

Discover and Learn New Objects from Documentaries

Kai Chen Hang Song Chen Change Loy Dahua Lin

Department of Information Engineering, The Chinese University of Hong Kong

ck015,hsong,ccloy,dhlin@ie.cuhk.edu.hk

Abstract

Despite the remarkable progress in recent years, detecting objects in a new context remains a challenging task. Detectors learned from a public dataset can only work with a fixed list of categories, while training from scratch usually requires a large amount of training data with detailed annotations. This work aims to explore a novel approach – learning object detectors from documentary films in a weakly supervised manner. This is inspired by the observation that documentaries often provide dedicated exposition of certain object categories, where visual presentations are aligned with subtitles. We believe that object detectors can be learned from such a rich source of information. Towards this goal, we develop a joint probabilistic framework, where individual pieces of information, including video frames and subtitles, are brought together via both visual and linguistic links. On top of this formulation, we further derive a weakly supervised learning algorithm, where object model learning and training set mining are unified in an optimization procedure. Experimental results on a real world dataset demonstrate that this is an effective approach to learning new object detectors.

1. Introduction

Recent years have witnessed a wave of innovation in object detection driven by the advances in deep learning [12, 29, 35]. Despite the great success reported on public benchmarks, practical applications of such techniques are impeded by a significant obstacle, namely, the lack of annotated data. Specifically, detectors pre-trained on public datasets can only cover a limited list of object categories, which are often not sufficient for real-world applications, where the objects of interest can go beyond such lists. On the other hand, training a new detector requires a large quantity of annotated images with bounding boxes provided for individual objects. Obtaining such a dataset is an investment that takes tremendous amount of time and resources.

These difficulties call for an alternative approach to learning object detectors. “There is no record of an Orca doing any harm in the wild.” – we were deeply touched by

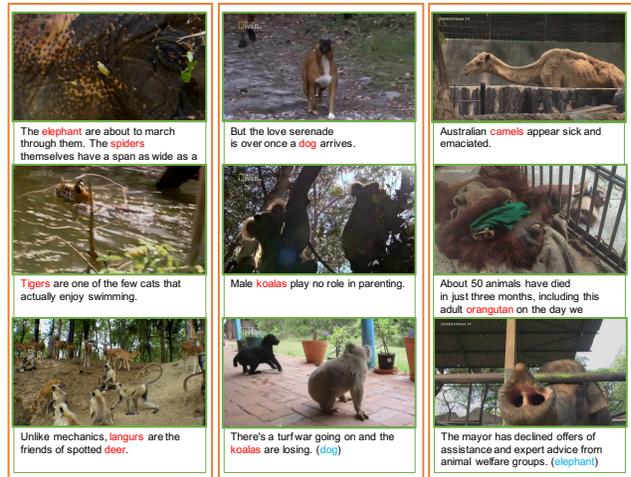


Figure 1: We wish to devise a method that can *automatically* learn object detectors as it watches a documentary, in a weakly supervised fashion. The figures show some examples of the proposed Wildlife Documentaries (WLD) dataset. The nouns are marked in red in there is a correspondence between subtitle and object in the video. Otherwise, we provide the corresponding noun in blue color when there is a mismatch.

this as we watched *Blackfish*, a famous documentary film directed by Gabriela Cowperthwaite. For many of us, the knowledge about various creatures on the planet are learned from documentaries as such. Now, as we move steadily towards the era of AI, we are motivated to ask: *Can computers learn in a similar way?*

From an intuitive perspective, the idea of exploiting documentaries to learn object detectors is promising. Thousands of documentary films that cover a broad range of topics are produced every year, and the number is rising rapidly. A documentary film usually comprises a number of visual presentations with subtitles that cover the entities being introduced from different perspectives. Clearly, this is a rich source of information for visual learning.

In this work, we aim to derive a method that can *automatically* learn object detectors as it watches a documentary, in a weakly supervised fashion (Fig. 1). Towards this goal, we are facing multiple challenges. First and foremost, the framework has completely no prior knowledge as

to what it will see from the videos, *e.g.* the number of object categories and how the story is going to be presented. Everything has to be discovered from the data – no human assistance is available. Second, associations between linguistic cues and visual cues are often ambiguous. For example, when multiple objects are present at the same time, it is difficult to identify the correspondence between the *nouns* in a subtitle and the *objects* that appear in the associated frames. The ubiquitous pronouns and synonyms further complicate the problem. – a *tiger* is often called a *cub* or a *big cat*, it can also be referred to using pronouns, such as “*it*” or “*this*”. Third, there are substantial variations of an objects appearance due to position changes, blurring, and partial occlusions. Even state-of-the-art recognition techniques may experience difficulties in handling all such circumstances.

The difficulties are clear – whereas the materials contained in documentaries are rich and relevant, the observations that we work with are filled with noises, ambiguities, and disconnected pieces. In tackling this challenge, we develop a *probabilistic framework*, with an aim to bridge individual pieces of information, turning them into a coherent picture. Specifically, the framework incorporates three kinds of factors as *bridges*: (1) the *appearance* factor that captures the common appearance pattern of each object category, (2) the *geometric* factor that stitches isolated tracklets of an object into a whole, and (3) the *grounding* factor that associates linguistic references with visual observations. In addition, the framework also identifies nouns and pronouns that describe the same entities via coreference analysis. On top of this formulation, we further derive an learning algorithm, where the learning of object detectors and the mining of training data are conducted along with each other within a unified optimization procedure.

