# **Complex Event Recognition from Images with Few Training Examples**

Unaiza Ahsan\* uahsan3@gatech.edu Chen Sun\*\* chensun@google.com James Hays\* hays@gatech.edu Irfan Essa\* irfan@cc.gatech.edu

\*Georgia Institute of Technology \*\*University of Southern California<sup>1</sup>

## Abstract

We propose to leverage concept-level representations for complex event recognition in photographs given limited training examples. We introduce a novel framework to discover event concept attributes from the web and use that to extract semantic features from images and classify them into social event categories with few training examples. Discovered concepts include a variety of objects, scenes, actions and event sub-types, leading to a discriminative and compact representation for event images. Web images are obtained for each discovered event concept and we use (pretrained) CNN features to train concept classifiers. Extensive experiments on challenging event datasets demonstrate that our proposed method outperforms several baselines using deep CNN features directly in classifying images into events with limited training examples. We also demonstrate that our method achieves the best overall accuracy on a dataset with unseen event categories using a single training example.

# 1. Introduction

The widespread adoption of smart-phones coupled with easy to use photo sharing services has resulted in a massive increase in images shared online. A large number of these images are snapshots of special occasions, which we refer to as social events such as birthdays, graduations, weddings *etc.* or news events such as political campaigns, natural disasters or marathons. Many of these images do not have clean textual labels hence identifying events from visual content alone is non-trivial. The recent success of deep Convolutional Neural Networks (CNNs) in object and scene recognition has resulted due to large labeled training databases such as ImageNet [39] and Places [56]. Current approaches which use pretrained CNNs and fine-tune on datasets also require significant number of labeled examples. Since creating huge labeled datasets from the con-



Figure 1: Event concepts as an intermediate feature representation for recognizing social events in photographs.

stantly evolving space of events is not realistic, we propose to learn an event concept-based representation and leverage that to identify rare events. Discovering web-driven concepts using Wikipedia and Flickr tags, we aim to categorize social events from static photographs when few labeled examples are available. Images of social events inherently consist of a combination of objects (*e.g.* 'banner'), scenes (*e.g.* 'ground'), actions (*e.g.* 'shouting slogans'), event subtypes (*e.g.* 'speech') and attributes (*e.g.* 'protest peacefully'). Object appearance significantly changes when combined with different objects, actions and attributes in cluttered backgrounds. Hence recognizing events from static images requires us to explicitly learn concept classifiers, each concept a combination of objects, scenes, actions and attributes. We call these *event concepts* (see Figure 1).

Event recognition approaches that use concepts or attributes have previously been applied to video-based events [31, 32, 42, 54] where temporal dynamics play an important role in recognizing what is happening in the video. This makes event recognition from a single photograph an interesting and challenging problem domain. Several attributebased recognition methods require datasets annotated with all the concepts [20, 23] which is a tedious process. To by-

<sup>&</sup>lt;sup>1</sup>The author currently works at Google.

pass the need for manual concept labeling and inspired by the recent 'webly supervised' learning approaches [5, 6, 10] that use web content to discover visual concepts, we propose an event concept learning framework using Wikipedia to generate event categories and Flickr tags as our initial pool of concepts. From noisy Flickr tags, we generate segments or phrases using a tweet segmentation algorithm proposed by Li et al. [25] which is a method designed specifically to extract event-centric phrases from noisy twitter streams. Finally, we project each event category on to a word embedding pretrained on the Google News Dataset using the popular word2vec [34] approach, extract nearest neighbors and add them to the pool of segmented phrases. We extract images related to each concept from MS Bing image search engine and compute deep CNN [24] features extracted from a pretrained network on all the images and train concept classifiers. The concept scores predicted on a given test image form the final features for event images.

#### Our primary contributions are:

- 1. A novel framework which involves using web data to discover event related concepts and employing efficient concept pruning strategies that result in clean, relevant and diverse event concepts.
- 2. A concept-based representation that not only improves single-shot event classification performance but can also be generalized to those categories which were not used during concept discovery.
- 3. A large scale Social Event Image Dataset (SocEID) comprising 37,000 general event images belonging to 8 event categories as well as a challenging Rare Events Dataset (RED) comprising 7,000 images belonging to 21 specific real world events.

# 2. Related Work

Attributes have been used to describe objects (both fixed [3, 12, 23, 48] and relative [36]), faces [20], scenes [37] and actions [14, 30]. These attribute detectors are then run on new images for high level recognition [46, 47]. Researchers have explored creating a set (or bank) of detectors pretrained on objects such as Object Banks [27], an ontology of abstract concepts such as Classemes [45] or scene attributes [8, 37].

