

Weakly-Supervised Image Annotation and Segmentation with Objects and Attributes

Zhiyuan Shi, Yongxin Yang, Timothy M. Hospedales, Tao Xiang

Abstract—We propose to model complex visual scenes using a non-parametric Bayesian model learned from weakly labelled images abundant on media sharing sites such as Flickr. Given weak image-level annotations of objects and attributes without locations or associations between them, our model aims to learn the appearance of object and attribute classes as well as their association on each object instance. Once learned, given an image, our model can be deployed to tackle a number of vision problems in a joint and coherent manner, including recognising objects in the scene (automatic object annotation), describing objects using their attributes (attribute prediction and association), and localising and delineating the objects (object detection and semantic segmentation). This is achieved by developing a novel Weakly Supervised Markov Random Field Stacked Indian Buffet Process (WS-MRF-SIBP) that models objects and attributes as latent factors and explicitly captures their correlations within and across superpixels. Extensive experiments on benchmark datasets demonstrate that our weakly supervised model significantly outperforms weakly supervised alternatives and is often comparable with existing strongly supervised models on a variety of tasks including semantic segmentation, automatic image annotation and retrieval based on object-attribute associations.

Index Terms—Weakly supervised learning, object-attribute association, semantic segmentation, non-parametric Bayesian model, Indian Buffet Process

1 INTRODUCTION

One of the many incredible features of the human visual system is that it is able to generate rich description of scene content after a glance at an image. Such a description typically contains nouns and adjectives, corresponding to objects and their associated attributes, respectively. For example, an image can be described as containing “a person in red and a shiny car”. In addition, humans can effortlessly delineate each object in the scene. One of the key objectives of computer vision research in the past five decades is to imitate this ability, resulting in intensive studies of a number of fundamental computer vision problems including recognising objects in the scene (object annotation¹) [3], [4], describing the objects using their attributes (attribute prediction and association) [5], [6], [7], and localising and delineating the objects (object detection and semantic segmentation) [8], [9], [10].

Although all of these problems are closely related, existing studies typically focus on one problem only, or two of them but solve them independently. Additionally, most studies employ fully supervised models learned from strongly labelled data. Specifically, in a conventional supervised approach (Fig. 1) images are strongly labelled with object bounding boxes or segmentation masks, and associated attributes, from which object detectors and attribute classifiers are learned. Given a new image, the learned object detectors are first applied to find object locations, where the attribute classifiers are then applied to produce the object descriptions. However, this conventional fully su-

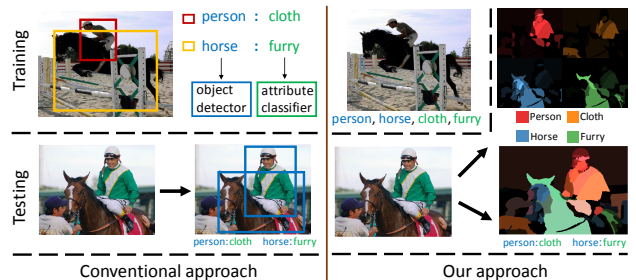


Fig. 1: Comparing our weakly supervised approach to object-attribute association learning to the conventional strongly supervised approach.

pervised and independent learning approach has two critical limitations: (1) Considering there are over 30,000 object classes distinguishable to humans [11], a large number of attributes to describe them, and a much larger number of object-attribute combinations, fully supervised learning is not scalable due to the lack of fully labelled training data. (2) Tackling closely related tasks jointly in a single model can be beneficial. In particular, recent studies have shown that modelling attributes boosts object prediction accuracy and vice-versa [12], and localising objects helps both automatic object annotation and attribute prediction [13].

We aim to overcome these limitations by solving all of these problems jointly with an object+attribute model learned from weakly labelled data, i.e., images with object and attribute labels but neither their associations nor their locations in the form of bounding boxes or segmentation masks (see Fig. 1). As such weakly labelled images are abundant on media sharing websites such as Flickr. There-

1. Note that while annotation sometimes refers to human created ground truth, here it refers to automatic tagging of an image with detected object categories [1], [2].

fore lack of training data would never be a problem.

However, learning strong semantics, i.e. explicit object-attribute association for each object instance from weakly labelled images is extremely challenging due to the label ambiguity: a real-world image with the tags “dog, white, coat, furry” could contain a furry dog and a white coat or a furry coat and a white dog. Furthermore, the tags/labels typically only describe the foreground/objects. There could be a white building in the background which is ignored by the annotator, and a computer vision model must infer that this is not what the tag ‘white’ refers to. A desirable model thus needs to jointly learn multiple objects, attributes and background clutter in a single framework in order to explain away ambiguities in each by knowledge of the other. In addition, a potentially unlimited number of attributes can co-exist on a single object (e.g. there are hundreds of different ways to describe the appearance of a person) which are almost certainly not labelled exhaustively. They also need to be modelled so that they do not act as distractors that have a detrimental effect on the understanding of objects and attributes of interest. For instance, annotators may label bananas in training images but not bother to label yellow. Even if yellow has never been used as an attribute in the training set, the model should be able to infer yellow as a latent attribute [14] and associate it with the bananas, so that other colours would not be assigned wrongly.

To this end, we develop a novel unified framework capable of jointly learning objects, attributes and their associations. The framework is illustrated schematically in Fig. 1, where weak annotations in the form of a mixture of objects and attributes are transformed into object and attribute associations with object segmentation. Under the framework, given a training image with image level labels of objects and attributes, the image is first over-segmented into superpixels; the joint object and attribute annotation and segmentation problem thus becomes a superpixel multi-label classification problem whereby each superpixel can have one object label but an arbitrary number of attribute labels. Treating each label as a factor, we develop a novel factor analysis solution by generalising the non-parametric Indian Buffet Process (IBP) [15].

The IBP is chosen because it is designed for explaining multiple factors that simultaneously co-exist to account for the appearance of a particular image or superpixel, e.g., such factors can be an object and its particular texture and colour attributes. Importantly, as an infinite factor model, it can automatically discover and model latent factors not defined by the provided training data labels, corresponding to latent object/attributes as well as structured background ‘stuff’ (e.g. sky, road). However, the conventional IBP is limited in that it is unsupervised and, as a flat model, applies to either superpixels or images, but not both; it thus cannot be directly applied to our problem. Furthermore, the standard IBP is unable to exploit cues critical for segmentation and object-attribute association by modelling the correlation of factors within and across superpixels in each image. Such within-superpixel correlation captures the co-occurrence relations such as cars are typically metal and bananas are

typically yellow, whilst the across-superpixel correlation dictates that neighbouring superpixels are likely to have similar labels. To overcome these limitations, we formulate a novel variant of IBP, termed Weakly Supervised Markov Random Field Stacked Indian Buffet Process (WS-MRF-SIBP). It differs from the conventional IBP in the following ways: (1) By introducing hierarchy into IBP, WS-MRF-SIBP is able to explain images as groups of superpixels, each of which has an inferred multi-label description vector corresponding to an object and its associated attributes. (2) It learns from weak image-level supervision, which is disambiguated into multi-label superpixel explanations. (3) Two types of Markov Random Field (MRF) over the hidden factors of an image are introduced to the model correlations: across-superpixel MRF to exploit spatial smoothness and within-superpixel MRF to exploit co-occurrence statistics of different attributes with objects.

2 RELATED WORK

Our work is related to a wide range of computer vision problems including image classification, object recognition, attribute learning, and semantic segmentation. It is thus beyond the scope of this paper to present a comprehensive review. Since our approach differs from most existing ones in that it attempts to address all of these problems jointly using a single model learned from weakly labelled data, we shall focus on reviewing studies that solve multiple problems simultaneously and/or use a weakly supervised learning approach.

Learning object-attribute associations Attributes have been used to describe objects [5], [16], people [17], clothing [18], [19], scenes [20], faces [21], [22], and video events [14]. However, most previous studies learn and infer object and attribute models separately, e.g., by independently training binary classifiers, and require strong annotations/labels indicating object/attribute locations and/or associations if the image is not dominated by a single object. A few recent studies have learned object-attribute association explicitly [23], [24], [6], [20], [25], [26], [27], [28]. Different from our approach, [23], [25], [26], [27] only train and test on unambiguous data, i.e. images containing a single dominant object, assuming object-attribute association is known at training; moreover, they allocate exactly one attribute per object. Kulkarni et al. [6] model the more challenging PASCAL VOC type of data with multiple objects and attributes co-existing. However, their model is pre-trained on object and attribute detectors learned using strongly annotated images with object bounding boxes provided. The work in [20] also does object segmentation and object-attribute prediction. But its model is learned from strongly labelled images in that object-attribute association are given during training; and importantly prediction is restricted to object-attribute pairs seen during training. In summary no existing studies learn flexible object-attribute association from weakly labelled data as we do in this work.

Some existing studies aim to perform attribute-based query [29], [30], [31], [21]. In particular, recent studies

have considered how to calibrate [31] and fuse [29] multiple attribute scores in a single query. We go beyond these by supporting object+multi-attribute conjunction queries. Moreover, existing methods either require bounding boxes or assume simple data with single dominant objects, and do not reason jointly about multiple attribute-object associations. This means that they would be intrinsically challenged in reasoning about (multi)-attribute-object queries on challenging data with multiple objects and multiple attributes in each image (e.g., querying furry brown horse, in a dataset with black horses and furry dogs in the same image). In other words, they cannot be directly extended to solve query by object-attribute association.

