

# Zero-shot Object Prediction using Semantic Scene Knowledge

Rene Grzeszick and Gernot A. Fink

*Department of Computer Science, TU Dortmund University, Dortmund, Germany*  
{rene.grzeszick, gernot.fink}@tu-dortmund.de

**Keywords:** Zero-shot, Object Prediction, Scene Classification, Semantic Knowledge.

**Abstract:** This work focuses on the semantic relations between scenes and objects for visual object recognition. Semantic knowledge can be a powerful source of information especially in scenarios with few or no annotated training samples. These scenarios are referred to as zero-shot or few-shot recognition and often build on visual attributes. Here, instead of relying on various visual attributes, a more direct way is pursued: after recognizing the scene that is depicted in an image, semantic relations between scenes and objects are used for predicting the presence of objects in an unsupervised manner. Most importantly, relations between scenes and objects can easily be obtained from external sources such as large scale text corpora from the web and, therefore, do not require tremendous manual labeling efforts. It will be shown that in cluttered scenes, where visual recognition is difficult, scene knowledge is an important cue for predicting objects.

## 1 INTRODUCTION

Much progress has been made in the field of image classification and object detection, yielding impressive results in terms of visual analysis. Latest results show that up to a thousand categories and more can be learned based on labeled instances (Simonyan and Zisserman, 2014). In comparison, it is estimated that humans recognize about 30.000 visual categories and even more sub-categories such as car brands or animal breeds (Palmer, 1999). Human learning is different from machine learning, although it can be based on visual examples, it is also based on external knowledge such as descriptions of entities or the context in which they appear. Recognition systems often omit basic knowledge on a descriptive level. This work will focus on the semantic relations between scenes and objects which can be an important cue for predicting objects. This is especially useful in cluttered scenes where visual recognition may be difficult. Knowing the scene context, which yields a strong prior on what to expect in a given image, the presence of objects in an image will be predicted.

It is known that contextual information can help in the task of recognizing objects (Divvala et al., 2009; Choi et al., 2012; Zhu et al., 2015). Various forms of context that can improve visual recognition tasks have already been investigated in (Divvala et al., 2009). Visual context can be obtained in a very local manner such as pixel context or in a global manner by image

descriptors like the Gist of a scene (Oliva and Torralba, 2006). Another form of visual context is the presence, appearance or location of different objects in a scene. External context cues are, for example, of photogrammetric, cultural, geographic or semantic nature (Divvala et al., 2009).

Especially the object level approaches that define context based on the dependencies and co-occurrences of different objects are pursued in several works (Felzenszwalb et al., 2010; Choi et al., 2012; Vezhnevets and Ferrari, 2015). In (Felzenszwalb et al., 2010) a stacked SVM classifier is applied that uses the maximal detection scores for each object category in an image in order to re-rank the prediction scores. The work presented in (Choi et al., 2012) uses a hierarchical tree structure in order to model the occurrence of objects as well as a spatial prior. In (Vezhnevets and Ferrari, 2015) the detection scores of a specific bounding box are re-ranked based on its position in the scene as well as the relative position of other bounding boxes. More global approaches use image level context definitions for object detection or image parsing (Liu et al., 2009; Tighe and Lazebnik, 2010; Modolo et al., 2015). Most of these context definitions follow the same approach: They retrieve a subset of training images which are similar to the given image and transfer the object information (Liu et al., 2009; Tighe and Lazebnik, 2010). A slightly different approach is pursued in (Modolo et al., 2015), where a Context Forest is trained that

learns the relation between a global image descriptor and the objects within this image. This allows to efficiently find related images based on the forest's leaf nodes and then transfer assumptions about object locations or classes. Most of these works have in common that a considerable effort went into training a state-of-the-art detector and the results of various detections are combined in order to obtain a context descriptor. In (Felzenszwalb et al., 2010; Choi et al., 2012; Modolo et al., 2015) these were deformable part based models that build on HOG features. Nowadays, Convolutional Neural Networks (CNNs), like very deep CNNs (Simonyan and Zisserman, 2014) and R-CNNs (Girshick et al., 2016; Ren et al., 2015) show state-of-the-art performance in object prediction and object detection respectively. While methods like data augmentation and pre-training have reduced the required number of samples and weakly supervised annotation schemes lower the required level of detail, still a considerable annotation effort is required to train these models.

A different idea is adding further modalities for context. The most prominent of these modalities is text, for example, image captions or additional tags (Lin et al., 2014). These multi-modal approaches allow for answering visual queries (Zhu et al., 2015; Wu et al., 2016; Krishna et al., 2016), the captioning of images or videos (Rohrbach et al., 2013; Donahue et al., 2015) and recognition with limited training samples supported by additional linguistic knowledge (Rohrbach et al., 2011; Lampert et al., 2014). Several of these approaches incorporate additional attributes that allow for transferring knowledge without explicitly annotating a specific class or object label (Rohrbach et al., 2011; Lampert et al., 2014; Zhu et al., 2015). For example, instead of recognizing an animal, visual attributes like its color, whether it has stripes or is shown in water are recognized. These attributes can then be used as additional visual cues.