Another line of work proposes learning visual concepts from the Web with minimal human supervision ('webly supervised approaches'). NEIL [6] uses image search engine results in a semi-supervised setting to learn and train visual concept detectors. LEVAN [10] uses Google NGram corpus to extract all possible words related to a given concept, extracts images from image search engine and learn visual concepts related to any given keyword. The authors of [29] use a multiple instance learning approach to learn concepts from image search results. Some approaches learn concepts from images and their labels [55], from image descriptions [44] or by using a deep network [5] using principles from curriculum learning. Our proposed work is inspired by web supervision but for a different domain. The key difference between our approach and other webly supervised concept learning approaches is that our methods are designed to obtain *event specific* concepts. We explain further in Section 3.

The earliest work addressing event classification from static images [28] classifies sports events (rowing, rock climbing *etc.*) using object and scene information. Some related approaches [2, 16] require scene geometry or temporal alignment between event images to identify individual events. Recently [52] propose to train two deep networks; one on images and the second one on spatial maps of detected people/objects at different scales for event recognition. Our work aims to learn relevant concepts from the web and uses pretrained CNNs for feature extraction thus saving training time. More importantly, training a deep network for identifying events in images requires a large labeled dataset. We attempt to eliminate that requirement by discovering and using general event concepts from the web.

Motivated by the need to learn with few labeled examples, vision researchers have addressed one-shot learning to learn object classifiers [1, 13, 21, 22, 35] and more recently used deep networks [15, 18] and part-based models [49]. Video event recognition community suffers from lack of labeled data for training too. Hence for recognizing events from few examples, many approaches use concepts as intermediate representations for event recognition [31, 50]. Ma et al. [32] use labels in external videos as concepts and jointly model concept classification and event detection. Chen et al. [4] use Flickr tags to discover concepts for an event and its associated text description. Cui et al. [9] propose Concept Bank which consists of events mined from WikiHow and concepts from Flickr tags. Ye et al. [54] extend the previous work and propose to arrange events and their concepts in a hierarchy learned from WikiHow articles and YouTube descriptions. Shao et al. [41] generate video event attributes using crowd-based annotation and use motion channels and appearance to train a deep model. Yang et al. [53] learn video concepts from YouTube descriptions and Flickr tags. They generate event concepts using a tweet segmentation algorithm along with other metrics and train multiple classifiers for a single concept. The major difference between their work and ours is that we begin from event labels instead of descriptions, target image-based event recognition when few labeled examples are available and integrate word2vec based concepts into our concept pool.

air shows	auto show	beauty pageants
all star games	ballet show	beer festivals
american football match	ballroom dance	birdwatching
annual protests	balls	black friday
art exhibition	barbecue	boating
arts festivals	baseball	bowling
astronomy events	basketball match	boxing

Table 1: Sample events mined from Wikipedia

# 3. Approach

We begin by asking the question, "What if all events had limited labeled examples?" In reality, there are several events for which labeled image datasets are available *e.g.* birthdays and weddings. We develop methods and test our approach on datasets containing these popular events by taking only a single labeled example from each category as training data. The rest of the dataset is used for testing. This approach enables us to determine whether our proposed methods work on popular events before taking on the harder task of identifying rare events (using our RED Dataset) from images.

Our proposed method is an adaptation of the webly supervised learning approaches to learn visual concepts relevant to specific event categories. Extracting all visual concepts related to a keyword such as 'birthday' from the web using the approach of [10] results in many concepts that have either little to do with a birthday event or is not generalizable to real world images. Those concepts include birthday settings advertised by event planners online and objects associated with birthday with a clean background and a canonical viewpoint [33]. Hence we argue for eventspecific concepts for complex event recognition from images with few labeled examples.

Our approach is divided into three main parts: Event Concept Discovery, Training Concept Classifiers and Prediction of Concept Scores for Event Classification.

#### **3.1. Event Concept Discovery**

We use Wikipedia to mine a list of events from its category 'Social Events.' This list contains general events (such as birthday) and specific events (such as royal wedding). We do *not* include specific events in our initial events list as the aim is to build a concept bank that is applicable for all events. If we build a concept bank for royal wedding, concepts such as 'Kate Middleton,' 'Buckingham Palace' *etc.* would be too specific to apply to generic event images. Thus we end up with 150 generic social events (see Table 1). This list of events was mined from Wikipedia but filtered (to remove specific events) manually.

Each event category in our list is used to query Flickr and we obtain the first 200 image results. We collect the tags (a set of words describing the image) of those images in the form of captions and this forms the noisy caption pool from which we aim to generate meaningful concepts that can describe general social events. We empirically observe that by increasing the number of images from which we mine Flickr tags, the noise in the data increases hence we limit ourselves to 200 images per event category.