The main **contributions** of this work lie in several aspects: (1) a novel approach to learning object detectors, namely, learning from documentary videos in a weakly supervised way, without any annotated object seeding or bag-level supervision as in multiple instance learning methods.; (2) a framework that can effectively integrate noisy pieces of information, including visual and linguistic cues; and (3) a new dataset with detailed annotations (Fig. 1)¹.

2. Related Work

Weakly-supervised object localization. The time-consuming annotation process in object detection can be sidestepped by weakly supervised learning [2, 3, 7, 9, 13, 30, 33, 36, 40, 41, 42, 43, 44]. In many cases, the supervised information is restricted to binary labels that indicate the absence/presence of object instances in the image, without their locations. Typically, a multiple instance learning [25] framework is adopted. Specifically, each im-

age is considered as a ‘bag’ of examples given by object proposals. Positive images are assumed to contain at least one positive instance window, while negative images do not have the object at all. A good overview is provided in [7]. Prest *et al.* [33] introduce an approach for learning object detectors from real-world web videos known only to contain objects of a target class. In other words, their study requires one label per video. Our problem is more challenging in that documentaries do not provide precise labels even at the ‘bag’ level, so we do not have definite positive and negative windows in videos. We thus require an effective framework to integrate noisy pieces of information. Joulin *et al.* [16] propose a method to localize objects of the same class across a set of distinct images or videos. The colocalization problem assumes each frame in a set of videos contains one object of the same category.

Other relevant studies include [22, 27] that perform semi-supervised learning to iteratively learn and label object instances from long videos. Annotated seeds are required in these approaches. Kumar *et al.* [19] transfer tracked object boxes from weakly-labeled videos to weakly-labeled images to automatically generate pseudo ground-truth boxes. Our approach works directly on video. Alayrac *et al.* [1] model narrated instruction videos to learn the sequence of main steps to complete a certain task. Many videos with transcript of the same task are required. Kwak *et al.* [20] formulate the problem of unsupervised spatio-temporal object localization as a combination of discovery and tracking. It locates one instance each frame and cannot align semantic labels with clusters of objects. Recent studies [28, 45] show that object detectors emerge within a convolutional neural network (CNN) trained with image-level labels. We leverage on this concept to generate candidate proposals in our approach.

Grounding objects from image descriptions. A number of approaches have been proposed to localize objects in an image given its description. For instance, Karpathy *et al.* [18] address the inter-modal alignment problem by embedding detection results from a pretrained object detector and the dependency tree from a parser with a ranking loss. Plummer *et al.* [32] learn a joint embedding of image regions and text snippets using Canonical Correlation Analysis (CCA) to localize objects mentioned in the caption. Recent studies [15, 23] build upon image captioning frameworks such as LRCN [10] or m-RNN [24], which are trained with ground-truth phrase-region pairs of known object classes. The idea is extended to image segmentation from natural language expressions [14]. Rohrbach *et al.* [37] present a latent attention approach that learns to attend to the right region of an image by reconstructing phrases. In contrast to aforementioned studies, our work neither assumes objects of seen classes nor any paired ground-truth phrase-image or phrase-region data. Instead,

¹Dataset and code are available at <https://github.com/hellock/documentary-learning>.

we focus on discovering and learning to detect unknown objects with the help of unstructured linguistic references.

Linguistic cues in vision. Subtitles have been exploited for assisting the learning of visual recognizer. Several studies [6, 8, 31] automatically learn British Sign Language signs from TV broadcasts. Their videos contain a single signer with a stable pose. Positive correspondence between subtitle and signing can be easily identified due to the more structured data. Ours contain multiple animals moving around, exhibiting various poses and scales. Our problem thus requires a more error-tolerant formulation for learning and linking the appearance, geometric, and grounding factors. Another study [11] reduces ambiguity in automatic character identification by aligning subtitles with transcripts that contain spoken lines and speaker identity. In our case, we do not have access to transcripts. Some studies [11, 34, 5] explore transcripts of movies and TV series to help identify characters. A similar idea is proposed in [21] for action recognition. The character names are provided in advance, but in our setting, we assume that categories are unknown so that new objects can be discovered.

3. Wildlife Documentaries (WLD) Dataset

Right from the beginning of this study we ruled out videos with relatively clean separation of objects and background for our study of unknown objects discovery. Instead, we wish to mine for more meaningful and richer information from complex videos. To facilitate our study, we collect a new dataset called Wildlife Documentaries (WLD) dataset. It contains 15 documentary films that are downloaded from YouTube. The videos vary between 9 minutes to as long as 50 minutes, with resolution ranging from 360p to 1080p. A unique property of this dataset is that all videos are accompanied with subtitles that are automatically generated from speech by YouTube. The subtitles are revised manually to correct obvious spelling mistakes. To facilitate evaluations, we annotate all the animals in the videos resulting in more than **4098** object tracklets of **60** different visual concepts, e.g., ‘tiger’, ‘koala’, ‘langur’, and ‘ostrich’. We show some examples in Fig. 1.

The WLD dataset differs from conventional object detection datasets in that it is mainly designed to evaluate an algorithm’s capability in *discovering object of unknown classes* given with rich but ambiguous visual and linguistic information in videos. The videos are much longer and are left as they are without manual editing, while existing datasets usually provide short video snippets. The total frame number is more than **747,000**. Object bounding box annotations are not designated for model training, but provided to evaluate how accurate an algorithm could associate the object tubes with the right visual concepts.