In (Patterson et al., 2014) a database that focuses on scene attributes is introduced. Each scene is described based on its visual attributes such as natural or manmade. Attributes are predicted independently of each other using one SVM per attribute. Following up on the idea of attribute prediction, it has been shown that a combined prediction of these attributes can be beneficial as they are often correlated. In (Grzeszick et al., 2016) a neural network is trained that predicts multiple attributes simultaneously and outperforms the traditional per class SVMs on scene attributes.

In (Zhu et al., 2015) a knowledge base system is built that builds on a similar idea and relates scenes with attributes and affordances. It is shown that the

association of scenes with attributes and affordances allows for improving the predictions of scenes and their attributes as well as answering visual queries. An even more complex system is presented in the Visual Genome (Krishna et al., 2016) where a complete scene graph of objects, attributes and their relations with each other is presented. This allows not only for predicting attributes and relations, but also for answering complex visual questions. Textual queries are parsed with respect to attributes so that the best matching images can be retrieved.

In (Lampert et al., 2014) attributes, which are associated with a set of images or classes, are used for uncovering unknown classes and describing them in terms of their attributes. For example, an animal with the attributes *black*, *white* and *stripes* will most likely be a zebra. Each of the attributes is recognized independently and without any knowledge about the actual object classes. The attribute vector is then used in order to infer knowledge about object classes. Such methods with no training samples for given objects are also referred to as zero-shot learning approaches. Similarly, given a very small set of training samples, attributes can be used in order to transfer class labels to unknown images (Rohrbach et al., 2011).

Attributes for images can either be learned directly via annotated training images or indirectly via additional sources of information such as Wikipedia or WordNet (Lampert et al., 2014; Rohrbach et al., 2011). As annotating images with attributes is tedious, especially, the latter allows for scaling recognizers to a larger number of classes and attributes. Furthermore, it has been shown that these attributes can also be derived in a hierarchical manner, i.e., based on the WordNet tree (Rohrbach et al., 2011). An important factor for incorporating such additional linguistic sources is the vast amount of text corpora that are available on the web. Information extraction systems like TextRunner (Banko et al., 2007) or Reverb (Etzioni et al., 2011) allow for analyzing these text sources and uncovering information, like nouns and the relations between them.

This work will show that additional textual information is beneficial for predicting object presence in an image with minimal annotation effort similar to (Lampert et al., 2014; Rohrbach et al., 2011; Zhu et al., 2015). Here, instead of attributes, a more direct way is proposed exploiting the semantic relations between scenes and objects in a zero-shot approach. The presence of an object is predicted based on two sources: visual knowledge about the scene and the relations between scenes and different objects. For example, a car can hardly be observed in the livingroom or a dining table in the garage. Such knowledge can

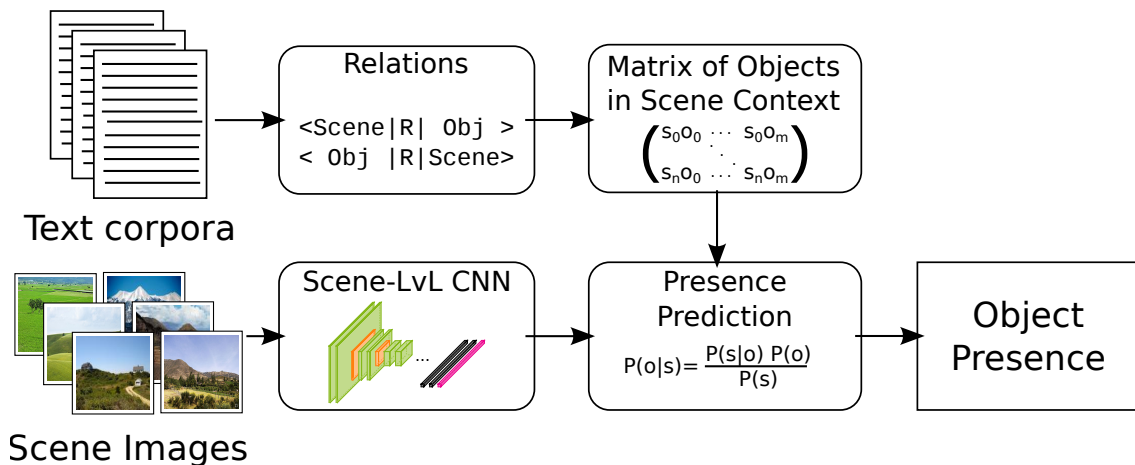


Figure 1: Given external text sources, these are analyzed with respect to possible scenes, objects and their relations, creating a matrix of objects in scene context. For a given set of images, a CNN is trained in order to predict scene labels. This information is then used in order to predict the presence of an object in scene.

easily be obtained from additional textual sources using methods like TextRunner (Banko et al., 2007) or Reverb (Etzioni et al., 2011). As a result the tremendous annotation effort that is required for annotating objects in images is no longer necessary. Similar to existing zero-shot approaches, it requires only a descriptive label, the scene name, and works completely unsupervised with respect to annotations of objects. In the experiments it will be shown that such high level knowledge allows for predicting the presence of objects, especially in very cluttered scenes.