**Tag Segmentation:** For the set of events  $E = \{e_1, e_2, ..., e_n\}$  where, for our work n = 150, we have a set of tags  $T = \{t_1, t_2, ..., t_N\}$ , N = 200n in our case. Our goal is to generate consecutive and non-overlapping segments  $S = \{s_1, s_2, ..., s_m\}$ . These segments can be a single word or phrases. We obtain a set of segments  $S_i = \{s_1, s_2, s_3, ..., s_m\} \subset S$  for each tag  $t_i \in T$ ,  $i \in \{1, ..., N\}$ , by applying a tweet segmentation method [26] which can be modeled as an optimization,

$$\underset{s_{1},s_{2},s_{3},\ldots,s_{m_{i}}}{\operatorname{argmax}}Stk(t_{i}) = \sum_{j=1}^{m_{i}}Stk(s_{j}), \quad (1)$$

where  $Stk(\cdot)$  is a function that computes the *stickiness* of a segment. Stickiness measures the probability that a segment of text is a 'named entity' (name of a person, place or object) in a large text corpus or a knowledge base like Wikipedia. The final score for each segment  $s_j$  is given by

$$\operatorname{score}(s_j) = Stk(s_j) \cdot V_{\operatorname{flickr}}(s_j).$$
<sup>(2)</sup>

In Equation 2  $Stk(s_j)$  refers to the stickiness score of the *jth* segment and  $V_{\text{flickr}}$  refers to visual representativeness which is a measure of how visually coherent a word or phrase is. Words like 'economy,' and 'public,' if queried to an image search engine, result in ambiguous images, whereas a phrase like 'birthday cake' generally returns very similar images. Hence it has a high visual representativeness score as compared to words like 'economy' or 'activity.' We obtain visual representativeness scores for each segment via a public dataset made available by Sun and Bhowmick [43] where they provide representativeness scores for the most popular tags on Flickr. After computing the final scores, we inspect the highest scoring segments to remove ambiguous or slang words. Figure 2 shows some sample tags and the returned high scoring segments. Further details on tag segmentation can be found in [26].

Generating segments from tags is not enough to cover all aspects of a complex event. To expand the taxonomy and find event-specific concepts we additionally project each

breakdancing, salsa dancing, argentinean tango, wows crowd, freerunning, reggae hiphop, dance disco fever, flash mobs, street dance, dancers, breakbeat, hop, dance craze, dancefest, bollywood bhangra, asian pop, dance workout, hip hop dance troupe, breakdancers

 Table 2: Nearest neighbors of the event 'dance' in word2vec space



Figure 2: Generated segments from Flickr tags for event label 'protest.'



Figure 3: Event concept discovery pipeline for generic social events.

event label to word2vec space [34]. We use an embedding trained on the Google News Dataset which consists of about 100 billion words. The model is available for public use (https://code.google.com/p/word2vec) and contains 300-dimensional vectors for 3 million words and phrases. The advantage of using Google News pretrained vectors is that we obtain semantically similar concepts for each event label. Table 2 shows the nearest neighbors of the event label 'dance' in word2vec space.

We extract top 20 nearest neighbors for each event label and add them to the tag segments generated via the segmentation scheme described above. This pool of event concepts is then filtered to remove duplicate concepts, slang words and foreign words. We finally end up with 856 event concepts. Our concepts not only include objects, scenes and actions but also include sub-events and their types. Our event concept discovery pipeline is shown in Figure 3.

#### **3.2. Training Concept Classifiers**

Now we describe our approach for selecting training images for concept classifiers. Given a set of concepts  $C = \{c_1, c_2, \ldots, c_m\}$ , we input each concept as an image search query to Microsoft Bing and retrieve the top 100

images returned for each concept. Clipart, duplicate images and images containing only text are removed from the search results. Several concepts in C are correlated with each other, resulting in similar images for different concepts. Training them independently without taking into account their correlation will lead to false negatives that will impact the concept classifier training negatively. The reason why we do not exclude correlated concepts in our concept pool C is that they capture different (and thus important) aspects of a social event. For example, 'birthday party,' 'birthday boy' and 'birthday celebrations' are three different concepts within the same event category (birthday) and they may have similar images. Thus when selecting training images for training the classifier for concept  $c_i$ , it is naive to sample negative training images from all concepts  $c_i$  where  $j \neq i$ . Figure 4 shows an example of correlated concepts and their associated images.

Hence, we first cluster all the concepts using their word2vec-based vector representations using minibatch kmeans clustering [40]. We set k = 150. Thus for *ith* concept  $c_i$ , belonging to a y-sized cluster, we obtain a list of concepts  $C_{pos}^i = \{c_{i,pos_1}, c_{i,pos_2}, \dots, c_{i,pos_y}\}$  that are highly correlated with the given concept  $c_i$ . We construct the concept classifier training set for concept  $c_i$  as follows:

- Let  $\xi^+ = \{I_i\}_{i=1}^u$  be the set of positive training images for classification where u is the number of images retrieved for concept  $c_i$ .
- Let  $\xi^{-} = \{I_{j}\}_{j=1}^{v}$  be the set of negative training im-



Figure 4: Examples of correlated event concepts



Figure 5: Selecting training images for 'boxing' classifier

ages for classification where v is the number of images retrieved for the set of concepts  $C_{neg}^i = \overline{C_{pos}^i} = C - C_{pos}^i$ .