The dataset is challenging in two aspects. Since videos are long documentaries, large variation in scale, occlusion

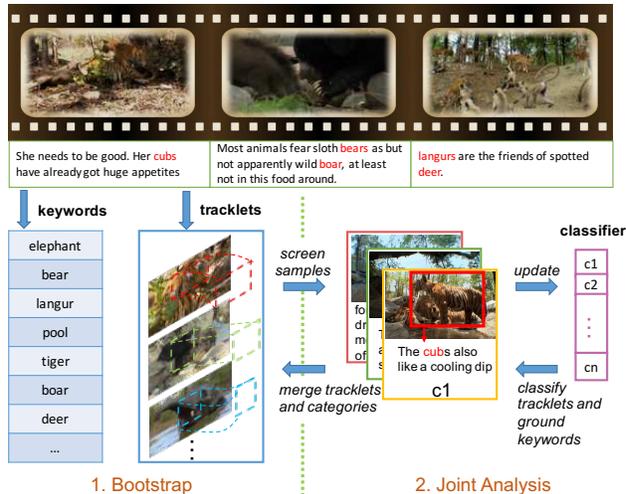


Figure 2: The proposed framework learns object detectors from documentary videos in an unsupervised manner. This is made possible through integrating noisy pieces of information, including visual and linguistic cues.

and background clutters is common. In many cases multiple objects co-exist in a frame. This adds difficulty in associating a target object with the correct nouns. Besides the visual challenges, highly unstructured subtitles also add to complexity. As can be seen from Fig. 1, meaningful nouns are overwhelmed by abundant of pronouns and synonyms. The occurrence of a noun does not necessarily imply the presence of the corresponding object due to temporal distance between object and subtitle. It is possible that a correspondence do not occur at all.

4. Framework Overview

Our primary goal in this work is to develop a framework that can *discover* new objects from a documentary video, and *learn* visual detectors therefrom. Note that the framework knows nothing about the given video a priori. The video itself and the associated subtitles, are the only data that it can rely on in the analysis.

As shown in Figure 2, our framework accomplishes this task through two major stages, namely the *bootstrap* stage and the *joint analysis* stage. The *bootstrap* stage is to prepare the materials for analysis. Specifically, it will acquire a collection of *tracklets* using a class-agnostic detector and tracker, and extract a list of *keywords* via low-level linguistic analysis. The *joint analysis* stage that ensues aims to assemble the tracklets and keywords into a coherent and meaningful picture. This analysis is devised on top of a CRF formulation, which unifies several tasks into an iterative optimization procedure. At each iteration, it selects a subset of *confident* samples to train or fine-tune the classifier for each object category, employs the updated classifiers to recognize the classes of individual tracklets, grounds the

keywords to them accordingly, and merges isolated tracklets based on the updated knowledge. In this way, object classifiers can be gradually improved with a growing and refined pool of training samples.

5. Bootstrap

The bootstrap stage is to prepare the inputs for joint analysis, which includes *candidate tracklets* and a list of *keywords* extracted from subtitles.

Obtain candidate tracklets. We employ an simplified version of the method proposed by Kang *et al.* [17] to generate candidate tracklets, without performing the actual detection task. Specifically, we first use a Region Proposal Network (RPN) [35] to generate class-agnostic object proposals from each video frame, and then apply a CNN-based classifier to evaluate the *objectness* score of each proposed bounding box. Then, we extend each proposal with high score to a *tracklet* via tracking. Note that we re-train the RPN by excluding any object of interests in the WLD dataset to make sure our approach is agnostic to the categories. The CNN for evaluating objectness provides an appearance feature for each bounding box, which will be retained and used in joint analysis.

Select keywords. Our linguistic analysis takes four steps to select keywords from subtitles:

(1) *Coreference resolution.* Pronouns are ubiquitous in typical narratives. To address this, we first use the method described in [26] to resolve the correspondences between pronouns and nouns and substitute all pronouns with the corresponding nouns, *e.g.* from “his paw” to “tiger’s paw”.

(2) *POS tagging.* Object categories are usually nouns. To identify them, we apply *Part-Of-Speech (POS)* tagging using a parser [4].

(3) *Lemmatization.* In English, a word can appear in a sentence in different forms, *e.g.* “bear”, “bears”, and “bear’s”. This step is to replace different forms of a word with its canonical form, so they will be considered as the same in later analysis.

(4) *Selection.* Finally, we want to select a subset of nouns as *keywords* that may indicate object categories. To avoid common words such as “animal” and “food”, we resort to *tf-idf scores* [39], motivated by the rationale that relevant keywords should frequently appear in only a small number of documents. To obtain reliable *tf-idf scores*, we collect a corpus that contains 2000 documentary transcripts and compute the scores thereon. We found empirically that the procedure outlined above can provide a good list of keywords that cover most of the animals appearing in our dataset.

6. Joint Analysis

Following the *bootstrap* stage, the framework will perform joint analysis based on a probabilistic model – clas-

sify each tracklet, infer the associations between tracklets and keywords, and as well obtain new object classifiers.

6.1. Task Statement

Specifically, given a video, the inputs at this stage include two parts: (1) **Tracklets.** We denote all tracklets as $\mathcal{T} = \{\tau_1, \dots, \tau_n\}$. Each tracklet τ_i is a sequence of bounding boxes, and can be described by a *visual feature* \mathbf{v}_i and a *geometric feature* \mathbf{u}_i . The former is formed by the appearance features extracted from a sub-sequence of frames, as $\mathbf{v}_i = (\mathbf{v}_i^{(1)}, \dots, \mathbf{v}_i^{(l_i)})$; while the latter captures the spatial/temporal characteristics of the bounding boxes. (2) **Keywords.** We consider each subtitle as a *bag of keywords* for simplicity. Putting all subtitles together, we can thus obtain a large collection of keywords, denoted by $\mathcal{W} = \{w_1, \dots, w_m\}$. Each keyword has a time span, which is the same as that of the parent subtitle.