Weakly supervised semantic segmentation Most existing semantic segmentation models are fully supervised requiring pixel-level labels [8], [9], [10], [12]. A few weakly supervised semantic segmentation methods have been presented recently, exploring a variety of models such as conditional random fields (CRF) [32], [33], label propagation [34] and clustering [35]. More recently, convolutional neural networks have been shown to work very well for this challenging task, either in a fully supervised fashion [36], [37] or a weakly supervised fashion [38], [39]. However, these methods require a large-scale annotated dataset (e.g. ImageNet) to train or pre-train a deep CNN model for feature representation. Another closely related problem is two- or multi-class co-segmentation [40], where the task is to segment shared objects from a set of images. Although co-segmentation does not require image labels *per-se*, it indeed assumes common objects across multiple training images. Like previous semantic segmentation methods, we focus on how to learn a model to segment unseen and unlabelled test images, rather than solely segmenting the training images as in co-segmentation. Importantly, all of these previous methods only focus on object labels (nouns). Our method provides a mechanism to jointly learn objects, attributes and their associations (adjective-noun pairs). We show that attribute labels provide valuable complementary information via inter-label correlation, especially under this more ambitious weakly supervised setting.

This work is not the first to exploit the benefit of joint modelling object and attributes for segmentation. Recently, Zheng et al [12] formulated joint visual attribute and object segmentation as a multi-label problem using a fully connected CRF. Similarly, a model was proposed in [13] to learn and extract attributes from segmented objects, which notably improves object classification accuracy. However, both of these methods are fully supervised, requiring pixel-level ground truth for training. In contrast, our proposed approach can cope with weakly labelled data to alleviate the burden of strong annotation.

Weakly supervised learning: our model vs. discriminative models Discriminative methods underpin many high performance recognition and annotation studies [5], [41], [17], [20], [42], [6]. Similarly existing weakly supervised learning (WSL) methods are also dominated by discriminative models. Apart from the mentioned conditional random field (CRF) [32], [33], label propagation

[34] and clustering [35] models, some discriminative multi-instance learning (MIL) models were proposed [43], [44]. Our model is a probabilistic generative model. Compared to a discriminative model, the strengths of a generative model for WSL is its abilities to infer latent factors corresponding to background clutter and un-annotated objects/attributes, and to model them jointly in a single model so as to explain away the ambiguity existing in the weak image-level labels. Very recently deep learning based image captioning has started to attract attention [45], [46]. Generating a natural sentence describing the an image is a harder task than listing nouns and adjectives - other words including verbs (action) and prepositions (where) need to be inferred and language syntax needs to be followed in the generated text description. However, these neural network models are essentially still discriminative models and have the same drawbacks as other discriminative models for WSL.

Weakly supervised learning: our model vs. other probabilistic generative models The flexibility of generative probabilistic models and their suitability particularly for WSL have seen them successfully applied to a variety of WSL tasks [2], [14], [47], [48]. These studies often generalise probabilistic topic models (PTM) [49]. However PTMs are limited for explaining objects and attributes in that latent topics are competitive - the fundamental assumption is that an object is a horse *or* brown *or* furry. They intrinsically do not account for the reality that it is *all* at once. In contrast, our model generalises Indian Buffet Process (IBP) [50], [15]. The IBP is a latent feature model that can independently activate each latent factor, explaining imagery as a weighted sum of active factor appearances.

Our Weakly Supervised Markov Random Field Stacked Indian Buffet Process (WS-MRF-SIBP) differs significantly from the standard flat and unsupervised IBP in that it is hierarchical to model grouped data (images composed of superpixels) and weakly supervised. This allows us to exploit image-level weak supervision, but disambiguate it to determine the best explanation in terms of which superpixels correspond to un-annotated background, which superpixels correspond to which annotated objects, and which objects have which attributes. In addition, a Markov random field (MRF) is integrated into the IBP to model correlations of factors both within and across superpixels. A few previous studies [51], [52], [53] generalise classic PTMs [49] by integrating a MRF to enforce spatial coherence across topic labels of neighbouring regions. Unlike these methods, we generalise the IBP by defining the MRF over hidden factors. Furthermore, beyond spatial coherence we also define a factorial MRF to capture attribute-attribute and attribute-object co-occurrences within superpixels.

Our contributions This paper makes the following key contributions: (i) We for the first time jointly learn all object, attribute and background appearances, object-attribute association, and their locations from realistic weakly labelled images including multiple objects with cluttered background; (ii) we formulate a novel weakly supervised Bayesian model by generalising the IBP to make it weakly

supervised, hierarchical, and integrate two types of hidden factor MRFs to learn and exploit spatial coherence and factor co-occurrence; (iii) Once learned from weakly labelled data, our model can be deployed for various tasks including semantic segmentation, image description and image query, many of which rely on predicting strong object-attribute association. Extensive experiments on benchmark datasets demonstrate that on all tasks our model significantly outperforms a number of weakly supervised baselines and in many cases is comparable to strongly supervised alternatives. A preliminary version of our work was described in [54].

3 METHODOLOGY

3.1 Image representation

Given a set of images labelled with image-level object and attribute labels, but without explicitly specifying which attribute is associated with which object, we aim to learn a model that, given a new image, segments each object in the image and assign both object and attribute labels to it. As in most previous semantic segmentation works [10], [8], [9], [32], [33], we first decompose each image into superpixels which are over-segmented image patches that typically contain object parts. The problem of joint object and attribute annotation and segmentation thus boils down to multi-label classification of each superpixel, from which various tasks such as automatic image-level annotation, object-attribute association, and object segmentation can be performed.

Each image i in a training set is decomposed into N_i superpixels using a recent hierarchical segmentation algorithm [55]². Each segmented superpixel is represented using two normalised histogram features: SIFT and Color. (1) SIFT: we extract regular grid (every 5 pixels) colorSIFT [56] at four scales. A 256 component GMM model is constructed on the collection of ColourSIFTs from all images. We compute Fisher Vector + PCA for all regular points in each superpixel following [57]. The resulting reduced descriptor is 512-D for every segmented region. (2) Colour: We convert the image to quantised LAB space $8 \times 8 \times 8$. A 512-D color histogram is then computed for each superpixel. The final normalised 1024-D feature vector concatenates SIFT and Colour features together.

3.2 Model formulation

We propose a non-parametric Bayesian model that learns to describe images composed of superpixels from weak image-level object and attribute annotation. In our model, each superpixel is associated with an infinite latent factor vector indicating if it corresponds to (an unlimited variety of) unannotated background clutter, or an object of interest, and what set of attributes are possessed by the object. Given a set of images with weak labels and segmented into superpixels, we need to learn: (i) which are the corresponding superpixels shared by all images with a particular label, (ii)

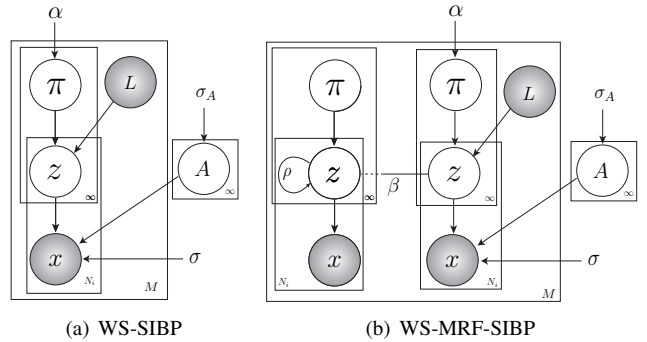


Fig. 2: The probabilistic graphical models representing our WS-SIBP and WS-MRF-SIBP. Shaded nodes are observed.

which superpixels correspond to unannotated background, and (iii) what is the appearance of each object, attribute and background type. Moreover, since multiple labels (attribute and object) can apply to a single superpixel, we need to disambiguate which aspects of the appearance of each superpixel are due to each of the (unknown) associated object and attribute labels. To address the weakly-supervised learning tasks we build on the IBP [50] and introduce in Sec. 3.2.1 a weakly-supervised stacked Indian Buffet process (WS-SIBP) to model data represented as bags (images) of instances (superpixels) with bag-level labels (image-level annotations). This is analogous to the notion of documents in topic models [49]. Furthermore, to fully exploit spatial and inter-factor correlation, two types of MRFs are integrated (see Sec. 3.2.2), resulting in the full model termed WS-MRF-SIBP³.

3.2.1 WS-SIBP

We aim to associate each image/superpixel with a latent factor vector whose elements will correspond to objects, attributes and/or unannotated attribute/background present in that image/superpixel. Let images i be represented as bags of superpixels $\mathbf{X}^{(i)} = \{\mathbf{X}_j^{(i)}\}$, where the notation \mathbf{X}_j means the vector of row j in matrix \mathbf{X} , i.e. the 1024-D feature vector representing the j -th superpixel, and $j \in \{1 \dots N_i\}$. Assuming there are K_o object categories and K_a attributes in the provided image-level annotations, they are represented by the first $K_{oa} = K_o + K_a$ latent factors. In addition, an unbounded number of further factors are available to explain away background clutter in the data, as well as discover unannotated latent attributes. At training time, we assume a binary label vector $L^{(i)}$ for objects and attributes is provided for each image i . So $L_k^{(i)} = 1$ if attribute/object k is present, and zero otherwise. Also $L_k^{(i)} = 1$ for all $k > K_{oa}$. That is, without any labels, we assume all background/latent attribute types can be present. With these assumptions, the generative process (see Fig. 2(a)) for the image i is as follows:

For each latent factor $k \in 1 \dots \infty$:

2. We set the segmentation threshold to 0.1 to obtain a single over-segmentation from the hierarchical segmentations for each image.

3. The codes for both models will be made available at <http://zhiyshi.com/publication/>

- 1) Draw an appearance distribution mean $\mathbf{A}_k \sim \mathcal{N}(0, \sigma_A^2 \mathbf{I})$.

