaluated the methods on both Wikipedia and GNews language models. In order to have conclusive experiments on these eight settings of our model compared to (Dalton, Allan, and Mirajkar 2013), we performed all of them on the four different sets of concepts (i.e. each has the same concept detectors; completely consistent comparison); see Table 1. Details about these concept sets are attached in SM.

There are a number of observations. (1) using GNews (the bigger text corpus) language model is consistently better than using the Wikipedia language model. This indicates when the word embedding model is trained with a bigger text corpus, it captures more semantics and hence more accurate in our setting. (2) max pooling $\rho_m(\cdot)$ behaves consistently better than average pooling $\rho_a(\cdot)$. (3) $s_p(\cdot, \cdot)$ similarity measure is consistently better than $s_t(\cdot, \cdot)$, which we see very interesting since this indicates that our hypothesis of using the vector operations on \mathcal{M}_s manifold better represent $p(e|c_i)$. Hence, we recommend finally to use the model trained on the larger corpus, $\rho_m(\cdot)$ for concept pooling, and use $s_p(\cdot, \cdot)$ to measure the performance on \mathcal{M}_s manifold. (4) our model's final setting is consistently better than (Dalton, Allan, and Mirajkar 2013). The final MED13 ROC AUC performance is 0.834. MAP for MED13 Events 31 to 40 (E31:40) is 5.97%. Detailed figures are attached in SM.

Our next experiment shows the final MAP performance using the recommended setting for our framework on the whole set of concepts, detailed earlier and in SM. Table 2 shows our final performance compared with (Dalton, Allan, and Mirajkar 2013) on the same concept detectors. It is not hard to see that our method performs more than double the

Table 1: MED2013 MAP performance on four concept sets (event title query)

TRECVID MED 2013	Ours-Gnews				Ours-Wiki				(Dalton etal, 2013)
	$\rho_m(\cdot)$		$\rho_a(\cdot)$		$\rho_m(\cdot)$		$\rho_a(\cdot)$		
	$s_p(\cdot, \cdot)$	$s_t(\cdot, \cdot)$	$s_p(\cdot, \cdot)$	$s_t(\cdot, \cdot)$	$s_p(\cdot, \cdot)$	$s_t(\cdot, \cdot)$	$s_p(\cdot, \cdot)$	$s_t(\cdot, \cdot)$	
Concepts G1 (152 concepts)	4.29	3.94%	2.39%	2.38%	3.14%	2.13%	1.85%	1.70%	2.57%
Concepts G2 (101 concepts)	1.74	1.20	1.56%	1.20%	1.09%	0.96%	0.66%	0.60%	1.17%
Concepts G3 (60 concepts)	1.72	1.33%	1.28%	1.16%	1.21%	0.88%	0.88%	0.74%	1.54%
Concepts G4 (56 concepts)	1.22	0.95	0.84%	0.69%	0.87%	0.76%	0.67%	0.56%	0.83%

Table 2: MED2013 full concept set MAP Performance (auto-weighted versus manually-weighted concepts)

Ours (auto)	Dalton et al,13(auto)		Dalton et al,13(manual)		Overfeat
8.36%	3.40%		7.4%		2.43%
SUN	Object Rank	Classeme	CD^{DT}	$WSC_{YouTube}^{D-SIFT}$	
0.48%	0.77%	0.84%	2.28%	3.48%	

MAP performance of (Dalton, Allan, and Mirajkar 2013) under the same concept set. Even when manual semantic editing is applied to (Dalton, Allan, and Mirajkar 2013), our performance is still better without semantic editing. We also show the performance on the same events of different concepts (i.e. SUN (Patterson and Hays 2012), Object Rank (Li et al. 2010), Classeme (Torresani, Szummer, and Fitzgibbon 2010)), and the best performing concepts in (Wu et al. 2014) (i.e., CD^{DT} , $WSC_{YouTube}^{D-SIFT}$). These numbers are as reported in (Wu et al. 2014). The results indicate the value of our concepts and approach compared to (Wu et al. 2014) and their concepts. We also report our performance using Overfeat concepts only to retrieve videos for the same events. This shows the value of involving action and scene concepts compared to only still image concepts like Overfeat for zero-shot event detection. We highlight that the results in (Wu et al. 2014) uses the whole event description which explicitly includes names of relevant concepts.