The purpose of this *joint analysis* stage is to accomplish three key tasks: (1) **Categorization.** An important goal of this work is to detect objects from a given video. This is accomplished by assigning a category label $z_i \in \mathcal{C}$ to each candidate tracklet τ_i . Here, \mathcal{C} is the set of all categories, including all object categories and a *background* category with label 0. (2) **Grounding.** Along with a subtitle, multiple tracklets may appear in the scene. To bridge the visual and the linguistic domains, we need to *ground* the keywords to individual tracklets, *i.e.* determine which keywords correspond to which tracklets. Generally, a keyword may be grounded to zero or multiple tracklets, and vice versa a tracklet may be referred to by multiple keywords. Here, we use $a_{ij} \in \{0, 1\}$ to indicate whether the tracklet τ_i is associated with the keyword w_j . (3) **Classifier Learning.** The detected tracklets with their inferred labels constitute a training set on which object classifiers can be learned. Specifically, we can select a *confident* subset of tracklets classified to each object category and train a classifier thereon.

6.2. Probabilistic Formulation

In this work, we propose a Conditional Random Field (CRF) that unifies all these tasks into a probabilistic formulation, which comprises the following potentials:

Appearance potential $\psi_{ap}(z_i|\mathbf{v}_i; \boldsymbol{\theta})$: This potential measures how well a tracklet τ_i matches a object category z_i based on its appearance feature \mathbf{v}_i . It is defined as

$$\psi_{ap}(z_i|\mathbf{v}_i, \boldsymbol{\theta}) = \sum_{t=1}^{l_i} \log p(z_i|\mathbf{v}_i^{(t)}; \boldsymbol{\theta}). \quad (1)$$

When a convolutional network is used, $p(z|\mathbf{v}; \boldsymbol{\theta})$ is simply the logarithm of the output of the softmax layer, and the parameters $\boldsymbol{\theta}$ are the network weights.

Keyword-tracklet potential $\phi_{kt}(z_i, a_{ij}|\boldsymbol{\eta})$: As mentioned, a_{ij} indicates whether the tracklet τ_i associates with the key-

word w_j . The value of this potential is determined by the object category z_i , as

$$\phi_{kt}(z_i, a_{ij}|\boldsymbol{\eta}) = \log p(z_i|w_j; \boldsymbol{\eta}). \quad (2)$$

Each object category can have multiple keywords, *e.g.* class *tiger* have keywords “*tiger*” and “*cub*”. Here, $p(z_i|w_j; \boldsymbol{\eta})$ is the probability that the keyword w_j belongs to the class z_i , and the parameter $\boldsymbol{\eta}$ is the conditional probability table. A restriction is enforced here – each keyword can only be grounded to a tracklet whose time span overlaps with its own. In other words, a_{ij} is forced to be zero when the time spans of τ_i and w_j have no overlap.

Geometric potential $\phi_{st}(r_{ii'}, z_i, z_{i'}|\mathbf{u}_i, \mathbf{u}_{i'})$: Here, $r_{ii'}$ indicates whether tracklets τ_i and $\tau_{i'}$ are two consecutive segments of an object trajectory and thus should be merged. The value of ϕ_{st} is defined to be

$$\phi_{st} = \begin{cases} \delta(z_i = z_{i'}) \cdot s(\mathbf{u}_i, \mathbf{u}_{i'}) & (r_{ii'} = 1), \\ 0 & (r_{ii'} = 0). \end{cases} \quad (3)$$

Here, $s(\mathbf{u}_i, \mathbf{u}_{i'})$ is the spatial/temporal consistency. This definition ensures that two tracklets can only be merged when they have the same class label and are consistent in both space and time.

Joint model. The joint CRF can then be formalized as

$$p(\mathbf{z}, \mathbf{a}, \mathbf{r}|\mathbf{o}; \Theta) = \frac{1}{Z(\Theta)} \exp(\Psi_{ap}(\mathbf{z}|\mathbf{o}; \boldsymbol{\theta}) + \Phi_{kt}(\mathbf{z}, \mathbf{a}|\mathbf{o}; \boldsymbol{\eta}) + \Phi_{st}(\mathbf{r}, \mathbf{z}|\mathbf{o})). \quad (4)$$

Here, \mathbf{z} , \mathbf{a} , and \mathbf{r} are the vectors that respectively comprise all tracklet labels (z_i), keyword-tracklet association indicators (a_{ij}), and tracklet link indicators (r_{ij}). \mathbf{o} denotes all observed features, and Θ are model parameters. The three terms are given by

$$\Psi_{ap}(\mathbf{z}|\mathbf{o}; \boldsymbol{\theta}) = \sum_{i=1}^n \psi_{ap}(z_i|\mathbf{v}_i; \boldsymbol{\theta}), \quad (5)$$

$$\Phi_{kt}(\mathbf{z}, \mathbf{a}|\mathbf{o}; \boldsymbol{\eta}) = \sum_{(i,j) \in \mathcal{G}} \phi_{kt}(z_i, a_{ij}|\boldsymbol{\eta}), \quad (6)$$

$$\Phi_{st}(\mathbf{r}, \mathbf{z}|\mathbf{o}) = \sum_{(i,i') \in \mathcal{R}} \phi_{st}(r_{ii'}, z_i, z_{i'}|\mathbf{u}_i, \mathbf{u}_{i'}). \quad (7)$$

Here, \mathcal{G} is the set of tracklet-keyword pairs that can possibly be associated, *i.e.* their time spans overlap; \mathcal{R} is the set of all tracklet-tracklet pairs which can possibly be merged, *i.e.* spatial-temporal consistency is sufficiently high.