For each image $i \in 1 \dots M$:

- 1) Draw a sequence of i.i.d. random variables $v_1^{(i)}, v_2^{(i)} \dots \sim \text{Beta}(\alpha, 1)$,
- 2) Construct an image prior $\pi_k^{(i)} = \prod_{t=1}^k v_t^{(i)}$,
- 3) Input weak annotation $L_k^{(i)} \in \{0, 1\}$,
- 4) For each superpixel $j \in 1 \dots N_i$:
- Sample state of each latent factor k : $z_{jk}^{(i)} \sim \text{Bern}(\pi_k^{(i)} L_k^{(i)})$,
 - Sample superpixel appearance: $\mathbf{X}_j^{(i)} \sim \mathcal{N}(\mathbf{Z}_j^{(i)} \mathbf{A}, \sigma^2 \mathbf{I})$.

where \mathcal{N} , Bern and Beta respectively correspond to Normal, Bernoulli and Beta distributions with the specified parameters. The Beta-Bernoulli and Normal-Normal conjugacy are chosen because they allow more efficient inference. α is the prior expected sparsity of annotations and σ^2 is the prior variance in appearance for each factor.

This generative process encodes the assumptions that the available factors for each superpixel are determined by the image level labels if given (generative model for \mathbf{Z}); and that multiple factors come together to explain each superpixel (generative model for \mathbf{X} given \mathbf{Z}).

Joint probability: Denote hidden variables by $\mathbf{H} = \{\boldsymbol{\pi}^{(1)}, \dots, \boldsymbol{\pi}^{(M)}, \mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(M)}, \mathbf{A}\}$, M images in a training set by $\mathbf{X} = \{\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(M)}\}$, and parameters by $\boldsymbol{\Theta} = \{\alpha, \sigma_A, \sigma, \mathbf{L}\}$. Then the joint probability of the variables and data given the parameters is:

$$p(\mathbf{H}, \mathbf{X} | \boldsymbol{\Theta}) = \prod_{i=1}^M \left(\prod_{k=1}^{\infty} \left(p(\pi_k^{(i)} | \alpha) \prod_{j=1}^{N_i} p(z_{jk}^{(i)} | \pi_k^{(i)}, L_k^{(i)}) \right) \cdot \prod_{j=1}^{N_i} p(\mathbf{X}_j^{(i)} | \mathbf{Z}_j^{(i)}, \mathbf{A}, \sigma) \right) \cdot \prod_{k=1}^{\infty} p(\mathbf{A}_k | \sigma_A^2). \quad (1)$$

Learning our model (detailed in Sec. 3.3) aims to compute the posterior $p(\mathbf{H} | \mathbf{X}, \boldsymbol{\Theta})$ for: disambiguating and localising all the annotated ($L^{(i)}$) objects and attributes among the superpixels (inferring $\mathbf{Z}_j^{(i)}$), inferring the attribute and background prior for each image (inferring $\boldsymbol{\pi}^{(i)}$), and learning the appearance of each factor (inferring \mathbf{A}_k).

3.2.2 WS-MRF-SIBP

Now we generalise the WS-SIBP to WS-MRF-SIBP by introducing two types of factor correlation.

Spatial MRF across superpixels: Each superpixel's latent factors are so far drawn from the image prior $\pi_k^{(i)}$ – independently of their neighbours (Eq. (1)). Thus spatial structure is ignored in WS-SIBP, even though adjacent superpixels are strongly correlated in real images [51]. Inspired by the successful use of random fields for capturing the spatial

coherence of image region labels [52], [51], [53], we introduce a MRF with connections between adjacent nodes (superpixels). Specifically, the following MRF potential [51], [58] is introduced to the generative process for \mathbf{Z} to correlate the superpixel factors drawn in image i spatially:

$$\Phi(\mathbf{Z}_{\cdot k}^{(i)}) = \exp \sum_{j, m \in N_i} \beta \mathbf{I}(z_{jk}^{(i)} = z_{mk}^{(i)}), \quad (2)$$

where $j, m \in N_i$ enumerates node pairs that are neighbours in image i . The indicator function \mathbf{I} returns one when its argument is true, i.e., when neighbouring superpixels have the same assignment for factor k . β is the coupling strength parameter of the MRF, which controls how likely they have the same label *a priori*. The initial WS-SIBP formulation can be obtained by setting $\beta = 0$. The spatial MRF is encoded for all given K_{oa} and newly discovered factors.

Factorial MRF within superpixel: Although individual factors are now correlated spatially, we do not yet model any inter-factor co-occurrence statistics within a single superpixel (as in most other MRF applications [51], [52]). However, exploiting this information (e.g., *person* superpixels more likely to share attribute *clothing*, than *metallic*) is important, especially in the ambiguous WSL setting. To represent these inter-factor correlations, we introduce a factorial MRF via the following potential on \mathbf{Z} :

$$\Psi(\mathbf{Z}_j^{(i)}) = \exp \sum_{k, l}^{\infty} \psi(z_{jk}^{(i)}, z_{jl}^{(i)}) \quad (3)$$

$$\psi(z_{jk}^{(i)}, z_{jl}^{(i)}) = \begin{cases} 0 & \text{if } k = l \\ \rho \mathcal{M}_{kl} & \text{otherwise,} \end{cases} \quad (4)$$

where ρ controls the importance of the factorial MRF, and \mathcal{M}_{kl} is an element of the factor correlation matrix \mathbf{M} that encodes the correlation between factor k and factor l . In the traditional strongly-supervised scenario, \mathbf{M} can be trivially learnt from the fully labelled annotations. In the WSL scenario, \mathbf{M} cannot be determined directly. We will discuss how to learn \mathbf{M} in Sec. 3.3.

WS-MRF-IBP Prior: Overall, combining the two MRFs, the latent factor prior

$$p(\mathbf{Z}^{(i)} | \boldsymbol{\pi}^{(i)}, L^{(i)}) = \prod_{k=1}^{\infty} \prod_{j=1}^{N_i} p(z_{jk}^{(i)} | \pi_k^{(i)}, L_k^{(i)})$$

used by Eq. (1), is now replaced by:

$$p(\mathbf{Z}^{(i)} | \boldsymbol{\pi}^{(i)}, L^{(i)}, \beta, \rho) \propto \exp \left(\sum_{k=1}^{\infty} \sum_{j=1}^{N_i} \log p(z_{jk}^{(i)} | \pi_k^{(i)}, L_k^{(i)}) + \sum_{k=1}^{\infty} \log \Phi(\mathbf{Z}_{\cdot k}^{(i)}) + \sum_{j=1}^{N_i} \log \Psi(\mathbf{Z}_j^{(i)}) \right).$$

and the list of model parameters $\boldsymbol{\Theta}$ is extended to $\boldsymbol{\Theta} = \{\alpha, \sigma_A, \sigma, \mathbf{L}, \rho, \beta, \mathbf{M}\}$.

3.3 Model learning

Exact inference for $p(\mathbf{H} | \mathbf{X}, \boldsymbol{\Theta})$ in our model is intractable, so an approximate inference algorithm in the spirit of [50]

```

while not converge do
  for k = 1 to K_max do
    
$$\phi_k = \left( \frac{1}{\sigma^2} \sum_{i=1}^M \sum_{j=1}^{N_i} \nu_{jk}^{(i)} (\mathbf{X}_j^{(i)} - \sum_{l:l \neq k} \nu_{jl}^{(i)} \phi_l) \right) \cdot \left( \frac{1}{\sigma_A^2} + \frac{1}{\sigma^2} \sum_{i=1}^M \sum_{j=1}^{N_i} \nu_{jk}^{(i)} \right)^{-1} \quad (5)$$

    
$$\Phi_k = \left( \frac{1}{\sigma_A^2} + \frac{1}{\sigma^2} \sum_{i=1}^M \sum_{j=1}^{N_i} \nu_{jk}^{(i)} \right)^{-1} \mathbf{I} \quad \triangleright \text{Update appearance model including mean and covariance.} \quad (6)$$

  end
  for i = 1 to M do
    for k = 1 to K_max do
      
$$\tau_{k1}^{(i)} = \alpha + \sum_{m=k}^{K_{max}} \sum_{j=1}^{N_i} \nu_{jm}^{(i)} + \sum_{m=k+1}^{K_{max}} (N_i - \sum_{j=1}^{N_i} \nu_{jm}^{(i)}) \left( \sum_{s=k+1}^m q_{ms}^{(i)} \right) \quad (7)$$

      
$$\tau_{k2}^{(i)} = 1 + \sum_{m=k}^{K_{max}} (N_i - \sum_{j=1}^{N_i} \nu_{jm}^{(i)}) q_{mk}^{(i)} \quad \triangleright \text{Update image prior for every factor k.} \quad (8)$$

      for j = 1 to N_i do
        
$$\eta = \sum_{t=1}^k (\varphi(\tau_{t1}^{(i)}) - \varphi(\tau_{t2}^{(i)})) - \mathbb{E}_{\mathbf{v}}[\log(1 - \prod_{t=1}^k v_t^{(i)})] - \frac{1}{2\sigma^2} (\text{tr}(\Phi_k) + \phi_k \phi_k^T - 2\phi_k (\mathbf{X}_j^{(i)} - \sum_{l:l \neq k} \nu_{jl}^{(i)} \phi_l)^T) \quad \triangleright \text{Top-down prior and bottom-up data cues.} \quad (9)$$

        
$$\eta' = \eta + \sum_{m \in N(j)} \beta \eta_{mk}^{(i)} + \sum_{n:n \neq k} \rho \mathcal{M}_{kn} \eta_{jn} \quad \triangleright \text{Influence of the two MRFs.} \quad (10)$$

        
$$\nu_{jk}^{(i)} = \frac{L_k^{(i)}}{1 + \exp[-\eta']} \quad \triangleright \text{Final posterior for each latent factor } z_{jk}^{(i)} \text{ state.} \quad (11)$$

      end
    end
  end
  for i = 1 to M do
    for j = 1 to N_i do
      
$$\mathbf{M} = \mathbf{M} + (\eta_j^i)^T \eta_j^i \quad \triangleright \text{Update the intra-superpixel correlation given inferred factors.} \quad (12)$$

    end
  end
end