In other words, we make sure that the set  $\xi^-$  does not include images retrieved for any concept in the set  $C_{pos}^i$  because those images are highly similar to the images in  $\xi^+$  as they belong to the same cluster as  $c_i$  (See Figure 5).

For each concept, we extract the CNN 'fc7' layer activations as features from all its images, select the training and test examples as described above and input them to logistic regression classifiers. We select the classifier parameters through 5-fold cross validation and for all of our concept classifiers, the cross validation accuracy is above 90%.

#### 3.3. Predicting Concept Scores for Classification

After training all concept classifiers, we compute the classifier scores on images belonging to our evaluation datasets. For each image I, its feature vector is a concatenation of all concept classifier scores predicted on the image. Thus  $f_I = \{x_i\}_{i=1}^m$  where m is the total number of concepts and  $x_i$  is the score predicted for *ith* concept classifier. Finally, we use these features to classify event images into social events using a linear SVM with default parameters fixed for all experiments. We outline our experimental setup in detail in the next section.

## 4. Experiments and Evaluations

We evaluate our approach on four labeled event datasets with one-shot learning, that is, we use a single positive and negative training example from each class. We also report classification results for all-shot learning (using a 70-30 split) for comparison.

## 4.1. Datasets

We evaluate our concept-based event recognition algorithm on the following four datasets:



Figure 6: Sample images of the SocEID Dataset for two events: birthday (top) and graduation (bottom).

1. Social Event Image Dataset (SocEID): This is the dataset we created in-house. We collected images of the following social events: birthdays, graduations, weddings, marathons/races, protests, parades, soccer matches and concerts. We queried Instagram and Flickr with a tag related to the event itself ('wedding day,' 'Graduation 2014' etc.) and downloaded public images in chronological order determined by post date. Our dataset includes some relevant images from the NUS-WIDE dataset [7] and the Social Event Classification subtask from MediaEval 2013 [38]. We passed all the images to 3 trusted coders (non-Turkers) and asked them to filter any image that did not depict a particular social event. We established the final scores on the images via majority vote and discarded all the rest. Finally, we ended up with nearly 37,000 images. Figure 6 shows sample images from the SocEID dataset.

**2. Web Image Dataset for Event Recognition (WIDER):** This dataset is introduced by [52] and consists of 50,574 images annotated with 61 classes.

**3. UIUC Sports Event Dataset:** This dataset [28] consists of 1579 images belonging to 8 sports events categories such as badminton, bocce, sailing *etc*.



Figure 7: Sample images of the Rare Events Dataset.

4. Rare Events Dataset (RED): This is another in-house dataset we collected by querying MS Bing image search engine with a set of 26 'rare' event categories. We call them rare not on the basis of how frequently they occur in the world but on how seldom they are found in large labeled event image datasets. These event categories comprise of recent news events such as: Justin Trudeau elected, election campaign Trump and natural disasters such as Hurricane Katrina, Hurricane Sandy and Nepal earthquake (see Figure 7). The whole dataset comprises nearly 7,000 images and we do not remove any image from any event category manually. Note that these are all *specific* events (the full list can be found in the supplementary section) and our main motivation behind collecting this dataset is twofold: i) Since few labeled examples are available for these events, it is a suitable test case for our claim that our learned concepts are a powerful intermediate representation to recognize events with few examples, ii) We want to test whether our discovered event concepts generalize to recognizing specific event images or not.

rowing	rowing championships, rowing regatta, recreational boating, canoe trip, junior rowing, canoe polo, standup pandling, swimming canoeing, recreational fishing, cable wakeboard
polo	horseback riding, horse show, mountain biking, bronc riding, bareback bronc,
	horseback ride, ranch rodeo, stampede rodeo, canoeing horseback, riding mountain

Table 3: Top 10 predicted concepts for sports events 'rowing' and 'polo'

# 4.2. Experimental Setup

We begin our experiments by training event concept classifiers. We have a total of 86,000 images associated with the concepts, retrieved from Microsoft Bing using the publicly available Bing crawler by Dengxin Dai.<sup>1</sup> We use the Caffe [17] deep learning framework to extract CNN layer 7 activations ('fc7') as features for all the images using HybridCNN which is a publicly available CNN model pretrained on 978 object categories from ImageNet database [39] and 205 scene categories from Places dataset [56] using the AlexNet deep architechture [19]. For each concept, we select the positive and negative training features as described in Section 3.2 and train L2-regularized logistic regression classifiers using the publicly available LIBLINEAR library [11]. Every image in our event datasets is input to each of the trained concept classifiers and a probabilistic concept score is computed on it. The fusing of concept scores form the final feature vector of that image.