6.3. Joint Learning and Inference

Given a video, we employ *variational EM* to estimate the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$, and as well infer the latent variables, including (z_i) , (a_{ij}) and $(r_{ii'})$.

Initialization. To begin with, we form the initial set of classes by considering each distinct keyword as a class – such classes may be merged later as the algorithm proceeds. Also, to provide initial labels, we cluster all tracklets into a number of small groups based on their appearance using Mean Shift. Each cluster will be assigned a label according to the most frequent keyword among all those that overlap with the tracklets therein. Our experiments will show this heuristic method, while simple and not very accurate, does provide a reasonable initialization for joint analysis.

Iterative Optimization. The main algorithm is an iterative procedure that alternates between the following steps:

1. **Screen samples.** We observed detrimental impacts if all tracklets are fed to the training of classifiers, especially in the initial iteration, where the inferred labels can be very noisy. To tackle this issue, we explicitly enforce a screening mechanism, a confidence value will be computed for each tracklet based on a number of metrics, *e.g.* the length, the mean objectness score, the stability of the objectness scores, as well as the classification margin, *i.e.* the difference between the highest classification score and the runner-up. These metrics are combined using an SVM trained on a subset held out for tuning design parameters.
2. **Update classifiers.** With the confidence values, all tracklets whose confidence is beyond a certain threshold will be gathered to train or fine-tune the object classifiers of the corresponding categories. For each tracklet, we consider the class with highest score as its “true class” for the current iteration. The tracklets are sub-sampled with a fixed interval in order to save training time. The sampling interval is determined in a way such that the number of samples in all classes are relatively balanced. In addition, we also train a classifier for the *background class*, so as to enhance the contrast between foreground objects and the background.
3. **Classify tracklets.** With the updated classifier, we can infer the posterior probabilities of the class label z_i for each tracklet τ_i , denoted by q_i , as

$$q_i(z_i) \propto \exp\left(\psi_{ap}(z_i|\mathbf{v}_i, \boldsymbol{\theta}) + \sum_{j \in \mathcal{G}_i} \phi_{kt}(z_i, a_{ij}|\boldsymbol{\eta})\right). \quad (8)$$

Here, $\mathcal{G}_i = \{j : (i, j) \in \mathcal{G}\}$ is the set of keywords that overlap with τ_i . Here, the inference of the label z_i considers not only the appearance (1st term) but also the keywords associated with it (2nd term).

4. **Ground keywords.** First, each keyword w_j can only be grounded to those tracklets with overlapping time spans. For one of such tracklets, whether w_j should be grounded to τ_i depends only on the class label z_i . Particularly, the posterior probabilities of a_{ij} of given by

$p(a_{ij}|z_i) \propto \exp(\phi_{kt}(z_i, a_{ij}|\boldsymbol{\eta}))$. We found empirically that for most of the cases the probabilities are close to either 0 or 1. Hence, we simply sample a_{ij} therefrom to determine the grounding relations.

5. **Merge tracklets.** For each pair of tracklets τ_i and $\tau_{i'}$, we sample $r_{ii'}$ based on $\phi_{st}(r_{ii'}, z_i, z_{i'}|\mathbf{u}_i, \mathbf{u}_{i'})$. If $r_{ii'} = 1$, they will be merged into a new single tracklet. From the next iteration, the newly merged tracklet will be treated as a whole. Over time, the pool of isolated tracklets will be gradually consolidated into longer trajectories.
6. **Merge categories.** At the end of each iteration, we compute the similarity between each pair of categories using the *Earth Mover's Distance* [38], and merge similar classes (*i.e.* the distance below a threshold) into a single one. Along with this, the tracklet labels of the merged classes will be re-mapped to the new labels accordingly.

As we can see, multiple tasks are integrated in this procedure. Through iterative updates, a coherent picture over the given video will be gradually formed, where tracklets are assigned to object categories, keywords are grounded to relevant tracklets, and more importantly, a set of new object classifiers will be derived. Note that this is not a strict optimization algorithm, as two additional steps are inserted, which include *screen sampling* and *category merging*. While such steps are not directly derived from the CRF formulation, they do play significant roles in guiding the procedure towards a desirable direction.

7. Experiments

We evaluated our framework on the WLD dataset, comparing the results from different iterations and those from a supervised counterpart. We also studied the contributions of different components.

Settings. Recall that we use the RPN [35] for proposals generation and a CNN for objectness scoring and feature extraction in the bootstrap stage. The CNN is later used as a classifier in the joint analysis stage by adding fully-connected and softmax layers. For efficient iterative optimization during joint analysis, we only update the fully-connected layers and keep the convolutional layers fixed. We use ResNet-269 for the RPN and ResNet-101 for the CNN. Both models are trained using images excluding the object of interests in WLD.

Evaluation metrics. Since our approach is unsupervised, we evaluate our method on the whole WLD dataset and the primary metric is precision-recall. Some video detection tasks such as ImageNet VID evaluate the performance the same way as in object detection in images, regardless of the temporal accuracy and tracklet instance matching. On the contrary, we evaluate the results at the tracklet-level instead of the box-level, which can reflect the detection results more

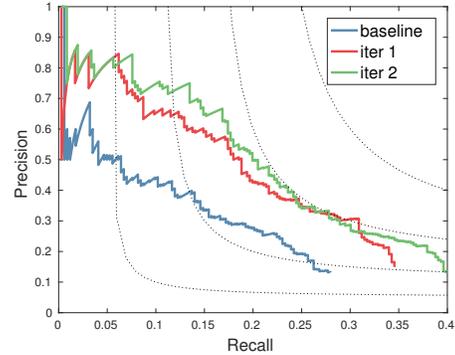


Figure 3: The overall performance of different iterations.