ASR and OCR based Retrieval

First, we compared our OCR and ASR retrieval trained on both Wikipedia and GoogleNews language model. Table 3

shows that the GoogleNews model MED13 MAP is better than the Wikipedia Model MAP in both ASR and OCR, which is consistent with our concept retrieval results. Fig. 5 shows the GoogleNews MED13 AP per event for both OCR and ASR. We further show our AP performance on MED14 events 31 to 40 in Fig. 5.

In order to show the value our semantic modeling, we computed the performance of string matching method as a baseline, which basically increment the score for every exact match in the the detected text to words in the query. While, both our model and the matching model use the same query words and ASR/OCR detection, semantic properties captured by \mathcal{M}_s boosts the performance compared to string matching; see table 3. This is since semantically relevant terms to the query have a high cosine similarity in \mathcal{M}_s (i.e., high $vec(w_i)^T vec(w_j)$ if w_i is semantically related to

Table 3: ASR & OCR Retrieval MAP on \mathcal{M}_s using GNews, Wikipedia, and using word matching

	GNews MED2013	Wiki MED2013	matching MED2013
OCR	4.81%	3.85%	1.8%
ASR	4.23%	1.50%	3.77%

Table 4: ASR & OCR MAP performance using GNews corpus compared to (Wu et al. 2014)(prefix E indicates Event)

	MED13	MED13 (Wu et al. 2014)	MED14(E31:40)
OCR	4.81%	4.30%	2.5%
	MED13	MED13 (Wu et al. 2014)	MED14 (E31:40)
ASR	4.23%	3.66%	5.97%

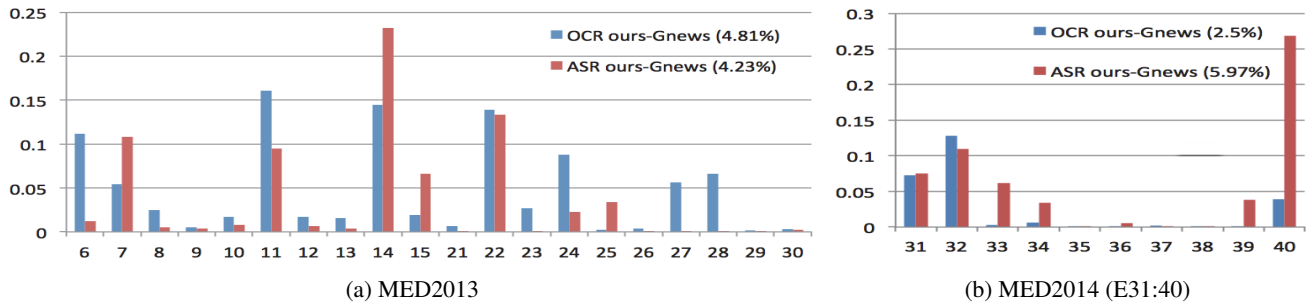


Figure 5: ASR & OCR AP Performance (Google News)




Figure 6: ASR & OCR AUCs on MED2013: Ours (GoogleNews) vs keyword Matching (the same query)

w_j). On the other hand, hard matching basically assumes that $\text{vec}(w_i)^\top \text{vec}(w_j) = 1$ if $w_i = w_j$, 0 otherwise. We also computed the ROC AUC metric for our method and the hard matching method on ASR and OCR; see Fig. 6. For ASR, average AUC is 0.623 for ours and 0.567 for Matching (9.9% gain). For OCR, average AUC is 0.621 for ours and 0.53 for Matching (17.1% gain). We report our GNews model results compared with (Wu et al. 2014) to indicate that, we achieve state-of-the-art MED13 MAP performance or even better for ASR/OCR; see table 4. The table also shows our ASR&OCR MED14 (E31:40) MAP.