```

Algorithm 1: Variational inference for learning WS-MRF-SIBP

is developed. The mean field variational approximation to the desired posterior $p(\mathbf{H}|\mathbf{X}, \Theta)$ is:

$$q(\mathbf{H}) = \prod_{i=1}^M (q_{\tau}(\mathbf{v}^{(i)}) q_{\nu}(\mathbf{Z}^{(i)}) q_{\phi}(\mathbf{A})) \quad (13)$$

where $q_{\tau}(v_k^{(i)}) = \text{Beta}(v_k^{(i)}; \tau_{k1}^{(i)}, \tau_{k2}^{(i)})$, $q_{\nu}(z_{jk}^{(i)}) = \text{Bernoulli}(z_{jk}^{(i)}; \nu_{jk}^{(i)})$, $q_{\phi}(\mathbf{A}_{k\cdot}) = \mathcal{N}(\mathbf{A}_{k\cdot}; \phi_k, \Phi_k)$ and the infinite stick-breaking process for latent factors is truncated at K_{max} , so $\pi_k = 0$ for $k > K_{max}$. A variational message passing (VMP) strategy [50] can be used to minimise the KL divergence of Eq. (13) to the true posterior. Updates are obtained by deriving integrals of the form $\ln q(\mathbf{h}) = E_{\mathbf{H} \setminus \mathbf{h}}[\ln p(\mathbf{H}, \mathbf{X})] + C$ for each group of hidden variables \mathbf{h} . These result in the series of iterative updates given in Algorithm 1, where $\varphi(\cdot)$ is the digamma function; and $q_{ms}^{(i)}$ and $\mathbb{E}_{\mathbf{v}}[\log(1 - \prod_{t=1}^k v_t^{(i)})]$ are given in [50].

Like [52], [51], the MRF influence is via Eqs. (10) and (11). However, while the works of [52], [51] only consider spatial coherence, we further model the inter-factor correlation, which we will see is very important for our weakly supervised tasks, especially in image annotation.

Factor correlation learning: The correlation matrix \mathbf{M} is non-trivial to estimate accurately in the WSL case, in contrast to the fully-supervised case where it is easy to obtain as the correlation of superpixel annotations. In the WSL case, it can only be estimated a priori from image-level tags. However, this is a very noisy estimate. For example, an image with tags *furry, horse, metal, car* will erroneously suggest *horse-car, furry-metal, horse-metal* as correlations.

To address this, we initialise \mathbf{M} coarsely with image-level labels as $\mathbf{M} = \sum_{i=1}^M (L_k^{(i)})^T (L_k^{(i)})$, and refine it with an EM process. During learning, we re-estimate \mathbf{M} at each iteration using the disambiguated superpixel-level factors inferred by the model, as in Eq. (12). Thus as the correlation estimate improves, the estimated factors become more accurate, and vice-versa. The effectiveness of this iterative learning procedure is demonstrated in the supplementary material.

Efficiency: In practice, the truncation approximation means that our WS-MRF-SIBP runs with a finite number of factors K_{max} which can be freely set so long as it is bigger than the number of factors needed by both annotations and

background clutter (K_{bg}), i.e., $K_{max} > K_o + K_a + K_{bg}$ ⁴. Despite the combinatorial nature of the object-attribute association and localisation problem, our model is of complexity $\mathcal{O}(MNDK_{max} + K_{max}^2)$ for M images with N superpixels, D feature dimension and K_{max} truncated factors.

3.4 Inference for test data

At testing time, the appearance of each factor k , now modelled by sufficient statistics $\mathcal{N}(\mathbf{A}_k; \phi_k, \Phi_k)$, is assumed to be known (learned from the training data), while annotations for each test image $L_k^{(i)}$ will need to be inferred. Thus Algorithm 1 still applies, but without the appearance update terms (Eqs. (5) and (6)) and with $L_k^{(i)} = 1 \forall k$, to reflect the fact that all the learned object, attribute, and background types could be present.

3.5 Applications of the model

Given the learned model applied to test data, we can perform the following tasks.

Free image annotation: This is to describe an image using a list of nouns and adjectives corresponding to objects and their associated attributes. To infer the objects present in image i , the first K_o latent factors of the inferred $\pi^{(i)}$ are thresholded or ranked to obtain a list of objects. This is followed by locating them via searching for the superpixels j^* maximising $\mathbf{Z}_{jk}^{(i)}$, then thresholding or ranking the K_a attribute latent factors in $\mathbf{Z}_{j^*k}^{(i)}$ to describe them. This corresponds to a “describe this image” task.

Annotation given object names: This is a more constrained variant of the free annotation task above. Given a named (but not located) object k , its associated attributes can be estimated by first finding the location as $j^* = \arg \max_j \mathbf{Z}_{jk}^{(i)}$, then the associated attributes by $\mathbf{Z}_{j^*k}^{(i)}$ for $K_o < k \leq K_o + K_a$. This corresponds to a “describe this (named) object in an image” task.

Object+attribute query: Images can be queried for a specified object-attribute conjunction $\langle k_o, k_a \rangle$ by searching for $i^*, j^* = \arg \max_j \mathbf{Z}_{jk_o}^{(i)} \cdot \mathbf{Z}_{j^*k_a}^{(i)}$. This corresponds to a “find images with a particular kind of object” task.

Semantic segmentation: In this application, we aim to label each superpixel j with one of K_o learned object factors. The label of superpixel j can be obtained by searching $k^* = \arg \max_k \mathbf{Z}_{jk}^{(i)}$, where $k \in K_o$. Although the annotation search space is solely objects, inference of the additional $k > K_o$ factors (including unannotated background or attribute annotation) can help detect objects $k \in K_o$ via disambiguation. Note that unlike most weakly supervised semantic segmentation methods [32], [35], our model can operate *without* access to the whole test set. But it can also operate in a transductive setting as those existing methods. Under this setting, the appearance distribution $\mathcal{N}(\mathbf{A}_k; \phi_k, \Phi_k)$ will be further updated by Eqs. (5) and

(6) based on the test images. The image-level labels of test data is assigned by the inferred factors of our model or alternatively by an image classifier (see Sec. 4.2.3).

4 EXPERIMENTS

Extensive experiments are carried out to demonstrate the effectiveness of our model on three real-world applications: automatic image annotation (see Sec. 4.1.2), object-attribute query (see Sec. 4.1.3) and semantic segmentation (see Sec. 4.2).

4.1 Image annotation and query

4.1.1 Datasets and settings

For the automatic image annotation and query tasks, various object and attribute datasets are available such as aPascal [5], ImageNet [41], SUN [59] and Awa [7]. We choose aPascal because it has multiple objects per image; and ImageNet because attributes are shared widely across categories.

aPascal: This dataset [5] is an attribute labelled version of PASCAL VOC 2008. There are 4340 images of 20 object categories. Each object is annotated with a list of 64 attributes that describe them by shape (e.g., isBoxy), parts (e.g., hasHead) and material (e.g., isFurry). In the original aPascal, attributes are strongly labelled for 12695 object bounding boxes, i.e. the object-attribute association are given. To test our weakly supervised approach, we merge the object-level category annotations and attribute annotations into a single annotation vector of length 84 for the entire image. This image-level annotation is much weaker than the original bounding-box-level annotation, as shown in Fig. 3. In all experiments, we use the same train/test splits provided by [5].

ImageNet Attribute: This dataset [41] contains 9600 images from 384 ImageNet synsets/categories. To study WSL, we ignore the provided bounding box annotation. Attributes for each bounding box are labelled as 1 (presence), -1 (absence) or 0 (ambiguous). We use the same 20 of 25 attributes as [41] and consider 1 and 0 as positive examples. Many of the 384 categories are subordinate categories, e.g. dog breeds. However, distinguishing fine-grained subordinate categories is beyond the scope of this study. That is, we are interested in finding a ‘black-dog’ or ‘white-car’, rather than ‘black-labrador’ or ‘white-ford-focus’. We thus convert the 384 ImageNet categories to 172 entry-level categories using [60] (see Fig. 4). We evenly split each class to create the training and testing sets.

We compare our WS-MRF-SIBP to two strongly supervised models and four weakly supervised alternatives:

Strongly supervised models: A strongly supervised model uses bounding-box-level annotation. Two variants are considered for the two datasets respectively. **DPM+s-SVM:** for aPascal, both object detector and attribute classifier are trained from fully supervised data (i.e. Bounding-Box-level annotation in Fig. 3). Specifically, we use the 20 pre-trained DPM detectors from [3] and 64 attribute classifiers from [5]. **GT+s-SVM:** for ImageNet attributes, there is

4. In this work, we set $K_{max} = K_o + K_a + 20$.

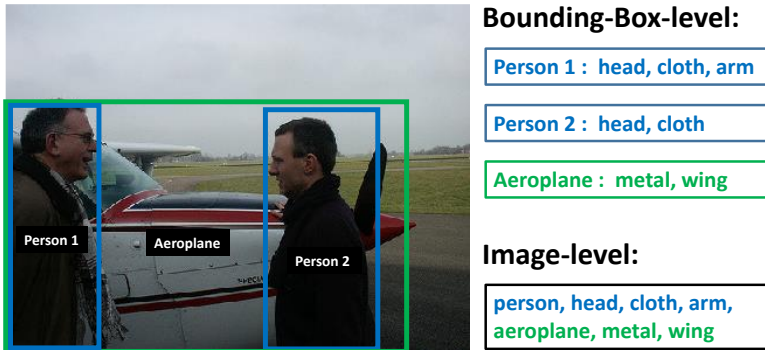


Fig. 3: Strong bounding-box-level annotation and weak image-level annotations for aPascal are used for learning strongly supervised models and weakly supervised models respectively.

not enough data to learn 172 strong DPM detectors as in aPascal. So we use the ground truth bounding box instead assuming we have perfect object detectors, giving a significant advantage to this strongly supervised model. We train attribute classifiers using our features (Sec. 3.1) and liblinear SVM [61]. These strongly supervised models are similar in spirit to the models used in [6], [20], [23] and can be considered to provide an upper bound for the performance of the weakly supervised models.