Table 1: Average precision (%) on a subset of WLD dataset.

	Iter 0	Iter 1	Iter 2
mAP	8.0	8.3	8.7
discovered mAP	20.6	28.9	30.7

intuitively and precisely. Specifically, the Intersection-over-Union (IoU) criterion is extended to spatio-temporal tracklets, the 3-D IoU of tracklet τ_p and τ_{gt} is calculated as $\frac{volume(\tau_p \cap \tau_{gt})}{volume(\tau_p \cup \tau_{gt})}$. We use 0.3 as the threshold in all of our experiments.

7.1. The Effectiveness of Iterative Optimization

Our learning method converges after 2 iterations and the results are shown in Fig. 3. The results suggest the effectiveness of our joint analysis step in incorporating noisy visual and linguistic information. It is observed that more iterations only give marginal improvements. This may due to that a majority of confident samples have been mined and utilized in the first two iterations, and little extra information can be found in further steps. We also measure the mAP of our results like other supervised detection methods, as shown in Table 1. Different from supervised and some weakly-supervised methods that have a prior of object categories, candidate categories are derived from subtle analysis in our framework, so some categories with only a few objects, *e.g.* *elephant*, may be hard for our framework to discover. Consequently, we propose *discovered mAP* besides the regular mAP, which means the mAP of all discovered categories. The results shows that object categories mentioned more often in the video are more likely to be discovered and learned.

7.2. Comparison with Supervised Method

We wish to examine further how our unsupervised trained detector compares against a strong fully supervised detector. To perform this experiment, we select categories that are available in both WLD and ImageNet. In total 7 classes are selected, which are available in a single docu-

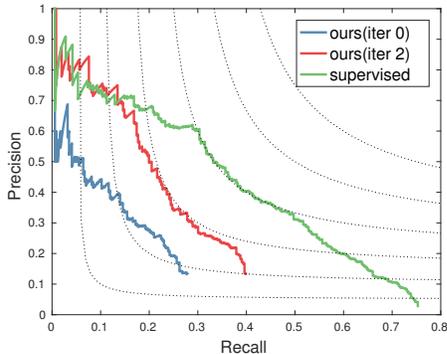


Figure 4: Comparison with the fully supervised method.

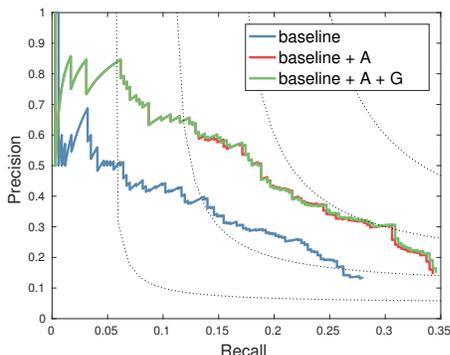


Figure 5: The results of introducing different potentials. We show the results of the first iteration in joint analysis.

mentary video. Since the fully supervised detector is trained using ImageNet object detection and localization dataset, it has prior knowledge on these selected categories, and its performance can be treated as an upper-bound of our unsupervised method. As shown in Fig. 4, after 2 iterations, our unsupervised detector is competitive to the supervised counterpart in high-precision region. The discovered mAP for the fully supervised method is 0.309, comparable to 0.307 of our method. It is also observed that our method has a comparable or even higher AP in major categories of the documentary, *e.g.*, *tiger* and *langur* in this case.

7.3. Ablation Study

To better understand how our joint model works, we study the necessity of appearance potential and geometric potential (keyword-tracklet potential is essential otherwise tracklet labels cannot be obtained). We compare the following results: (1) Baseline, (2) Baseline + Appearance, (3) Baseline + Appearance + Geometry, as shown in Fig. 5.

Appearance potential. From the results of (1) and (2), we can observe that the appearance potential contributes considerably to the performance increase. There are two possibilities for the gain: the classifier corrected the false negative samples back to positive, or the classifier disambiguated confusing labels across different foreground categories. To further examine the role of this potential, we

Table 2: The results before and after ‘classify tracklets’ and ‘ground keywords’ steps (see Sec.6.3 for details).

	F1-binary	F1-multiple	average top-3 size
before	0.409	0.240	96
after	0.367	0.305	277

Table 3: A comparison of our grounding approach with word counting baseline accuracy.

	Iter 0	Iter 1	Iter 2	Ground truth
Word counting	0.414	0.776	0.855	0.744
Ours	0.343	0.888	0.937	0.935

treat all the foreground categories as one class and calculate the F1 score of classifier prediction compared with the baseline. We also compute the average size of top 3 largest clusters. As shown in Table 2, despite the classifier cannot recall more foreground samples compared with the baseline (reflected by the F1-binary metric), it shows great improvement on grounding results (reflected by the F1-multiple metric). When we consider the improved F1-multiple metric and the increase of cluster size, the observation suggests that the appearance model captures meaningful visual similarity and strengthen some dominant categories.

Keyword-tracklet potential. In our grounding method, only tracklets that are predicted as positive can be associated with keywords and all negative ones are ignored. The positive tracklets can be noisy since the predicted labels may not be accurate. A good grounding strategy will associate correct words with true positive samples, as well as assign no word to false positive samples. As a baseline for grounding, we employ a straightforward approach which assigns category labels based on word frequencies. We compare the accuracy on true positive samples of two methods using class labels of different iterations as well as ground truth labels. The results are shown in Table 3, which clearly shows the significant contribution of this potential.