## 2 METHOD

In the proposed method for object prediction, the relations between scenes and objects which are obtained from text sources are used for modeling top-down knowledge. They replace the visual information that is typically used for object prediction. An overview is given in Fig. 1. The image is solely analyzed on scene level, which requires minimal annotation effort and no visual knowledge about the objects within the scene. More importantly, a text corpus is analyzed with respect to possible scenes and objects, extracting the relations between them and creating a matrix of objects in a scene context. This information is then used in combination with the scene classification in order to predict the presence of an object in an image.

### 2.1 Relations between Objects and Scenes

Given a large enough set of text, relations extracted from these texts can be assumed to roughly repre-

sent relations that are observed in the real world and henceforth may also be observed on images. Rich text corpora can, for example, be obtained by crawling Wikipedia or any other source of textual information from the web (Rohrbach et al., 2011). Here, sentences including possible scene or object categories and their relations will be of further interest. In the following extractions based on Reverb (Etzioni et al., 2011) are used. Reverb extracts relations and their arguments from a given sentence. Therefore, two steps are performed: First, for each verb  $v$  in the sentence, the longest sequence of words  $r_v$  is uncovered so that  $r_v$  starts at  $v$  and satisfies both a syntactical and a lexical constraint. For the lexical constraint, Reverb uses a dictionary of 1.7 million relations. The syntactical constraint is based on the following regular expression:

$$V|VP|VW^*P$$

$$V = \text{verb particle? adv?}$$

$$W = (\text{noun|adj|adv|pron|det})$$

$$P = (\text{prep|particle|inf. marker}) \quad (1)$$

Overlapping sets of relations are merged into a single relation. Second, given a relation  $r$ , the nearest noun left and right of the relation  $r$  are extracted. If two nouns can be observed for a relation  $r$ , this results in a triplet

$$t = (\text{arg1}, r, \text{arg2}) . \quad (2)$$

Simple examples relating scenes and objects may be 'A car drove down the street' or 'Many persons were on the streets'. Both relate an entity that is typically used in object detection (*car* or *person*) with a scene label (*street*). In this example, the triplets

$$t_0 = (\text{car}, \text{drove down}, \text{street}) ,$$

$$t_1 = (\text{person}, \text{were on}, \text{street})$$

would be extracted from the two sentences.

Given a set of relations that were extracted from a text corpus and a vocabulary that defines a set of  $S$  scene names and  $O$  object names, a matrix  $C$  describing the objects in scene context is created. In a general setting a vocabulary could be derived from frequently occurring words in a text corpus, the WordNet tree (Miller and Others, 1995) or just a set of objects and scenes that are of interest and known beforehand. At the index  $s, o$  the matrix  $C$  contains the number of relations between the respective scene  $s$  and object  $o$ :

$$C_{s,o} = \sum_r n(o, s) + \sum_r n(s, o) \quad \text{with} \quad (3)$$

$$n(i, j) = \#\{t = (i, \cdot, j)\} \quad . \quad (4)$$

Hence, the type of relation is discarded as only the number of relations between an object and a scene will be of further interest. In the experiments, some rare cases of self-similarity were observed, in which a scene and an object name are the same (e.g. a scene in a street may also show the object road/street, among others). In these cases the self-similarity is set to the maximum count observed.

## 2.2 Presence Prediction

The task of presence detection is concerned with the question whether an object can be observed one or more times in a given image. Under the assumption that a large diverse text corpus is representative for the real world, it can further be assumed that the likelihood of an object to occur in a given scene is correlated with the number of textual relations between those entities. Since multiple objects can occur in a single scene image, the presence predictions for different object classes are typically evaluated independently of each other. Due to this multiplicity, the probability  $P(o|s)$  of the object  $o$  to be shown in the scene  $s$  cannot be computed directly using the counts. However, since an image can only depict one scene,  $P(s|o)$  can be estimated from the counts. This allows to compute  $P(o|s)$  based on Bayes theorem as

$$P(o|s) = \frac{P(s|o)P(o)}{P(s)} \quad \text{with}$$

$$P(s|o) = \frac{C_{s,o}}{\sum_{s'} C_{s',o}} \quad \text{and} \quad P(o) = \frac{\sum_s C_{s,o}}{\sum_{s'} \sum_{o'} C_{s',o'}} \quad . \quad (5)$$

The prior probability  $P(o)$  can be approximated, assuming that one relation count represents the presence of at least one object  $o$ .