Weakly supervised models: w-SVM [5], [41]: In this weakly-supervised baseline, both object detectors and attribute classifiers are trained on the weak image-level labels as for our model (see Fig. 3). For aPascal, we train object and attribute classifiers using the feature extraction and model training codes (which is also based on [61]) provided by the authors of [5]. For ImageNet, our features are used, without segmentation. **MIML**: This is the multi-instance multi-label (MIML) learning method in [62]. Our model can also be seen as a MIML method with each image a bag and each superpixel an instance. The MIML model provides a mechanism to use the same superpixel based representation for images as our model, thus providing the object/attribute localisation capability as our model does. **w-LDA**: Weakly-supervised Latent Dirichlet Allocation (LDA) approaches [63], [48] have been used for object localisation. We implement a generalisation of LDA [49], [48] that accepts continuous feature vectors (instead of bag-of-words). Like MIML this method can also accept superpixel based representation, but w-LDA is more related to our WS-SIBP than MIML since it is also a generative model. **WSDC** [35]: Weakly supervised dual clustering is a recently proposed method for semantic segmentation that estimates pixel-level annotation given only image-level labels. This semantic segmentation method can be repurposed to our image annotation setting by considering the same input (superpixel representation + image-level label) followed by the same method as in our framework to first infer superpixel level labels and then aggregate them to compute image-level annotations (see Sec. 3.5).



Fig. 4: 43 subordinate classes of dog are converted into a single entry-level class ‘dog’.

4.1.2 Automatic Image annotation

An image description (annotation) can be automatically generated by predicting objects and their associated attributes. To comprehensively cover all aspects of performance of our method and competitors, we perform three annotation tasks with different constraints on test images: (1) *free annotation*, where no constraint is given to a test image, (2) *annotation given object names*, where named but not located objects are given, and (3) *annotation given locations*, where object locations are given in the form of bounding boxes, and the attributes are predicted.

	aPascal [5]			ImageNet [41]		
	AP@2	AP@5	AP@8	AP@2	AP@5	AP@8
w-SVM [5]	24.8	21.2	20.3	46.3	41.1	37.5
MIML [62]	28.7	22.4	21.0	46.6	43.2	38.3
w-LDA [48]	30.7	24.0	21.5	48.4	43.1	38.4
WSDC [35]	29.8	25.1	21.3	48.0	42.7	36.5
Ours	40.1	29.7	25.0	60.7	54.2	50.0
D/G+s-SVM	40.6	30.3	23.8	65.9	60.7	53.2

TABLE 1: Free annotation performance (AP@t) evaluated on t attributes per object.

Free annotation: For WS-MRF-SIBP, w-LDA and MIML the procedure in Sec. 3.5 is used to detect objects and then describe them using the top t attributes. For the strongly supervised model on aPascal (DPM+s-SVM), we use DPM object detectors to find the most confident objects and their bounding boxes in each test image. Then we use the 64 attribute classifiers to predict top t attributes in each bounding box. In contrast, w-SVM trains attributes and objects independently, and cannot associate objects and attributes. We thus use it to predict only one attribute vector per image regardless of which object label it predicts.

Since there are a variable number of objects per image in aPascal, quantitatively evaluating free annotation is not straightforward. Therefore, we evaluate only the most confident object and its associated top t attributes in each image, although more could be described. For ImageNet, there is only one object per image. We follow [1], [64] in evaluating annotation performance by average precision (AP@t), given varying numbers (t) of predicted attributes per object. If




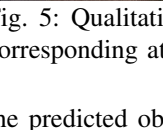
	w-SVM		MIML		w-LDA		WSDC		Ours		DPM+s-SVM			
	bicycle	motorbike	person	bicycle	motorbike	bicycle	motorbike	person	person	motorbike	person	motorbike		
	metal row wind shiny text wool	metal row wind shiny text wool	skin cloth shiny leather foot/shoe	headlight window wheel arm screen	taillight label engine shiny glass	wheel rein beak metal sail	metal wheel saddle furn.seat taillight	head plastic screen leg leather	cloth head nose face leaf	wheel exhaust shiny wool text	cloth skin hair head leg	person motorbike	engine metal label shiny taillight	
	furry cloth snout leg head	sheep furry cloth snout leg head	furry horn ear occluded leg	mouth furry torso taillight shiny	person	nose furry arm jet engine foot	skin glass face label hand	furry cloth wool head 2D Boxy	skin cloth arm hand rein	head leg furry hand Skin	person	cat	person	dog cloth torso skin clear head
	car	person	person	train	car	cow	horse	bus	horse	car	car	horse	horse	
	window glass pedal 3D boxy metal	window glass pedal 3D boxy metal	hand screen cloth arm flower	metal vert cyl wing door leg	wheel metal propeller label door	furn.leg head furry cloth wool	leg head metal ear hair	glass exhaust rein head 3D Boxy	furry head ear arm leg	metal window wheel 3D boxy door	metal 3D boxy plastic engine side mirror	furry torso leaf saddle feather	person	

Fig. 5: Qualitative results on free annotation. False positives are shown in red. If the object prediction is wrong, the corresponding attribute box is shaded.

the predicted object is wrong, all associated attributes are considered wrong.

Table 1 reports the free annotation performance of the compared models. We have the following observations: (1) Our WS-MRF-SIBP, despite learned from weak image-level annotation, yields comparable performance to the strongly supervised model (DPM/GT+s-SVM). The gap is particularly small for the more challenging aPascal dataset, whilst for ImageNet, the gap is bigger as the strongly supervised GT+s-SVM has an unfair advantage by using the ground truth bounding boxes during testing. (2) WS-MRF-SIBP consistently outperforms the four weakly supervised alternatives. The margin is especially large for $t = 2$ attributes per object, which is closest to the true number of attributes per object. For bigger t , all models must generate some irrelevant attributes thus narrowing the gaps. (3) As expected, the w-SVM model obtains the weakest results, confirming that the ability to locate objects is important for modelling object-attribute association. (4) Compared to the two generative models (ours and w-LDA), MIML has worse performance because a generative model is more capable of utilising weak labels [48]. The other discriminative model WSDC fares better than MIML due to its ability to exploit superpixel appearance similarity to disambiguate image-level labels, but it is still inferior to our model. (5) Between the two generative models, the advantage of our framework over w-LDA is clear; due to the ability of IBP to explain each superpixel with multiple non-competing factors⁵.

Fig. 5 shows some qualitative results on aPascal via the two most confident objects and their associated attributes. This is a challenging dataset – even the strongly supervised DPM+s-SVM makes mistakes for both attribute and object prediction. Compared to the other weakly supervised models, WS-MRF-SIBP has more accurate predictions – it jointly and non-competitively models objects and their attributes so object detection benefits from attribute detection and vice versa. The competitors are also more likely to mismatch attributes with objects, e.g. MIML detects a

shiny person rather than the correct shiny motorbike.

To gain insight into what has been learned by our model and why it is better than the weakly supervised alternatives, Fig. 6 visualises the attribute and object factors learned by the competing models which use superpixels as input. It is evident that without explicit background modelling, MIML suffers by trying to explain the background superpixel using the weak labels. In contrast, both w-LDA and WS-SIBP have good segmentation of foreground objects, showing that both the learned foreground and background topics are meaningful. However, for w-LDA, since object and attributes topics compete for the same superpixel, each superpixel is dominated by either an object or attribute topic. In contrast, the object factors and attribute factors co-exist happily in WS-SIBP as they should do, e.g. most person superpixels have the clothing attribute as well.

Annotation given object names (GN): In this experiment, we assume that object labels are given and we aim to describe each object by attributes, corresponding to tasks such as: “Describe the car in this image”. For the strongly supervised model on aPascal, we use the object’s DPM detector to find the most confident bounding box. Then we predict attributes for that box. Here, annotation accuracy is the same as attribute accuracy, so the performance of different models is evaluated following [65] by mean average precision (mAP) under the precision-recall curve. Note that for aPascal, w-SVM reports the same list of attributes for all co-existing objects, without being able to localise and distinguish them. Its result is thus not meaningful and is excluded. The same set of conclusions can be drawn from Table 2 as in the free annotation task: our WS-MRF-SIBP is on par with supervised models and outperforms weakly supervised ones.

		w-SVM	MIML	w-LDA	WSDC	Ours	SS
GN	aPascal	–	32.1	35.5	36.3	39.3	41.8
	ImageNet	32.4	33.5	39.6	44.2	52.8	56.8
GL	aPascal	33.2	35.1	35.8	38.4	43.6	42.1
	ImageNet	37.7	39.1	46.8	48.2	53.9	56.8

TABLE 2: Results on annotation given object names (GN) or locations (GL). SS stands for Strongly Supervised.

Given object location (GL): If we further know the

5. Training two independent w-LDA models for objects and attributes respectively is not a solution: the problem would re-occur for multiple competing attributes.