Geometric potential. Comparing (2) and (3), geometric potential also shows its effectiveness, although much weaker than the appearance potential. The effects of geometric link potential are of two aspects named “weak link” and “strong link”, which are decided by the potential value. “Weak link” refers to two tracklets (belonging to the same category) with weak geometric relation based on some threshold. “Strong link”, on the other hand, refers to two tracklets not only belong to the same category, but also strong geometric relation (which will be merged together). We compare the following variants: without geometric link, with weak geometric link, with both weak and strong geometric link, results are shown in Figure 6. “Weak link” contributes little compared with “strong link” (completely overlapped with baseline), suggesting that the merging strategy with strong link is the most effective part in geometric potential.

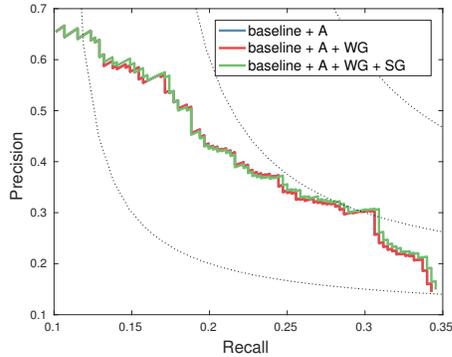


Figure 6: The results of weak and strong geometric link. “A”, “WG” and “SG” refer to “appearance”, “weak link” and “strong link” respectively.

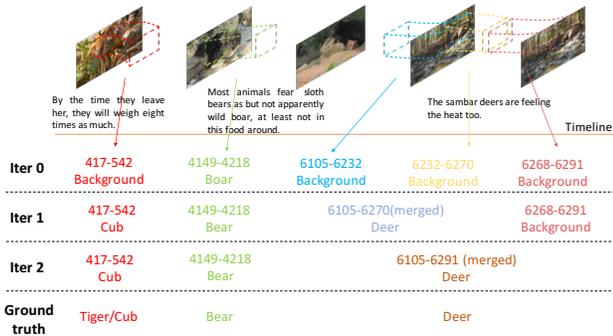


Figure 7: Tracklet examples in different iterations.

7.4. Examples and Failure Cases

We demonstrate some examples in Fig. 7, showing the results in different processes. The initialization (iter 0) is not satisfying for the some wrong labels, and the long trajectory is split into three short tracklet proposals. As the iteration proceeds, the learned appearance model classifies the tracklets into correct groups, and the strong geometric link merges separated tracklets into longer ones, which helps to obtain the correct labels in turn.

We also show some failure cases of our method. Figure 8 illustrates three typical types of failure: (1) contributed by the very ambiguous and challenging visual appearance of the object. These objects are either too small or heavily occluded by other object. (2) caused by the omission of keywords, *i.e.*, if the *tf-idf* score is below the threshold, the category is not included in the keyword list. (3) caused by mismatch in visual and linguistic evidences. Specifically, if two categories usually appear at the same time or the subtitles and frames are not synchronized, it is likely to be mismatched.



Figure 8: Examples of failure cases.

8. Conclusion

This paper presented a framework that can jointly discover object categories, detect object tracks, and learn object classifiers from documentary videos in an unsupervised fashion. Experimental results obtained on a real world datasets have demonstrated the effectiveness of this framework – most animal categories are discovered automatically. Also, as the joint analysis algorithm proceeds, the object classifiers will be gradually improved with a growing training set, raising the mAP from 8.0% to 8.7% and the discovered mAP from 20.6% to 30.7%. Our framework also present comparable results with fully supervised learning for the major categories of documentaries.

Acknowledgment This work is partially supported by the Early Career Scheme (ECS) grant (No. 24204215), and the collaboration grant from SenseTime Group. We would like to thank Dian Chen, Lingzhi Li, Miao Fan, Pik Kei Tse, Tsz Ching Fan, Yule Li, Zhehao Jiang for contributing to the annotation efforts.

References

- [1] J.-B. Alayrac, P. Bojanowski, N. Agrawal, J. Sivic, I. Laptev, and S. Lacoste-Julien. Unsupervised learning from narrated instruction videos. In *IEEE International Conference on Computer Vision*, pages 4575–4583, 2016. 2
- [2] S. Bagon, O. Brostovski, M. Galun, and M. Irani. Detecting and sketching the common. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2010. 2
- [3] H. Bilen, M. Pedersoli, and T. Tuytelaars. Weakly supervised object detection with convex clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2