However, for the same reason  $P(o|s)$  cannot be derived from the matrix of objects in scene context, the prior probability  $P(s)$  for a certain scene cannot be derived. Assume the Matrix  $C$  contains  $N$  relation

counts which relate the presence of at least  $N$  objects from  $O$  categories to the set of  $S$  scenes. Given the one to many relation between a scene and objects, the true count of scenes cannot be recovered. Therefore,  $P(s)$  is assumed to be uniformly distributed.

In order to be able to predict an object in a scene where no relations have been previously observed, unobserved events need to be handled. Therefore, the probability of an object  $o$  to occur in a scene  $s$  is smoothed by

$$P^*(o|s) = (1 - \alpha) P(o|s) + \alpha P(o), \quad (6)$$

similar to the smoothing of probability distributions for statistical natural language processing (cf. (Manning and Schütze, 1999)). The process is based on an interpolation factor  $\alpha$  which is estimated based on the number of relations with only a single occurrence:

$$\alpha = \frac{\#\{C_{o,s} | C_{o,s} = 1\}}{\sum_{o'} \sum_{s'} C_{o',s'}} \quad (7)$$

Furthermore, the counts are obtained from a text source that is unrelated to the visual tasks so that there is a remaining degree of uncertainty. The matrix representation does also not cover intra-scene variability (i.e. all scenes would have exactly the same relation to objects). In order to model these two issues,  $\hat{P}(o|s)$  is sampled by  $D$  draws from a normal distribution

$$\hat{P}(o|s) = \frac{1}{D} \sum_1^D n \quad \text{with} \quad n \sim \mathcal{N}(P^*(o|s), \sigma(C)) \quad . \quad (8)$$

The variance  $\sigma$  is estimated based on the variance within the matrix  $C$ . In order to estimate the probability of an object in a given image  $I$ , the presence is then predicted by:

$$P(o|I) = \sum_s P(s|I) \cdot \hat{P}(o|s) \quad . \quad (9)$$

The probability  $P(s|I)$  can be predicted by a classifier. Assuming a perfect classification  $P(s|I)$  would equal to one for the true scene  $s$  and the probability  $P(o|I)$  would solely be computed based on the relations between this scene and the objects. However, in practice there will be a distribution over a set of scene labels. This also takes into account the ambiguity between different scenes. In this work, a CNN that is based on a VGG16 network architecture is used for predicting the scene category (Simonyan and Zisserman, 2014). The network is pre-trained on ImageNet. It is then adapted to the scene classification using a set of scene images depicting  $S$  scenes categories.

Note that the requirement for a single scene label is very easy to fulfill. The annotation effort for a single scene label is much lower than labeling various

object classes in an image or even annotating the position of an object in a scene which has to be done for most supervised object detectors. Similar to attribute-based zero-shot learning, it is a descriptive abstraction that does not imply any visual knowledge about the objects within the scene.

### 3 EVALUATION

In the following the experimental setup and the evaluation of the proposed object prediction are described. Ideally, the evaluation requires a dataset that offers both scene and object labels. Hence, the different branches of the SUN dataset (Xiao et al., 2014) have been chosen for the evaluation: The SUN2012 Scene and Object Dataset in Pascal VOC format and the SUN2009 Context dataset. Both branches of the SUN dataset show a broad set of different scene and object categories.

**SUN2012 Scene and Object Dataset in Pascal VOC format:** the dataset contains a set of images taken from the SUN image corpus. While the more prominent SUN397 dataset is annotated with 397 different scenes labels, this set contains annotations for an additional 4,919 different object classes (Xiao et al., 2014). Of all 16,873 images, 11,426 are a subset of the SUN397 dataset for which both annotations, scene and object labels, are available. For the remaining 5,447 no scene annotations are provided.

**SUN2009 Context:** the dataset contains only about 200 different object categories, of which 107 were used for supervised detection experiments in (Choi et al., 2012). The same diversity as for the SUN2012 Scene and Object dataset can be observed with respect to the scene and object categories.

In contrast to traditional object detection tasks, like the Pascal VOC challenge (Everingham et al., 2015), there is a great variability with respect to the objects properties. While some of them are well defined (e.g. cars, person), some others describe regions (sky, road, buildings) or highly deformable objects (river, curtain). Moreover, the annotations in all versions of the SUN dataset are very noisy. Some of them contain descriptive attributes, like *person walking*, *table occluded*, *tennis court outdoor* others mix singular and plural.

In order to relate the scene and object labels with natural text, these descriptive attributes were removed and all objects and scene labels were lemmatized based on the WordNet tree (Miller and Others, 1995). This leaves 3,390 unique object labels in 377 different scenes labels. Although the lemmatized scene names may be semantically similar, they may be vi-

sually different (i.e. for *tennis court indoor* and *outdoor*). Therefore, the scene labels for all 397 labels will be predicted and the prediction results will be summarized. The objects are then predicted based on the lemmatized names.