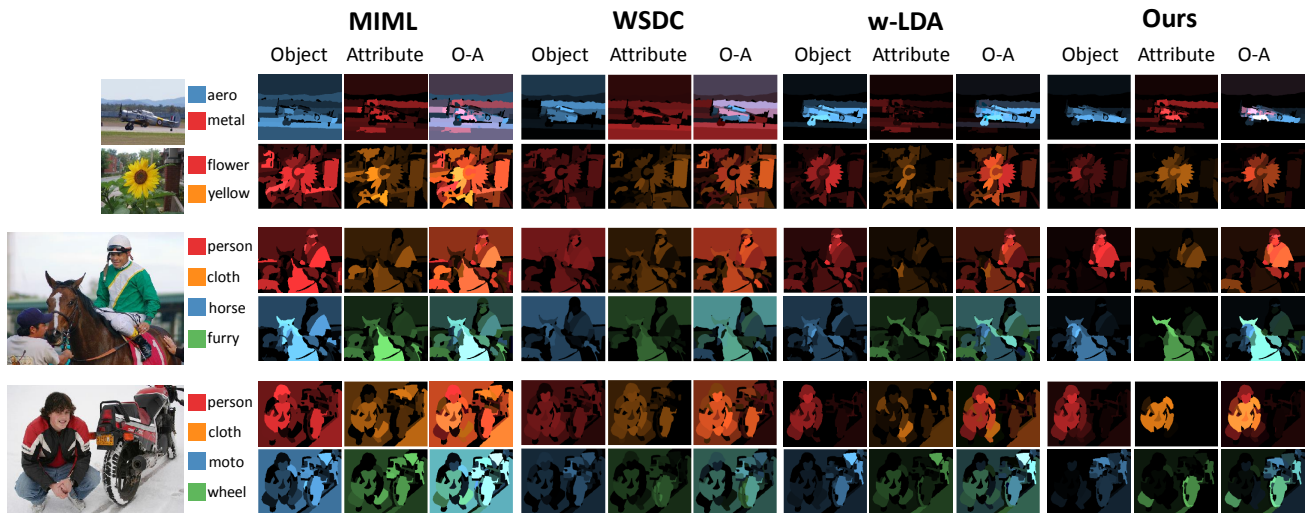


Fig. 6: Illustrating the inferred superpixel annotation. Object and attributes are coloured, and multi-label annotation blends colours. Each of the bottom two groups has two rows corresponding to the two most confident objects detected.

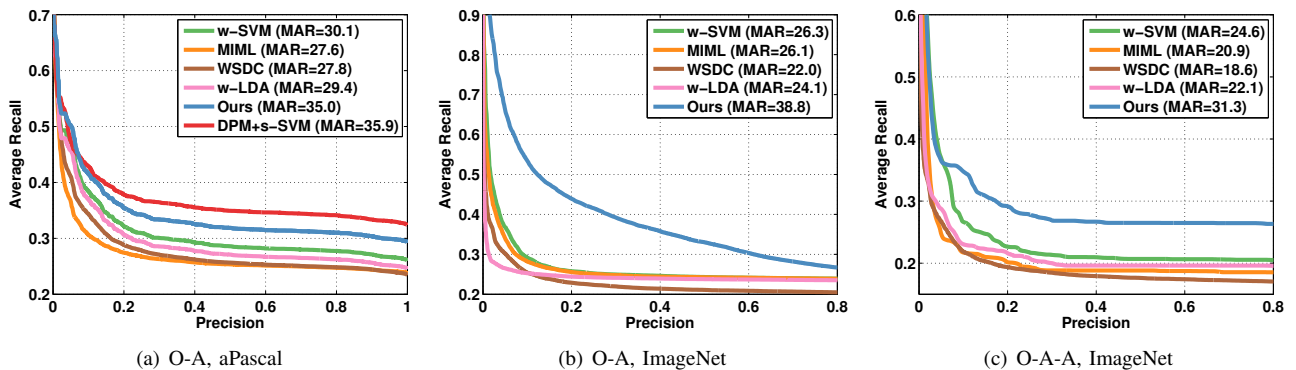


Fig. 7: Object-attribute query results as precision-average recall curve.

bounding box of an object in a test image, we can simply predict attributes inside each bounding box. This becomes the conventional attribute prediction task [5], [41] for describing an object. Table 2 shows the results, where similar observations can be made as in the other two tasks above. Note that in this case the strongly supervised model is the method used in [5]. The mAP obtained using our weakly supervised model is even higher than the strongly supervised model (though our area-under-ROC-curve value of 81.5 is slightly lower than the 83.4 reported in [5]).

4.1.3 Object-attribute query

In this task object-attribute association is used for image retrieval. Following work on multi-attribute queries [29], we use mean average recall over all precisions (MAR) as the evaluation metric. Note that unlike [29] which requires each queried *combination* to have enough (100) training examples to train conjunction classifiers, our method can query novel never-previously-seen combinations. Three experiments are conducted. We generate 300 random object-attribute combinations for aPascal and ImageNet respectively and 300 object-attribute-attribute queries for ImageNet. For the strongly supervised model, we normalise and multiply object detector with attribute classifier scores.

No object detector is trained for ImageNet so no result is reported there. For w-SVM, we use [31] to calibrate the SVM scores for objects and attributes as in [29]. For the three WS models, the procedure in Sec. 3.5 is used to compute the retrieval ranking.

Quantitative results are shown in Fig. 7 and some qualitative examples in Fig. 8. Our approach has very similar MAR values to the strongly supervised DPM+s-SVM, while outperforming all the other models. w-SVM calibration [31] helps it outperform MIML and w-LDA. However, the lack of object-attribute association and background modelling still causes problems for w-SVM. This is illustrated in the ‘dog-black-white’ example shown in Fig. 8 where a white background caused an image with a black dog retrieved at Rank 2 by w-SVM.

4.2 Semantic Segmentation

4.2.1 Datasets and settings

We evaluate semantic segmentation performance on the aPascal Segmentation (Sec. 4.2.2) and LabelMe Outdoor datasets (Sec. 4.2.3) under the weakly supervised setting. **aPascal Segmentation:** This dataset [12] is a subset of PASCAL VOC 2008 [5] where both pixel-level segmentation and attributes annotation are available. It contains 639

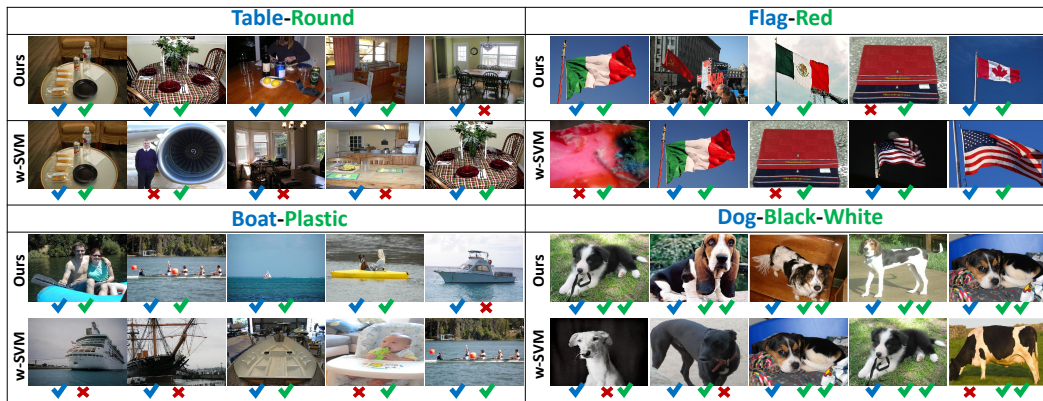


Fig. 8: Object-attribute query: qualitative comparison. Given a query, such as “Table+Round”, we list the top-5 results predicted by our method and w-SVM.

images from 20 classes. The 64 attribute annotation for each image is the same as the aPascal dataset used in the annotation experiments. We use the training (326 images) and testing (313 images) split provided by [12].

LabelMe Outdoor Segmentation: Also known as SIFT Flow [66], this widely used dataset contains 2688 images densely labelled with 33 object classes at pixel-level using the LabelMe online annotation tool. Every pixel in each image is assigned a label meaning that background ‘stuff’ such as sky, sea, street are also labelled as objects. Most images contain outdoor scenes. We use the standard training (2488 images) and testing (200 images) split provided in [66]. Note that no attribute labels are available for this dataset.

Evaluation metrics: The evaluation metrics used for semantic segmentation are often dataset dependent. Past works [10], [67] on the LabelMe dataset typically report results in both total per-pixel accuracy (defined as the number of correctly labelled pixels over the total number of pixels), and per-class accuracy (defined as the number of correctly labelled pixels for a class over the number of ground truth pixels of this class and then averaged over all object classes). Both metrics are necessary because for any model, some parameters can typically be tuned so that one metric is favoured at the expense of the other. The VOC images have very different characteristics compared with LabelMe. In particular, the images often contain large portions of unannotated background (stuff) and the 20 objects of interest are relatively small. The intersection-over-union (IOU) score is thus typically used for semantic segmentation performance evaluation on the Pascal VOC dataset [12], [35], [68] and adopted here on aPascal.

4.2.2 Results on aPascal

To our knowledge, no previous work models objects and attributes jointly for semantic segmentation in the weakly supervised setting. We therefore apply the weakly supervised (object only) segmentation method WSDC [35] as an alternative⁶ (see Sec. 4.1.1). With our WS-MRF-SIBP, the association of objects and available attributes can be

6. We do not have access to the codes of other weakly supervised segmentation methods and implementing them is non-trivial.

leveraged to improve performance. We explore three different sets of attribute annotations: (1) 8 attributes: material attributes used by [12] for object-attribute segmentation. (2) 64 attributes: the original attributes provided by [5]. (3) 74 attributes: we add 10 more color attributes based on aPascal sentence descriptions [6], [69].