Fusion Experiments and Related Systems

In table 5, we start by presenting a summary of our earlier ASR/OCR results on MED13 Test. Comparing OCR and ASR performances to Concepts performance, it is not hard to see that OCR/ASR have much lower average AUC zero-shot performance compared to concepts which are visual in our work. This indicates that OCR/ASR produces much higher false negatives compared to visual concepts. When we fused our all OCR and ASR confidences, we achieved 10.7% MAP performance, however, the average AUC performance is as low as 0.67. We achieved lower MAP for our concepts 8.36% MAP but the average AUC performance is as high as 0.834. This indicates that measuring retrieval performance on MAP performance only is not informative, so one approach might achieve a high MAP but lower average AUC and vice versa. We further achieved the best performance of our system by fusing all Concepts, OCR, and ASR to achieve 13.1% MAP and 0.830 average AUC. We found our system achieves better than the state of the art system (Wu et al. 2014) 4.0% gain in MAP, but significantly in average AUC; see 13.6% gain to (Wu et al. 2014) in table 5.

We also discuss CPRF (Yang and Hanjalic 2010), MM-PRF (Jiang et al. 2014b), and SPaR (Jiang et al. 2014a) reranking systems in contrast to our system that does not involve reranking. The initial retrieval performance is 3.9% MAP without reranking. Interestingly, we achieved a performance of 13.1% MAP also without reranking. The reranking methods assumes high top 5-10 precision of the initial ranking and that all test videos are available. Without any of these assumptions, our system without reranking performs 6.7%, 3.0%, and 0.2% better than CPRF (Yang and Hanjalic 2010), MM-PRF (Jiang et al. 2014b), and SPaR (Jiang et al. 2014a) re-ranking systems; see table 5. Unfortunately, ROC AUC performances are not available for these method to compare with. Regarding efficiency, given v_c representation of videos, our concept retrieval experiment on our whole concept set it takes ≈ 270 seconds on a 16 cores Intel Xeon processor (64GB RAM) to the retrieval task on 20 events altogether. This is more than the time that SPaR (Jiang et al. 2014a) takes to rerank one event on an Intel Xeon processor (16GB RAM); see (Jiang et al. 2014a). Since, we detect the MED13 events in ≈ 270 given v_c representation of videos and as reported in (Jiang et al. 2014a), their average detection time per event for MED13 is ≈ 5 minutes assuming feature representation of videos (i.e., 360 seconds per event = 7200 seconds per 20 events). This indicates that our system

Table 5: Fusion Experiments and Comparison to State of the Art Systems

Method	MAP	AUC
Our Concept retrieval (event title query)	8.36%	0.834
Concept retrieval (Dalton et al. 2013) (event title query)	3.4 %	-
Concept retrieval (Dalton et al. 2013) (manual concepts)	7.4%	-
Our ASR GNews	4.81%	0.623
Our OCR GNews	4.23%	0.621
Our ASR Matching	2.77%	0.567
Our OCR Matching	1.8%	0.536
Our ASR and OCR all fused	10.6	0.670
Our Full (Concepts+ASR+OCR) (No reranking)	13.1%	0.830
Our Full + SPaR reranking (Jiang et al. 2014a)	13.5%	0.790
Full system (Wu et al. 2014)	12.6	0.730
Reranking Systems		
Without Reranking (Jiang et al. 2014b)	3.9%	-
CPRF (Yang and Hanjalic 2010)	6.4%	-
Full Reranking system MM-PRF (Jiang et al. 2014b)	10.1%	-
Full Reranking system SPaR (Jiang et al. 2014a)	12.9%	-

is 26.67X faster than (Jiang et al. 2014a) in MED13 detection. Finally, when we applied SPaR on our output as an initial ranking, we found that it improves MAP (from 13.1% to 13.5%) but hurts ROC AUC (from 0.83 to 0.79). This indicates that reranking has a limited/harmful effect on the performance of our method. We think is since our method already achieve a high performance without re-ranking; see SM for details about the features in this experiment.