#### 4.2.1 Training with a Single Positive Image

We conduct our one-shot learning experiment on the event datasets as follows: For event category E with P positive training features and N negative training features we randomly sample a positive training feature  $f_p^E$  and a negative training feature  $f_n^E$ . We concatenate the two and feed this into a binary linear SVM as training features. For testing, we simply take the rest of the positive features ( $P - f_p^E$ ) and sample an equal number of negative features from the rest of the event categories in the dataset. Hence for all our experiments, the random baseline is 50%. We run all experiments five times and average the per-class classification accuracies.

### 4.2.2 Training with a 70%-30% split

In this experiment we take all of the labeled data into account and for each class, randomly select 70% of images for training and test on the remaining images. This experiment shows the maximum accuracy our method can achieve given all of the training data available. It provides a nice comparison against the case where only a single labeled image is available for each class. We compare our results with several powerful baselines:

• AlexNet [19] pretrained on ImageNet [39] and

<sup>&</sup>lt;sup>1</sup>http://www.vision.ee.ethz.ch/~daid



Figure 8: One-shot learning results on UIUC Sports Dataset and SocEID.

Features	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
AlexNet-fc7	57.51 57.97	55.76 55.54	50.94 53.05	<b>53.26</b>	50.45 57.43	<b>65.23</b>	52.70 54.32	51.56 54.26	53.66 55.46	55.50 60 19	53.01 56.03	52.38	54.53 55 13	<b>57.69</b>	58.86
Event concepts	<b>60.46</b>	56.62	<b>63.64</b>	52.16	<b>59.63</b>	64.96	51.71	51.59	54.80	56.87	57.82	59.29	49.00	57.05	59.47

Table 4: Result of one-shot learning on WIDER Dataset

Places [56] databases, from which we extract 4096dimensional layer fc7's activations and use them as features. We refer to this baseline as AlexNet-fc7 in the results.

- Chen *et al.* [5] which is a recently proposed webly supervised CNN trained on about 2.1 million images downloaded from Google Images using popular vision datasets' labels as search queries. The authors use 2,240 objects, 89 attributes, and 874 scene labels from ImageNet [39], SUN database [51] and NEIL knowledge base [6] and use principles from curriculum learning to train the network with easy examples first and then hard examples from Flickr. They show state of the art performance compared to AlexNet for objection detection and scene recognition. We use their 'GoogleA' network which is trained on 2.1 million Google images. We refer to this baseline as WEBLY-fc7.
- AlexNet [19] finetuned on WIDER database [52]. This baseline model is provided by the authors. We want to see how fc7 features extracted from this model perform with limited training data on our evaluation datasets (except WIDER) and whether our concept-level features are comparable to it. We refer to this baseline as WIDER-fc7.

## 5. Results and Discussion

Our one-shot learning result on UIUC Sports Events dataset (Figure 8 left) shows that the event concept features significantly outperform all the baselines in 6 out of 8 events. From 1-8, the UIUC Sports event categories are: badminton, bocce, croquet, polo, rock climbing, rowing, sailing and snowboarding. The main reason why this occurs is that our initial event list and hence our discovered event concepts include several sports events and their subtypes. For example, for all the images labeled with the event "rowing" and "polo" in the UIUC Sports Dataset, we count the top 10 most frequently predicted concepts. Table 3 qualitatively shows that our method extracts relevant concepts consistently accross the set of UIUC Sports events images.

Our one-shot learning experiment on the SocEID dataset (Figure 8 right) results in event concepts outperforming the baselines in 4 out of 8 categories. From 1-8, the categories are: birthday, concert, graduation, marathons, parade, protest, soccer and wedding. This is very likely due to the nature of images found in our dataset. Our dataset contains very clean images from the Web with significant visual cues of popular events such as birthdays. Thus the webly supervised network of [5] and WIDER [52] finetuned on event images is able to discriminate between the different event classes based on cues such as graduation caps and bridal gowns in the graduation and wedding pictures respectively even if limited training examples are present.