		Method	Avg. IOU (%)
Fully		Zheng et al [12]	37.1
		Krähenbühl et al [70]	36.9
Weakly		WSDC [35]	18.2
		Ours_0attribute	23.6
		Ours_8attributes	27.3
		Ours_64attributes	28.9
		Ours_74attributes	29.4

TABLE 3: Quantitative semantic segmentation comparison versus state-of-the-art on the aPascal dataset.

Our model is compared with one weakly supervised [35] and two fully-supervised alternatives [12], [70] in Table 3. It can be observed that our method outperforms the alternative weakly-supervised model WSDC [35], even without attribute annotation (Ours_0attribute). Moreover, its performance gradually improves as more attribute annotation becomes available, and eventually its performance with 74 attributes is not far off from the two fully-supervised models.

Some qualitative results are shown in Fig. 9. In each example, we show the color-coded segmentation output of Ours_0attribute, Ours_64attribute, Krähenbühl et al [70] and ground truth segmentation. Coloured regions are identified as the same foreground objects. Background clutter with similar appearance to foreground object can confuse Ours_0attribute as shown in the horse example (the third row of Fig. 9). However, by exploiting the additional (weak) attribute annotation, the segmentation performance is greatly improved (Our_64attribute) through disambiguation, and by capturing object-attribute co-occurrence.

4.2.3 Results on LabelMe

Table 4 compares the performance of our model with a number of state-of-the-art fully supervised [71], [10], [8],



Fig. 9: Qualitative illustration of (attribute-enhanced) semantic segmentation results on aPascal.

		Method	Per-pixel (%)	Per-class (%)
Fully		Tighe et al [71]	77.0	30.1
		Tighe et al [10]	78.6	39.2
		Sigh and Kosecka [8]	79.2	33.8
		Yang et al [9]	79.8	48.7
		Gould et al [67]	78.4	25.7
Weakly		Vezhnevets et al [72]	-	14
		Vezhnevets et al [33]	-	21
		Xu et al [32]	21.9	27.9
		WSDC [35]	19.3	25.0
		Ours	46.2	23.8
		Ours_transductive	52.5	31.2
	Ours_predict	48.1	26.7	

TABLE 4: Quantitative comparison of semantic segmentation performance on the LabelMe dataset.

[9], [67] and weakly supervised [33], [35], [32] models⁷. Three variants of our models are evaluated: Ours and Ours_transductive differ in whether the test set images are used for model update (see Sec. 3.5), whilst for Ours_predict, we follow [32] and use a pre-trained multi-label image classifier (Linear SVM with ImageNet-trained CNN features as input) to predict image-level object labels and use those labels for transductive learning.

The results show that our model outperforms the alternative weakly supervised models [33], [35], [32] particularly in the per-pixel accuracy which reflects more on the performance on the large classes such as sky and sea. Note that a number of recent weakly supervised learning methods including [35] are transductive. But our model, even without accessing the whole test set (Ours), can double the per-pixel accuracy of the alternative models whilst being comparable to them in per-class accuracy. When our model operates in the transductive mode (Ours_transductive), the

⁷ Very recently the weakly supervised segmentation results have been significantly improved by deep convolutional neural network-based models [73], [39]. However, these models use much more data to train therefore having an unfair advantage; e.g., the model in [73] is pre-trained using 1M ImageNet 1K images and fine-tuned with an additional set of 60K background images.

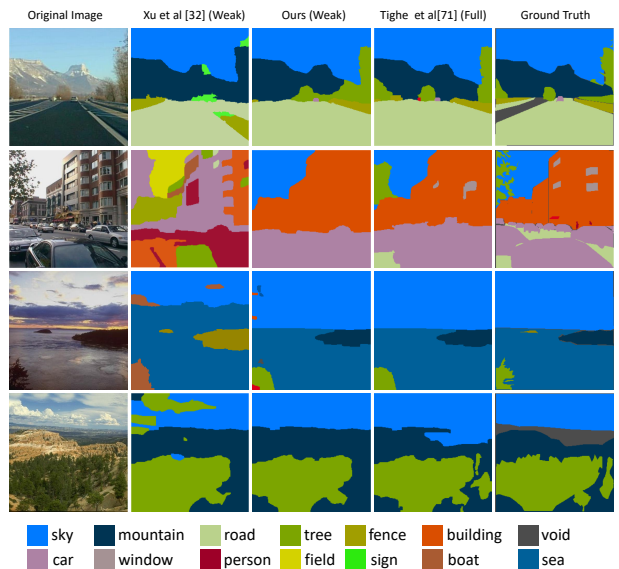


Fig. 10: Qualitative comparison of our semantic segmentation versus alternatives on the LabelMe dataset (best viewed in colour).

margin over the other models including [35], [32] gets even bigger. It is worth mentioning that the result of Xu et al [32] is obtained using the predicted image-level labels on the test set. Our result (Ours_transductive vs. Ours_predictive) suggests that this additional step is not necessary using our model – as demonstrated in the image annotation experiments earlier, our model itself can predict image-level labels and does not require assistance from another model. Table 4 also shows the performance of a number of state-of-the-art strongly supervised learning models [71], [10], [8], [9], [67] which require pixel-level annotation of the training images. As can be seen, there is still a fairly big gap between theirs and the best result achieved by our model (Ours_transductive), although on the per-class metric, it is much closer.

Fig. 10 qualitatively compares two weakly supervised methods ([32] and Ours) and one fully supervised method ([71]). We note that the advantage of our model over that in [32] is particularly pronounced in the cluttered street scene (second row of Fig. 10). For a scene like this, the ability to infer latent factors which correspond to latent attributes for describing object appearance is critical. For example, our model seems to be able to capture the fact that sky can have different types of appearance: clear and blue in the first row, overcast in the second, and cloudy in the bottom. Without accounting for these variations of appearance, the model in [32] struggled and assigned wrong labels to sky in the second and bottom row images.

4.3 Running cost:

Our unoptimised single-core implementation was run on a PC with an Intel 3.47 GHz CPU and 16GB RAM. The computational time of our approach is comparable to existing methods. Table 5 compares the per iteration time spent on training. These methods require similar numbers

of iterations to converge (around 1500). During testing, our method takes about 0.2 seconds per image, comparing to 0.15 using w-LDA. Transductive WSDC does not have a separate train and test stage, and thus there are no directly comparable figures.

Method	w-LDA	WSDC [35]	Ours
Time	65	70	80

TABLE 5: Computation time (seconds per iteration) of different methods on aPascal training set (2113 images).

5 CONCLUSION

We have presented an effective model for weakly-supervised learning of objects, attributes, and their locations and associations. Learning object-attribute association from weak image-level labels is non-trivial but critical for learning from ‘natural’ data, and scaling to many classes and attributes. We achieve this for the first time through a novel weakly-supervised IBP model that simultaneously disambiguates superpixel annotation correspondence, and learns the appearance of each annotation and superpixel-level annotation correlation. Our results show that on a variety of tasks, our model often performs comparably to strongly supervised alternatives that are significantly more costly to supervise, and is consistently better than weakly supervised alternatives.

The presented model can be improved in a number of ways. First, although the two MRFs integrated in our model capture the spatial label coherence and within-superpixel factorial correlation, other correlations can be considered. One is the image-level correlation capturing object-object co-occurrences in each image, for example, car and road typically co-exist in a street scene. Another is the spatial correlation capturing relative location between objects, for example, sky often being at the top of the road. Modelling this correlation provides additional constraints to obtain better image and thus superpixel-level labels. However, including this in the current model is non-trivial and a more complex learning algorithm may be needed [74]. Another prior knowledge the current model does not exploit is the fact that each superpixel should only be explained by a single object label. Although the within-superpixel MRF can implicitly model that, the current model can be extended to suppress the co-occurrence of multiple object labels explicitly. Finally, the study of automated scene understanding has evolved from single object, multiple object, multiple objects+attributes, towards automated image captioning with full sentences [45], [46], and visual question answering (VQA) [75], [76]. These models are discriminative, being hybrids of deep Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). Our framework provides a partial answer to captioning – (it provides a list of objects and their attribute associations) and VQA (it can return the attributes of a queried object, or vice-versa) via an ontology rather than natural language. One future research direction is to integrate to our model with deep learning based language models to tackle the

full image captioning/VQA tasks whilst keeping the advantages of the introduced generative non-parametric Bayesian model for weakly supervised learning.