Conclusion

We proposed a method for zero shot event detection by distributional semantic embedding of video modalities and with only event title query. By fusing all modalities, our method outperformed the state of the art on the challenging TRECVID MED benchmark. Based on this notion, we also showed how to automatically determine relevance of concepts to an event based on the distributional semantic space.

Acknowledgements. This work has been supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center contract number D11-PC20066. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/NBC, or the U.S. Government. This work is also partially funded by NSF-USA award # 1409683.

References

- Bengio, Y.; Schwenk, H.; Senécal, J.-S.; Morin, F.; and Gauvain, J.-L. 2006. Neural probabilistic language models. In *Innovations in Machine Learning*. 137–186.
- Chen, J.; Cui, Y.; Ye, G.; Liu, D.; and Chang, S.-F. 2014. Event-driven semantic concept discovery by exploiting weakly tagged internet images. In *ICMR*.
- Dalal, N., and Triggs, B. 2005. Histograms of oriented gradients for human detection. In *CVPR*.

- Dalton, J.; Allan, J.; and Mirajkar, P. 2013. Zero-shot video retrieval using content and concepts. In *CIKM*.
- Davis, S., and Mermelstein, P. 1980. Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. *Acoustics, Speech and Signal Processing, IEEE Transactions on* 28(4):357–366.
- Elhoseiny, M.; Saleh, B.; and Elgammal, A. 2013. Write a classifier: Zero-shot learning using purely textual descriptions. In *ICCV*.
- Farhadi, A.; Endres, I.; Hoiem, D.; and Forsyth, D. 2009. Describing objects by their attributes. In *CVPR*.
- Felzenszwalb, P.; McAllester, D.; and Ramanan, D. 2008. A discriminatively trained, multiscale, deformable part model. In *CVPR*.
- Felzenszwalb, P.; McAllester, D.; and Ramanan, D. 2013. Trecvid 2013 an overview of the goals, tasks, data, evaluation mechanisms and metrics. In *TRECVID 2013*, 1–8. NIST.
- Frome, A.; Corrado, G. S.; Shlens, J.; Bengio, S.; Dean, J.; Mikolov, T.; et al. 2013. Devise: A deep visual-semantic embedding model. In *NIPS*.
- Google2014. Youtube, howpublished = "https://www.youtube.com/yt/press/statistics.html", year = 2014, note = "[online; 11/06/2014]".
- Habibian, A.; Mensink, T.; and Snoek, C. G. 2014. Composite concept discovery for zero-shot video event detection. In *ICMR*.
- Jiang, L.; Meng, D.; Mitamura, T.; and Hauptmann, A. G. 2014a. Easy samples first: Self-paced reranking for zero-example multimedia search. In *ACM Multimedia*.
- Jiang, L.; Mitamura, T.; Yu, S.-I.; and Hauptmann, A. G. 2014b. Zero-example event search using multimodal pseudo relevance feedback. In *ICMR*.
- Lampert, C. H.; Nickisch, H.; and Harmeling, S. 2009. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*.
- Li, L.-J.; Su, H.; Fei-Fei, L.; and Xing, E. P. 2010. Object bank: A high-level image representation for scene classification & semantic feature sparsification. In *NIPS*.
- Liu, J.; Yu, Q.; Javed, O.; Ali, S.; Tamrakar, A.; Divakaran, A.; Cheng, H.; and Sawhney, H. 2013. Video event recognition using concept attributes. In *WACV*.
- Liu, J.; Kuipers, B.; and Savarese, S. 2011. Recognizing human actions by attributes. In *CVPR*.
- Logan, B., et al. 2000. Mel frequency cepstral coefficients for music modeling. In *ISMIR*.
- Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. *IJCV*.
- Mensink, T.; Gavves, E.; and Snoek, C. G. 2014. Costa: Co-occurrence statistics for zero-shot classification. In *CVPR*.
- Metzler, D., and Croft, W. B. 2005. A markov random field model for term dependencies. In *ACM SIGIR*.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013a. Efficient estimation of word representations in vector space. *ICLR*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013b. Distributed representations of words and phrases and their compositionality. In *NIPS*.
- Mikolov, T.; Le, Q. V.; and Sutskever, I. 2013. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Miller, G. A. 1995. Wordnet: a lexical database for english. *Communications of the ACM* 38(11):39–41.
- Myers, G. K.; Bolles, R. C.; Luong, Q.-T.; Herson, J. A.; and Aradhye, H. B. 2005. Rectification and recognition of text in 3-d scenes. *IJDAR* 7.
- Norouzi, M.; Mikolov, T.; Bengio, S.; Singer, Y.; Shlens, J.; Frome, A.; Corrado, G. S.; and Dean, J. 2014. Zero-shot learning by convex combination of semantic embeddings. In *ICLR*.
- Parikh, D., and Grauman, K. 2011. Interactively building a discriminative vocabulary of nameable attributes. In *CVPR*.
- Patterson, G., and Hays, J. 2012. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *CVPR*.
- Rohrbach, M.; Stark, M.; Szarvas, G.; Gurevych, I.; and Schiele, B. 2010. What helps where—and why? semantic relatedness for knowledge transfer. In *CVPR*.
- Rohrbach, M.; Ebert, S.; and Schiele, B. 2013. Transfer learning in a transductive setting. In *NIPS*.
- Rohrbach, M.; Stark, M.; and Schiele, B. 2011. Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In *CVPR*.
- Salton, G., and Buckley, C. 1988. Term-weighting approaches in automatic text retrieval. *Information processing & management*.
- Sermanet, P.; Eigen, D.; Zhang, X.; Mathieu, M.; Fergus, R.; and LeCun, Y. 2014. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *ICLR*.
- Shen, Y.; He, X.; Gao, J.; Deng, L.; and Mesnil, G. 2014. A convolutional latent semantic model for web search. Technical report, Technical Report MSR-TR-2014-55, Microsoft Research.
- Socher, R.; Ganjoo, M.; Manning, C. D.; and Ng, A. 2013. Zero-shot learning through cross-modal transfer. In *NIPS*.
- Torki, M., and Elgammal, A. 2010. Putting local features on a manifold. In *CVPR*.
- Torresani, L.; Szummer, M.; and Fitzgibbon, A. 2010. Efficient object category recognition using classemes. In *ECCV*.
- van Hout, J.; Akbacak, M.; Castan, D.; Yeh, E.; and Sanchez, M. 2013. Extracting spoken and acoustic concepts for multimedia event detection. In *ICASSP*.
- Wu, S.; Bondugula, S.; Luisier, F.; Zhuang, X.; and Natarajan, P. 2014. Zero-shot event detection using multi-modal fusion of weakly supervised concepts. In *CVPR*.
- Yang, L., and Hanjalic, A. 2010. Supervised reranking for web image search. In *ACM Multimedia*.