Our experiments on the WIDER dataset [52] yield interesting insights. There are 61 classes in the dataset (the full list can be found in our supplementary material). For brevity, we only show the first 15 classes (we follow the order given by the authors). We test one-shot learning using our event concept features and compare them against the baselines, WEBLY-fc7 and AlexNet-fc7. Our event concept features outperform AlexNet-fc7 and WEBLY-fc7 in 31 classes when training with a single image. Table 4 shows our results. From 1-15, the classes

Event classes	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	Avg. Acc
AlexNet-fc7	54.0	53.4	60.0	56.8	53.3	54.0	53.0	50.3	59.9	56.0	52.5	57.5	65.3	54.0	53.6	53.2	55.5	55.5	64.0	50.7	62.6	56.0
WEBLY-fc7	52.2	58.1	60.9	56.2	53.9	56.3	50.4	55.5	67.8	54.9	61.0	61.3	59.6	55.6	47.8	61.0	53.2	52.1	56.0	52.6	69.0	57.0
WIDER-fc7	50.4	55.5	57.2	52.2	55.8	55.1	51.1	52.8	55.5	52.6	52.6	58.9	63.3	53.7	47.3	49.8	53.4	55.6	61.3	52.8	59.3	54.6
Event concepts	57.5	58.9	67.7	55.5	54.8	53.1	53.9	58.6	75.1	54.2	60.8	55.4	73.4	56.2	52.0	58.0	56.2	52.4	58.7	53.3	64.8	58.6

Table 5: Result of one-shot learning on RED Dataset

Features

AlexNet-fc7

WEBLY-fc7

	Overall Average Accuracy (%)									
Features	UIUC Sports	SocEID	WIDER	RED						
AlexNet-fc7	59.79	52.57	55.40	55.94						
WEBLY-fc7	55.39	60.89	58.14	56.92						
WIDER-fc7	70.81	58.11	N/A	54.58						
Event concepts	82.09	62.98	58.29	58.59						

Table 6: Overall accuracy for one-shot learning on th	ie
evaluation datasets	

are: parade, handshaking, demonstration, riot, dancing, car-accident, funeral, cheering, election-campaign, pressconference, people-marching, meeting, group, interview and traffic. We note that the WIDER dataset contains not only events but also individual actions such as 'cheering'. Our proposed approach uses concepts such as 'cheering' to identify events which typically involve cheering such as dances, games or graduations (as the cheering concept will result in high probabilistic prediction on such events). However, the model is not trained to identify cheering alone. This can be verified by noting our performance scores in classes such as parade (1), demonstration (3), dancing (5) etc. We score well on other categories (not shown in table) such as basketball, soccer, running, aerobics etc. Categories where our scores are comparable but not above the baselines are: sports coach trainer, greeting, surgeons, spa and stock market, to name a few. These categories are recognized more effectively by deep CNNs pretrained to recognize objects and scenes.

Finally, we evaluate our method on the RED dataset (see Table 5) which is the most challenging because the images are highly diverse and consist of specific real world events. In one-shot learning on RED, for 10 out of 21 classes, our proposed method outperfoms the baselines. Thus our method is able to generalize to unseen event categories as our initial event list does not contain any of the rare event categories in the RED dataset. From left to right, categories are: 'Russian airstrikes Syria,' 'Boston Bombing,' 'Nepal Earthquake,' 'Arab Spring,' *etc.* (the full list can be found in the supplementary material).

The overall classification accuracies are shown in Table 7. For all the datasets, event image classification using a single training example with our proposed event concepts as features outperform the baselines which shows the strength of our approach to recognize complex real world events when limited labeled examples are available.

WIDER-fc7	93.85	80.42	N/A	76.64						
Event concepts	96.68	85.39	78.59	77.57						
Table 7: Overall accuracy for all shot learning on the										

**UIUC Sports** 

96.47

95.16

Overall Average Accuracy (%)

SocEID

86.42

83.66

WIDER

77.94

77.85

RED

77.86

79.39

Table 7: Overall accuracy for all-shot learning on the evaluation datasets

We also evaluate our method against the baselines using all the available training data (70%-30% split). Table 7 shows the overall classification accuracies on our evaluation datasets. Even when using all training examples, our event concept features are comparable to the state of the art in recognizing events from images. For two datasets, (UIUC Sports and WIDER) our method actually outperforms the state of the art.

# 6. Conclusion

In this paper we propose to discover event-specific concepts from the web to recognize complex events from images with few labeled examples. Our proposed framework discovers relevant concepts by combining segmented Flickr tags and word2vec nearest neighbors of event categories resulting in a compact intermediate representation which identifies real world events with only a single training example. We show the strength of our proposed method by evaluating on challenging datasets against powerful baselines which directly use CNN features pretrained on objects, scenes, attributes and events. It is interesting to note that in the problem domain of event recognition from visual content where only a few training examples are available, webdriven concept discovery and web images for training can result in highly discriminative intermediate representations which outperform directly using deep CNNs trained on millions of images and even deep CNNs finetuned on a large event dataset.