#### 3.1 Creating a Matrix of Objects in Scene Context

In order to obtain a matrix of objects in scene context, the OpenIE database has been queried. It contains over 5 billion extractions that have been obtained using Reverb on over a billion web pages<sup>1</sup>. Hence, a very diverse dataset that captures the relations between a huge set of nouns has been used. The vocabulary has been defined based on the task of the SUN2012 Scene and Object dataset so that the vocabulary consists of the  $S = 377$  scene names and  $O = 3390$  object names. All possible combinations of scenes and objects were queried for which a total relation count of 1,375,559 has been extracted<sup>2</sup>. Note that the distribution of these relation counts is very long tailed, leaving a large set of unobserved events.

#### 3.2 Scene Prediction

For recognizing objects based on the scene context, the probability of the given image to depict the scene  $s$  needs to be computed. In the following, two different setups are evaluated.

**Perfect Classifier:** it is assumed that the scene label is known beforehand (i.e. given by a human in the loop) or that training a perfect scene classifier with respect to the annotated scene labels would be possible. In order to simulate this case,  $P(s|I)$  is set to 1 for the annotated scene label and to 0 otherwise. Note that these labels might be ambiguous and even human annotators deviate in their decision from the ground truth labels (Xiao et al., 2010).

**Scene-level CNN:** For recognizing scene labels a CNN is evaluated. A VGG16 network architecture has been pre-trained on ImageNet. It has then been adapted to the task of scene classification using all scene images from the SUN397 dataset that are not included in the SUN2012 Scene and Object dataset. The exclusion of the images from the SUN2012 Scene

<sup>1</sup>For a demo see <http://openie.allenliu.org>

<sup>2</sup>The Matrix of Objects in Scene Context will be made publicly available together with the detailed experimental setup containing the training/test split and the lemmatized annotations at <http://patrec.cs.tu-dortmund.de/cms/en/home/Resources/index.html>.

Table 1: Recognition rate for the  $k$  highest scoring prediction of the scene label using a CNN.

k-best	Recognition rate
1	62.2%
3	82.8%
5	88.9%

and Object dataset leaves a set of 97,304 training images. The training images have been augmented using random translations (0 – 5%), flipping (50% chance) and Gaussian noise ( $\sigma = 0.02$ ) in order to achieve a better generalization. In total 500,000 training images have been created. The learning rate has been set to  $\alpha = 0.0001$  using 25,000 training iterations of batch size 39 (= 10% of the number of classes).

The CNN recognition rates for the  $k$  highest scoring predictions on the SUN2012 Scene and Object dataset are shown in Tab. 1. The highest scoring prediction yields an accuracy of 62.2%. When considering that the correct result must be within the five highest scoring predictions an accuracy of 88.9% is achieved. This emphasizes that the probabilistic assignment to a set of scene categories based on the CNNs predictions is a meaningful input for the proposed object prediction.

### 3.3 Object Prediction

In the following experiments, the presence detection is evaluated. Note that only scene labels were used for training the CNN so that it is completely unsupervised with respect to object occurrences.

#### SUN 2012 Scene and Object Dataset

For the SUN2012 Scene and Object dataset, the most frequently occurring 100 up to all 3390 objects categories from the dataset were considered, some of them being comparably rare. The accuracy of the top  $k$  predictions is evaluated. Note that multiple objects can occur in a single image and, therefore, the accuracy of all predictions is evaluated (i.e. for  $k = 2$ , both predictions are compared to the ground truth so that each one can be a correct or false prediction). Only images with at least  $k$  annotated objects were used for the evaluation. Figure 2 shows the results when simulating a perfect classifier based on the annotated scene labels and Fig. 3 shows the results when predicting  $P(s|I)$  using the CNN. For both experiments the sampling parameter  $D$  has been set to 10.

Interestingly, the simulation of a perfect scene classifier does not yield superior results compared to those achieved by predicting the scene label using the CNN. The reason for this is two fold: scenes are often




Figure 2: Accuracy for the top  $k$  object predictions on the SUN2012 Scene and Object Dataset. The scene label is predicted by simulating a perfect classifier. Hence, the object prediction is only based on the number of relations between the annotated scene and the set of objects.