Appendix A: Proof $p(e_c|v)$ for $s(\cdot, \cdot) = s_p(\cdot, \cdot)$

We start by Eq. 5 while replacing $s(\cdot, \cdot)$ as $s_p(\cdot, \cdot)$.

$$\begin{aligned}
 p(e_c|v) &\propto \sum_i s_p(\theta(e_c), \theta(c_i)) p(c_i|v) \\
 &\propto \sum_i \frac{\overline{\theta(e_c)}^T \overline{\theta(c_i)}}{\|\theta_{e_c}\| \|\theta_{c_i}\|} v_c^i \propto \frac{\overline{\theta(e_c)}^T}{\|\theta_{e_c}\|} \left(\sum_i \frac{\overline{\theta(c_i)}}{\|\theta_{c_i}\|} v_c^i \right) \quad (7)
 \end{aligned}$$

which is the dot product between $\frac{\overline{\theta(e_c)}^T}{\|\theta_{e_c}\|}$ representing the event embedding, and $\sum_i \frac{\overline{\theta(c_i)}}{\|\theta_{c_i}\|} v_c^i$ representing the video embedding, which is a function of $\psi(v_c^i) = \{\theta_v(c_i) = \theta(c_i)v_c^i\}$. This equation clarifies our notion of distributional semantic embedding of videos and relating it to event title