# References

 E. Bart and S. Ullman. Cross-generalization: Learning novel classes from a single example by feature replacement. In *Computer Vision and Pattern Recognition, 2005. CVPR* 2005. *IEEE Computer Society Conference on*, volume 1, pages 672–679. IEEE, 2005.

- [2] L. Bossard, M. Guillaumin, and L. Van. Event recognition in photo collections with a stopwatch hmm. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 1193–1200. IEEE, 2013.
- [3] H. Chen, A. Gallagher, and B. Girod. Describing clothing by semantic attributes. In *Computer Vision–ECCV 2012*, pages 609–623. Springer, 2012.
- [4] J. Chen, Y. Cui, G. Ye, D. Liu, and S.-F. Chang. Event-driven semantic concept discovery by exploiting weakly tagged internet images. In *Proceedings of International Conference* on Multimedia Retrieval, page 1. ACM, 2014.
- [5] X. Chen and A. Gupta. Webly supervised learning of convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1431–1439, 2015.
- [6] X. Chen, A. Shrivastava, and A. Gupta. Neil: Extracting visual knowledge from web data. In *Computer Vision (ICCV)*, 2013 IEEE International Conference on, pages 1409–1416. IEEE, 2013.
- [7] T.-S. Chua, J. Tang, R. Hong, H. Li, Z. Luo, and Y.-T. Zheng. Nus-wide: A real-world web image database from national university of singapore. In *Proc. of ACM Conf. on Image and Video Retrieval (CIVR'09)*, Santorini, Greece., July 8-10, 2009.
- [8] G. Ciocca, C. Cusano, S. Santini, and R. Schettini. Halfway through the semantic gap: Prosemantic features for image retrieval. *Information Sciences*, 181(22):4943–4958, 2011.
- [9] Y. Cui, D. Liu, J. Chen, and S.-F. Chang. Building a large concept bank for representing events in video. arXiv preprint arXiv:1403.7591, 2014.
- [10] S. K. Divvala, A. Farhadi, and C. Guestrin. Learning everything about anything: Webly-supervised visual concept learning. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 3270–3277. IEEE, 2014.
- [11] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *The Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [12] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. In *Computer Vision and Pattern Recognition*, 2009. CVPR 2009. IEEE Conference on, pages 1778–1785. IEEE, 2009.
- [13] L. Fei-Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *Pattern Analysis and Machine Intelligence*, *IEEE Transactions on*, 28(4):594–611, 2006.
- [14] Y. Fu, T. M. Hospedales, T. Xiang, and S. Gong. Attribute learning for understanding unstructured social activity. In *Computer Vision–ECCV 2012*, pages 530–543. Springer, 2012.
- [15] D. Held, S. Thrun, and S. Savarese. Deep learning for singleview instance recognition. arXiv preprint arXiv:1507.08286, 2015.
- [16] V. Jain, A. Singhal, and J. Luo. Selective hidden random fields: Exploiting domain-specific saliency for event classification. In *Computer Vision and Pattern Recognition*, 2008. *CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [17] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolu-

tional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.

- [18] G. Koch. Siamese Neural Networks for One-Shot Image Recognition. PhD thesis, University of Toronto, 2015.
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [20] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar. Attribute and simile classifiers for face verification. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 365–372. IEEE, 2009.
- [21] B. M. Lake, R. Salakhutdinov, J. Gross, and J. B. Tenenbaum. One shot learning of simple visual concepts. In Proceedings of the 33rd Annual Conference of the Cognitive Science Society, volume 172, page 2, 2011.
- [22] B. M. Lake, R. R. Salakhutdinov, and J. Tenenbaum. Oneshot learning by inverting a compositional causal process. In *Advances in neural information processing systems*, pages 2526–2534, 2013.
- [23] C. H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *Computer Vision and Pattern Recognition*, 2009. *CVPR 2009. IEEE Conference on*, pages 951–958. IEEE, 2009.
- [24] B. B. Le Cun, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Handwritten digit recognition with a back-propagation network. In *Advances in neural information processing systems*. Citeseer, 1990.
- [25] C. Li, A. Sun, and A. Datta. Twevent: segment-based event detection from tweets. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 155–164. ACM, 2012.
- [26] C. Li, J. Weng, Q. He, Y. Yao, A. Datta, A. Sun, and B.-S. Lee. Twiner: named entity recognition in targeted twitter stream. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 721–730. ACM, 2012.
- [27] F. Li, J. Carreira, and C. Sminchisescu. Object recognition as ranking holistic figure-ground hypotheses. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1712–1719. IEEE, 2010.
- [28] L.-J. Li and L. Fei-Fei. What, where and who? classifying events by scene and object recognition. In *Computer Vision*, 2007. ICCV 2007. IEEE 11th International Conference on, pages 1–8. IEEE, 2007.
- [29] Q. Li, J. Wu, and Z. Tu. Harvesting mid-level visual concepts from large-scale internet images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 851–858, 2013.
- [30] J. Liu, B. Kuipers, and S. Savarese. Recognizing human actions by attributes. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3337–3344. IEEE, 2011.
- [31] J. Liu, Q. Yu, O. Javed, S. Ali, A. Tamrakar, A. Divakaran, H. Cheng, and H. Sawhney. Video event recognition using concept attributes. In *Applications of Computer Vision*

(WACV), 2013 IEEE Workshop on, pages 339-346. IEEE, 2013.