- [4] S. Bird, E. Klein, and E. Loper. *Natural language processing with Python*. "O'Reilly Media, Inc.", 2009. 4
- [5] P. Bojanowski, F. Bach, I. Laptev, J. Ponce, C. Schmid, and J. Sivic. Finding actors and actions in movies. In *IEEE International Conference on Computer Vision*, pages 2280–2287, 2013. 3
- [6] P. Buehler, A. Zisserman, and M. Everingham. Learning sign language by watching tv (using weakly aligned subtitles). In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 3
- [7] R. G. Cinbis, J. Verbeek, and C. Schmid. Weakly supervised object localization with multi-fold multiple instance learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016. 2
- [8] H. Cooper and R. Bowden. Learning signs from subtitles: A weakly supervised approach to sign language recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 3
- [9] T. Deselaers, B. Alexe, and V. Ferrari. Weakly supervised localization and learning with generic knowledge. *International Journal of Computer Vision*, 100(3):275–293, 2012. 2
- [10] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [11] M. Everingham, J. Sivic, and A. Zisserman. "hello! my name is... Buffy"—automatic naming of characters in tv video. In *British Machine Vision Conference*, 2006. 3
- [12] R. Girshick. Fast r-cnn. In *IEEE International Conference on Computer Vision*, 2015. 1
- [13] J. Hoffman, D. Pathak, T. Darrell, and K. Saenko. Detector discovery in the wild: Joint multiple instance and representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [14] R. Hu, M. Rohrbach, and T. Darrell. Segmentation from natural language expressions. In *European Conference on Computer Vision*, 2016. 2
- [15] R. Hu, H. Xu, M. Rohrbach, J. Feng, K. Saenko, and T. Darrell. Natural language object retrieval. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [16] A. Joulin, K. Tang, and L. Fei-Fei. Efficient image and video co-localization with frank-wolfe algorithm. In *European Conference on Computer Vision*, 2014. 2
- [17] K. Kang, W. Ouyang, H. Li, and X. Wang. Object detection from video tubelets with convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 4
- [18] A. Karpathy, A. Joulin, and F.-F. Li. Deep fragment embeddings for bidirectional image sentence mapping. In *Advances in Neural Information Processing Systems*, 2014. 2
- [19] K. Kumar Singh, F. Xiao, and Y. Jae Lee. Track and transfer: Watching videos to simulate strong human supervision for weakly-supervised object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [20] S. Kwak, M. Cho, I. Laptev, J. Ponce, and C. Schmid. Un-supervised object discovery and tracking in video collections. In *IEEE International Conference on Computer Vision*, 2015. 2
- [21] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2008. 3
- [22] X. Liang, S. Liu, Y. Wei, L. Liu, L. Lin, and S. Yan. Towards computational baby learning: A weakly-supervised approach for object detection. In *IEEE International Conference on Computer Vision*, 2015. 2
- [23] J. Mao, J. Huang, A. Toshev, O. Camburu, A. Yuille, and K. Murphy. Generation and comprehension of unambiguous object descriptions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [24] J. Mao, W. Xu, Y. Yang, J. Wang, Z. Huang, and A. Yuille. Deep captioning with multimodal recurrent neural networks (m-RNN). 2015. 2
- [25] O. Maron and T. Lozano-Pérez. A framework for multiple-instance learning. *Advances in Neural Information Processing Systems*, 1998. 2
- [26] S. Martschat and M. Strube. Latent structures for coreference resolution. *Transactions of the Association for Computational Linguistics*, 3:405–418, 2015. 4
- [27] I. Misra, A. Shrivastava, and M. Hebert. Watch and learn: Semi-supervised learning for object detectors from video. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [28] M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Is object localization for free?-weakly-supervised learning with convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [29] W. Ouyang, X. Wang, X. Zeng, S. Qiu, P. Luo, Y. Tian, H. Li, S. Yang, Z. Wang, C.-C. Loy, et al. Deepid-net: Deformable deep convolutional neural networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2403–2412, 2015. 1
- [30] M. Pandey and S. Lazebnik. Scene recognition and weakly supervised object localization with deformable part-based models. In *IEEE International Conference on Computer Vision*, pages 1307–1314, 2011. 2
- [31] T. Pfister, J. Charles, and A. Zisserman. Domain-adaptive discriminative one-shot learning of gestures. In *European Conference on Computer Vision*, 2014. 3
- [32] B. A. Plummer, L. Wang, C. M. Cervantes, J. C. Caicedo, J. Hockenmaier, and S. Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *IEEE International Conference on Computer Vision*, 2015. 2
- [33] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Learning object class detectors from weakly annotated video. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012. 2
- [34] V. Ramanathan, A. Joulin, P. Liang, and L. Fei-Fei. Linking people in videos with "their" names using coreference resolution. In *European Conference on Computer Vision*, pages 95–110. Springer, 2014. 3

- [35] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, 2015. [1](#), [4](#), [6](#)
- [36] W. Ren, K. Huang, D. Tao, and T. Tan. Weakly supervised large scale object localization with multiple instance learning and bag splitting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):405–416, 2016. [2](#)
- [37] A. Rohrbach, M. Rohrbach, R. Hu, T. Darrell, and B. Schiele. Grounding of textual phrases in images by reconstruction. In *European Conference on Computer Vision*, 2016. [2](#)
- [38] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, 2000. [6](#)
- [39] G. Salton and M. J. McGill. Introduction to modern information retrieval. 1986. [4](#)
- [40] Z. Shi, T. M. Hospedales, and T. Xiang. Bayesian joint topic modelling for weakly supervised object localisation. In *IEEE International Conference on Computer Vision*, 2013. [2](#)
- [41] P. Siva, C. Russell, and T. Xiang. In defence of negative mining for annotating weakly labelled data. In *European Conference on Computer Vision*, 2012. [2](#)
- [42] H. O. Song, R. B. Girshick, S. Jegelka, J. Mairal, Z. Harchaoui, T. Darrell, et al. On learning to localize objects with minimal supervision. In *International Conference on Machine Learning*, 2014. [2](#)
- [43] H. O. Song, Y. J. Lee, S. Jegelka, and T. Darrell. Weakly-supervised discovery of visual pattern configurations. In *Advances in Neural Information Processing Systems*, 2014. [2](#)
- [44] C. Wang, W. Ren, K. Huang, and T. Tan. Weakly supervised object localization with latent category learning. In *European Conference on Computer Vision*, 2014. [2](#)
- [45] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Object detectors emerge in deep scene cnns. *arXiv preprint arXiv:1412.6856*, 2014. [2](#)