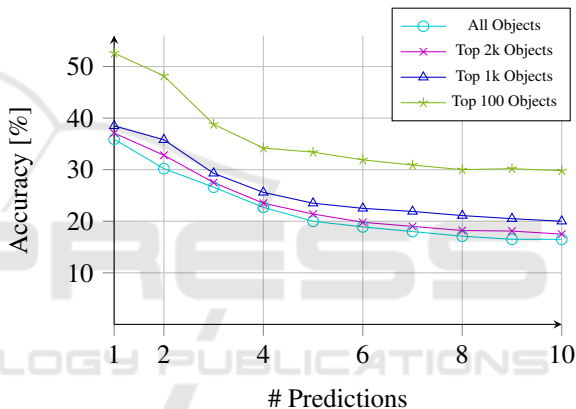


Figure 3: Accuracy for the top  $k$  object predictions on the SUN2012 Scene and Object Dataset. The scene labels are predicted using a CNN and the relations between the predicted scenes and the set of objects is used for predicting the object presence.

ambiguous (cf. (Xiao et al., 2010)) and the prediction using the CNN computes a probabilistic assignment to a set of scenes. For example, a scene depicting a *cathedral*, *church* or *chapel* may not only be visually similar, but they are also similar on a semantic level. As the distribution of objects in scene context is obtained from a very general external text source, it does not accurately match the ground truth distribution that can be observed in the dataset. Hence, a mixture of many scenes is more robust.

The results in Fig. 3 also show that even without any knowledge about the visual appearance of an object the highest ranking object predictions have a precision of up to 52.6% when considering a set of 100 objects and 35.9% when considering as many as 3,390 different object categories. However, as mentioned before, some of the objects describing regional

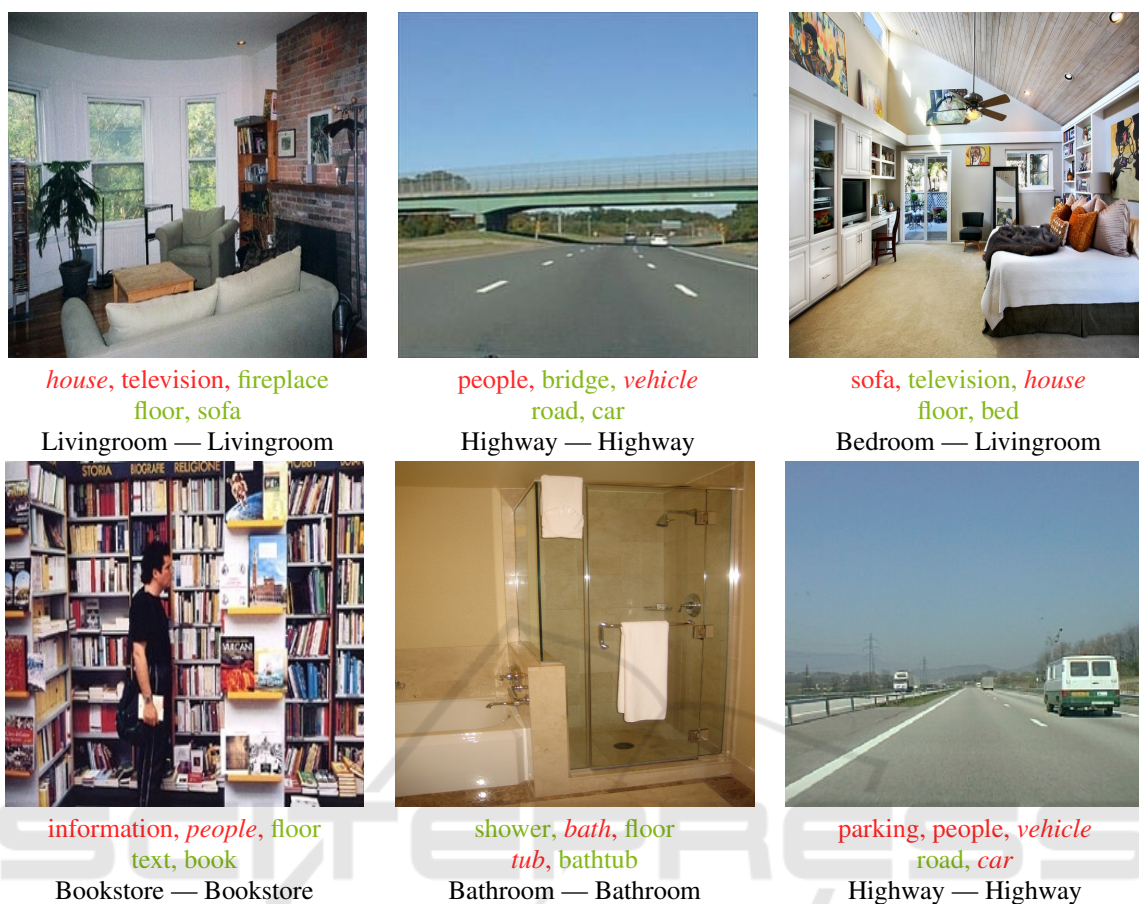


Figure 4: Further exemplary results showing the five highest scoring object predictions: (green) correct (red) wrong (red & italic) wrong according to annotations, but can be seen in the image. In the bottom row: (left) Annotation (right) highest scoring CNN prediction.

objects tend to be very general and can safely be assumed to occur in most scenes.