- [32] Z. Ma, Y. Yang, Z. Xu, S. Yan, N. Sebe, and A. G. Hauptmann. Complex event detection via multi-source video attributes. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 2627–2633. IEEE, 2013.
- [33] E. Mezuman and Y. Weiss. Learning about canonical views from internet image collections. In Advances in Neural Information Processing Systems, pages 719–727, 2012.
- [34] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781, 2013.
- [35] E. G. Miller, N. E. Matsakis, and P. A. Viola. Learning from one example through shared densities on transforms. In *Computer Vision and Pattern Recognition*, 2000. Proceedings. IEEE Conference on, volume 1, pages 464–471. IEEE, 2000.
- [36] D. Parikh and K. Grauman. Relative attributes. In Computer Vision (ICCV), 2011 IEEE International Conference on, pages 503–510. IEEE, 2011.
- [37] G. Patterson and J. Hays. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2751–2758. IEEE, 2012.
- [38] T. Reuter, S. Papadopoulos, G. Petkos, V. Mezaris, Y. Kompatsiaris, P. Cimiano, C. de Vries, and S. Geva. Social event detection at mediaeval 2013: Challenges, datasets, and evaluation. In *Proceedings of the MediaEval 2013 Multimedia Benchmark Workshop Barcelona, Spain, October 18-19,* 2013, 2013.
- [39] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, pages 1–42, April 2015.
- [40] D. Sculley. Web-scale k-means clustering. In Proceedings of the 19th international conference on World wide web, pages 1177–1178. ACM, 2010.
- [41] J. Shao, K. Kang, C. C. Loy, and X. Wang. Deeply learned attributes for crowded scene understanding. In *Proc. CVPR*, 2015.
- [42] B. Singh, X. Han, Z. Wu, V. I. Morariu, and L. S. Davis. Selecting relevant web trained concepts for automated event retrieval. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4561–4569, 2015.
- [43] A. Sun and S. S. Bhowmick. Quantifying tag representativeness of visual content of social images. In *Proceedings of the international conference on Multimedia*, pages 471–480. ACM, 2010.
- [44] C. Sun, C. Gan, and R. Nevatia. Automatic concept discovery from parallel text and visual corpora. In *Proceedings* of the IEEE International Conference on Computer Vision, pages 2596–2604, 2015.
- [45] L. Torresani, M. Szummer, and A. Fitzgibbon. Efficient object category recognition using classemes. In *Computer Vision–ECCV 2010*, pages 776–789. Springer, 2010.

- [46] J. Vogel and B. Schiele. Semantic modeling of natural scenes for content-based image retrieval. *International Journal of Computer Vision*, 72(2):133–157, 2007.
- [47] G. Wang, D. Hoiem, and D. Forsyth. Learning image similarity from flickr groups using stochastic intersection kernel machines. In *Computer Vision*, 2009 IEEE 12th International Conference on, pages 428–435. IEEE, 2009.
- [48] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-ucsd birds 200. 2010.
- [49] A. Wong and A. L. Yuille. One shot learning via compositions of meaningful patches. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1197– 1205, 2015.
- [50] S. Wu, S. Bondugula, F. Luisier, X. Zhuang, and P. Natarajan. Zero-shot event detection using multi-modal fusion of weakly supervised concepts. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 2665–2672. IEEE, 2014.
- [51] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *Computer vision and pattern recognition (CVPR)*, 2010 IEEE conference on, pages 3485–3492. IEEE, 2010.
- [52] Y. Xiong, K. Zhu, D. Lin, and X. Tang. Recognize complex events from static images by fusing deep channels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1600–1609, 2015.
- [53] X. Yang, T. Zhang, C. Xu, and M. S. Hossain. Automatic visual concept learning for social event understanding. *Multimedia*, *IEEE Transactions on*, 17(3):346–358, 2015.
- [54] 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 *Proceedings of the 23rd Annual* ACM Conference on Multimedia Conference, pages 471– 480. ACM, 2015.
- [55] B. Zhou, V. Jagadeesh, and R. Piramuthu. Conceptlearner: Discovering visual concepts from weakly labeled image collections. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1492–1500, 2015.
- [56] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. In *Advances in Neural Information Processing Systems*, pages 487–495, 2014.