Exemplary results showing the five highest scoring object predictions for a given image are shown in Fig. 4. It can be seen that although a large set of objects is annotated in the SUN dataset, the annotations are noisy and not at all complete. Some of the predictions that cannot be found in the ground truth annotations might be deemed as correct. Predictions that are not shown in the image, are often at least plausible guesses what else could be found in the scene. For example, people might be related to a highway or street scene and a sofa might be related to an indoor scene. The third example in the top row is a typical case where the highest scoring prediction of the CNN is incorrect. However, the prediction of livingroom instead of bedroom is not only visually but also semantically related. The wall on the left with the TV could also easily be placed in a livingroom. Furthermore, due to the probabilistic assignment to a set of scenes, the object class bed is still the fifth best

scoring prediction although no relation between livingroom and bed has been found in the external text sources. The example of a bathroom in the bottom row shows a typical example of ambiguity in natural language as well as in the provided annotations.

In order to provide a more detailed analysis, different sets of object categories are evaluated based on the VOC mean average precision (mAP) criterion (Everingham et al., 2015). The results for the 20 to 100 most common objects in the dataset is given in Tab. 2. The mAP of predicting an object by chance is indicating how frequently these objects occur in the dataset. For comparison, the ground truth distribution of all scenes and objects, which has been observed in the dataset, has been evaluated using the same model. This can be seen as an indication of an upper bound for the performance that could be obtained by solely using the proposed model of scene and object relations. It can clearly be observed that the relations obtained from the text sources do not model the distribution in the dataset perfectly. The proposed method

Table 2: Mean average precision for different sets of objects on the SUN2012 dataset in Pascal format. The presence predictions are based on the number of relations between the scenes and the objects. (1<sup>st</sup> col.) Simulation of a perfect classifier using the ground truth scene labels (2<sup>nd</sup> col.) Scene labels are predicted using a CNN. For comparison the results of a prediction by chance (3<sup>rd</sup> col.) and the results of the proposed method when using the true number of scene-object relations derived from the dataset (4<sup>th</sup> col.) are also shown.

Objects	Perfect Classifier mean AP [%]	Scene-level CNN mean AP [%]	Chance mean AP [%]	GT Distribution mean AP [%]
Top 20	34.7	38.5	21.6	54.1
Top 40	29.5	33.1	13.9	47.1
Top 60	24.2	27.3	10.3	43.2
Top 80	21.5	24.2	8.5	39.2
Top 100	19.4	22.1	7.2	35.9

Table 3: Mean average precision for the object presence in the SUN2009 Context dataset. (\*) The CNN predicts one of the 397 scenes from the SUN397 dataset without any knowledge about the objects.

Method	Annotations	# Objects	m AP[%]
Part based Models (Felzenszwalb et al., 2010)	Cropped objects	107	17.9
PbM + Tree Context (Choi et al., 2012)	Cropped objects	107	26.1
PbM + Context SVM (Choi et al., 2012)	Cropped objects	107	23.8
Scene-level CNN + Objects in Context	- (*)	107	19.1
Scene-level CNN + Objects in Context	- (*)	104	19.8

shows promising results given the fact that the relations are obtained from arbitrary websites. The results also show that there is potential for improvement if the relation can be estimated in a more accurate manner.

### SUN2009 Context

In order to emphasize the difficulty of detecting objects with a huge variability, as depicted in the SUN dataset, the approach has also been evaluated on the 107 object categories of the SUN2009 Context dataset. On this dataset, different object detectors based on deformable part based models were trained in a fully supervised manner and evaluated in (Choi et al., 2012). The presence detection of the proposed zero-shot method after predicting a scene label for each of the scenes using the CNN is compared to the supervised object detectors of (Choi et al., 2012).

The results are shown in Table 3. As the original evaluation protocol contained three classes that were filtered by the stemming (*bottles*, *stones* and *rocks*) and can therefore not be recognized by the proposed method, the results for all 107 classes as well as the results for the remaining 104 classes after the stemming are displayed. It is surprising that a model that is solely based on scene-level predictions can achieve comparable results to deformable part based models that are trained completely supervised. Only with additional contextual information the part based models are able to outperform the proposed approach. This clearly shows the requirement for contextual informa-

tion, especially since the since visual information in cluttered scenes may be limited.

Note that although there is no evaluation of R-CNNs on this task, they have surpassed deformable part based models as the state-of-the-art in object detection (Girshick et al., 2016). It can be assumed that they outperform the part based models on this task as well. Nevertheless, deformable part based models are powerful object detectors and it is interesting that an unsupervised approach can achieve similar results. This shows that the relations between scenes and objects provide important cues for object prediction.

## 4 CONCLUSION

In this work a novel approach for predicting object presence in an image has been presented. The method works in a zero-shot manner and only relies on scene level annotations from which a probability for an object’s presence is derived. The probability is based on the relations between scenes and objects that were obtained from additional text corpora. As a result the proposed method is completely unsupervised with respect to objects and allows for predicting objects without any visual information about them. In the experiments it has been shown that it is possible to predict the occurrences for as many as 3390 objects. On tasks that are very difficult for visual classifiers, such as cluttered scenes with not very well structured objects, the approach yields similar performance to



visual object detectors that were trained in a fully supervised manner.

## ACKNOWLEDGEMENTS

This work has been supported by the German Research Foundation (DFG) within project Fi799/9-1. The authors would like to thank Kristian Kersting for his helpful comments and discussions.

## REFERENCES

- Banko, M., Cafarella, M. J., Soderland, S., Broadhead, M., and Etzioni, O. (2007). Open information extraction for the web. In *IJCAI*, volume 7, pages 2670–2676.
- Choi, M. J., Torralba, A., and Willsky, A. S. (2012). A tree-based context model for object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(2):240–252.
- Divvala, S. K., Hoiem, D., Hays, J. H., Efros, A. A., and Hebert, M. (2009). An empirical study of context in object detection. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1271–1278.
- Donahue, J., Anne Hendricks, L., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., and Darrell, T. (2015). Long-term recurrent convolutional networks for visual recognition and description. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Etzioni, O., Fader, A., Christensen, J., Soderland, S., and Mausam, M. (2011). Open information extraction: The second generation. In *IJCAI*, volume 11, pages 3–10.
- Everingham, M., Eslami, S. M. A., Van Gool, L., Williams, C. K. I., Winn, J., and Zisserman, A. (2015). The Pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136.
- Felzenszwalb, P. F., Girshick, R. B., McAllester, D., and Ramanan, D. (2010). Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2016). Region-based convolutional networks for accurate object detection and segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):142–158.
- Grzeszick, R., Sudholt, S., and Fink, G. A. (2016). Optimistic and pessimistic neural networks for scene and object recognition. *CoRR*, abs/1609.07982.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L., Shamma, D. A., Bernstein, M., and Fei-Fei, L. (2016). Visual genome: Connecting language and vision using crowdsourced dense image annotations.
- Lampert, C. H., Nickisch, H., and Harmeling, S. (2014). Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):453–465.
- Lin, T., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In *Proc. European Conference on Computer Vision (ECCV)*, pages 740–755. Springer.
- Liu, C., Yuen, J., and Torralba, A. (2009). Nonparametric scene parsing: Label transfer via dense scene alignment. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1972–1979. IEEE.
- Manning, C. D. and Schütze, H. (1999). *Foundations of statistical natural language processing*, volume 999. MIT Press.
- Miller, G. A. and Others (1995). WordNet: a lexical database for English. *Communications of the ACM*, 38(11):39–41.
- Modolo, D., Vezhnevets, A., and Ferrari, V. (2015). Context forest for object class detection. In *Proc. British Machine Vision Conference (BMVC)*.
- Oliva, A. and Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. *Progress in brain research*, 155:23.
- Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. MIT press Cambridge, MA.
- Patterson, G., Xu, C., Su, H., and Hays, J. (2014). The sun attribute database: Beyond categories for deeper scene understanding. *International Journal of Computer Vision*, 108(1-2):59–81.
- Ren, S., He, K., Girshick, R. B., and Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99.
- Rohrbach, M., Stark, M., and Schiele, B. (2011). Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1641–1648. IEEE.
- Rohrbach, M., Wei, Q., Titov, I., Thater, S., Pinkal, M., and Schiele, B. (2013). Translating video content to natural language descriptions. In *IEEE International Conference on Computer Vision*.
- Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556.
- Tighe, J. and Lazebnik, S. (2010). Superparsing: scalable nonparametric image parsing with superpixels. In *European conference on computer vision*, pages 352–365. Springer.
- Vezhnevets, A. and Ferrari, V. (2015). Object localization in imagenet by looking out of the window. *arXiv preprint arXiv:1501.01181*.
- Wu, Q., Shen, C., Hengel, A. v. d., Wang, P., and Dick, A. (2016). Image captioning and visual question an-

- swering based on attributes and their related external knowledge. *arXiv preprint arXiv:1603.02814*.
- Xiao, J., Ehinger, K. A., Hays, J., Torralba, A., and Oliva, A. (2014). SUN Database: Exploring a Large Collection of Scene Categories. *International Journal of Computer Vision (IJCV)*, pages 1–20.
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. (2010). Sun database: Large-scale scene recognition from abbey to zoo. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3485–3492. IEEE.
- Zhu, Y., Zhang, C., Ré, C., and Fei-Fei, L. (2015). Building a large-scale multimodal knowledge base system for answering visual queries. *arXiv preprint arXiv:1507.05670*.