REFERENCES

- [1] Z. Feng, R. Jin, and A. Jain, “Large-scale image annotation by efficient and robust kernel metric learning,” in *ICCV*, 2013.
- [2] L.-J. Li, R. Socher, and L. Fei-Fei, “Towards total scene understanding: classification, annotation and segmentation in an automatic framework,” in *CVPR*, 2009.
- [3] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, “Object detection with discriminatively trained part based models,” *TPAMI*, 2010.
- [4] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, “The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results,” <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- [5] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth, “Describing objects by their attributes,” in *CVPR*, 2009.
- [6] G. Kulkarni, V. Premraj, S. Dhar, S. Li, Y. Choi, A. Berg, and T. Berg, “Baby talk: Understanding and generating simple image descriptions,” in *CVPR*, 2011.
- [7] C. H. Lampert, H. Nickisch, and S. Harmeling, “Attribute-based classification for zero-shot visual object categorization,” *IEEE TPAMI*, 2013.
- [8] G. Singh and J. Kosecka, “Nonparametric scene parsing with adaptive feature relevance and semantic context,” in *CVPR*, 2013.
- [9] J. Yang, B. Price, S. Cohen, and M.-H. Yang, “Context driven scene parsing with attention to rare classes,” in *CVPR*, 2014.
- [10] J. Tighe, M. Niethammer, and S. Lazebnik, “Scene parsing with object instance inference using regions and per-exemplar detectors,” *International Journal of Computer Vision*, 2015.
- [11] I. Biederman, “Recognition-by-components: A theory of human image understanding,” *Psychological Review*, 1987.
- [12] S. Zheng, M.-M. Cheng, J. Warrell, P. Sturgess, V. Vineet, C. Rother, and P. Torr, “Dense semantic image segmentation with objects and attributes,” in *CVPR*, 2014.
- [13] Z. Li, E. Gavves, T. Mensink, and C. G. Snoek, “Attributes make sense on segmented objects,” in *ECCV*, 2014.
- [14] Y. Fu, T. Hospedales, T. Xiang, and S. Gong, “Learning multi-modal latent attributes,” *PAMI*, 2013.
- [15] T. L. Griffiths and Z. Ghahramani, “The indian buffet process: An introduction and review,” *JMLR*, 2011.
- [16] N. Turakhia and D. Parikh, “Attribute dominance: What pops out?” in *ICCV*, 2013.
- [17] L. Bourdev, S. Maji, and J. Malik, “Describing people: A poselet-based approach to attribute classification,” in *ICCV*, 2011.
- [18] H. Chen, A. Gallagher, and B. Girod, “Describing clothing by semantic attributes,” in *ECCV*, 2012.
- [19] T. L. Berg, A. C. Berg, and J. Shih, “Automatic attribute discovery and characterization from noisy web data,” in *ECCV*, 2010.
- [20] S. Wang, J. Joo, Y. Wang, and S.-C. Zhu, “Weakly supervised learning for attribute localization in outdoor scenes,” in *CVPR*, 2013.
- [21] B. Siddiquie, R. Feris, and L. Davis, “Image ranking and retrieval based on multi-attribute queries,” in *CVPR*, 2011.
- [22] N. Kumar, A. Berg, P. N. Belhumeur, and S. Nayar, “Describable visual attributes for face verification and image search,” *TPAMI*, 2011.
- [23] G. Wang and D. Forsyth, “Joint learning of visual attributes, object classes and visual saliency,” in *ICCV*, 2009.
- [24] A. Kovashka, S. Vijayanarasimhan, and K. Grauman, “Actively selecting annotations among objects and attributes,” in *ICCV*, 2011.
- [25] X. Wang and Q. Ji, “A unified probabilistic approach modeling relationships between attributes and objects,” in *ICCV*, 2013.
- [26] Y. Wang and G. Mori, “A discriminative latent model of object classes and attributes,” in *ECCV*, 2010.
- [27] D. Mahajan, S. Sellamanickam, and V. Nair, “A joint learning framework for attribute models and object descriptions,” in *ICCV*, 2011.
- [28] J. Johnson, R. Krishna, M. Stark, J. Li, M. Bernstein, and L. Fei-Fei, “Image retrieval using scene graphs,” in *CVPR*, 2015.
- [29] M. Rastegari, A. Diba, D. Parikh, and A. Farhadi, “Multi-attribute queries: To merge or not to merge?” in *CVPR*, 2013.
- [30] A. Kovashka and K. Grauman, “Attribute adaptation for personalized image search,” in *ICCV*, 2013.

- [31] W. Scheirer, N. Kumar, P. N. Belhumeur, and T. E. Boult, "Multi-attribute spaces: Calibration for attribute fusion and similarity search," in *CVPR*, 2012.
- [32] J. Xu, A. G. Schwing, and R. Urtasun, "Tell me what you see and I will show you where it is," in *CVPR*, 2014.
- [33] A. Vezhnevets, V. Ferrari, and J. Buhmann, "Weakly supervised structured output learning for semantic segmentation," in *CVPR*, 2012.
- [34] M. Rubinstein, C. Liu, and W. T. Freeman, "Annotation propagation in large image databases via dense image correspondence," in *ECCV*, 2012.
- [35] Y. Liu, J. Liu, Z. Li, J. Tang, and H. Lu, "Weakly-supervised dual clustering for image semantic segmentation," in *CVPR*, 2013.
- [36] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *CVPR*, 2015.
- [37] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Semantic image segmentation with deep convolutional nets and fully connected crfs," in *ICLR*, 2015.
- [38] D. Pathak, E. Shelhamer, J. Long, and T. Darrell, "Fully convolutional multi-class multiple instance learning," in *ICLR Workshop*, 2015.
- [39] J. Xu, A. G. Schwing, and R. Urtasun, "Learning to segment under various forms of weak supervision," in *CVPR*, 2015.
- [40] M. Rubinstein, A. Joulin, J. Kopf, and C. Liu, "Unsupervised joint object discovery and segmentation in internet images," in *CVPR*, 2013.
- [41] O. Russakovsky and L. Fei-Fei, "Attribute learning in large-scale datasets," in *ECCV*, 2010.
- [42] M. Sadeghi and A. Farhadi, "Recognition using visual phrases," in *CVPR*, 2011.
- [43] N. Nguyen, "A new svm approach to multi-instance multi-label learning," in *ICDM*, 2010.
- [44] T. Deslaers, B. Alexe, and V. Ferrari, "Weakly supervised localization and learning with generic knowledge," *IJCV*, vol. 100, 2012.
- [45] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and tell: A neural image caption generator," in *CVPR*, 2015.
- [46] A. Karpathy and L. Fei-Fei, "Deep visual-semantic alignments for generating image descriptions," in *CVPR*, 2015.
- [47] R. Socher and L. Fei-Fei, "Connecting modalities: Semi-supervised segmentation and annotation of images using unaligned text corpora," in *CVPR*, 2010.
- [48] Z. Shi, T. M. Hospedales, and T. Xiang, "Bayesian joint topic modelling for weakly supervised object localisation," in *ICCV*, 2013.
- [49] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *JMLR*, 2003.
- [50] F. Doshi-Velez, K. T. Miller, J. V. Gael, and Y.-W. Teh, "Variational inference for the indian buffet process," University of Cambridge, Tech. Rep., 2009.
- [51] J. Verbeek and B. Triggs, "Region classification with markov field aspect models," in *CVPR*, 2007.
- [52] B. Zhao, L. Fei-Fei, and E. P. Xing, "Image segmentation with topic random field," in *ECCV*, 2010.
- [53] B. Zhou, X. Wang, and X. Tang, "Random field topic model for semantic region analysis in crowded scenes from tracklets," in *CVPR*, 2011.
- [54] Z. Shi, Y. Yang, T. M. Hospedales, and T. Xiang, "Weakly supervised learning of objects, attributes and their associations," in *ECCV*, 2014.
- [55] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, "Contour detection and hierarchical image segmentation," *TPAMI*, 2011.
- [56] K. E. A. van de Sande, T. Gevers, and C. G. M. Snoek, "Evaluating color descriptors for object and scene recognition," *TPAMI*, 2010.
- [57] H. Jégou, F. Perronnin, M. Douze, J. Sánchez, P. Pérez, and C. Schmid, "Aggregating local image descriptors into compact codes," *TPAMI*, 2011.
- [58] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer-Verlag New York, Inc., 2006.
- [59] G. Patterson and J. Hays, "Sun attribute database: Discovering, annotating, and recognizing scene attributes," in *CVPR*, 2012.
- [60] V. Ordonez, J. Deng, Y. Choi, A. C. Berg, and T. L. Berg, "From large scale image categorization to entry-level categories," in *ICCV*, 2013.
- [61] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, "LIBLINEAR: A library for large linear classification," *JMLR*, 2008.
- [62] Z.-H. Zhou, M.-L. Zhang, S.-J. Huang, and Y.-F. Li, "Multi-instance multi-label learning," *Artificial Intelligence*, 2012.
- [63] N. Rasiwasia and N. Vasconcelos, "Latent dirichlet allocation models for image classification," *TPAMI*, 2013.
- [64] L. Wu, R. Jin, and A. K. Jain, "Tag completion for image retrieval," *TPAMI*, 2013.
- [65] N. Zhang, R. Farrell, F. Iandola, and T. Darrell, "Deformable part descriptors for fine-grained recognition and attribute prediction," in *ICCV*, 2013.
- [66] C. Liu, J. Yuen, and A. Torralba, "Nonparametric scene parsing via label transfer," *TPAMI*, 2011.
- [67] S. Gould, J. Zhao, X. He, and Y. Zhang, "Superpixel graph label transfer with learned distance metric," in *ECCV*, 2014.
- [68] J. Dong, Q. Chen, S. Yan, and A. Yuille, "Towards unified object detection and semantic segmentation," in *ECCV*, 2014.
- [69] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier, "Collecting image annotations using amazon's mechanical turk," in *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, 2010.
- [70] P. Krähenbühl and V. Koltun, "Efficient inference in fully connected CRFs with gaussian edge potentials," in *NIPS*, 2011.
- [71] J. Tighe and S. Lazebnik, "Superparsing: Scalable nonparametric image parsing with superpixels," *IJCV*, 2013.
- [72] A. Vezhnevets, V. Ferrari, and J. Buhmann, "Weakly supervised semantic segmentation with a multi-image model," in *ICCV*, 2011.
- [73] P. O. Pinheiro and R. Colloberti, "From image-level to pixel-level labeling with convolutional networks," in *CVPR*, 2015.
- [74] F. Doshi-Velez and Z. Ghahramani, "Correlated non-parametric latent feature models," in *UAI*, 2009.
- [75] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh, "Vqa: Visual question answering," in *ICCV*, 2015.
- [76] R. Hu, H. Xu, M. Rohrbach, J. Feng, K. Saenko, and T. Darrell, "Natural language object retrieval," in *CVPR*, 2016.



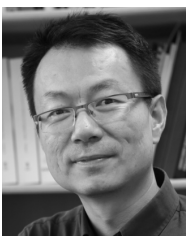
Zhiyuan Shi received the PhD degree in computer science from Queen Mary University of London in 2016. He is currently post-doctoral associate in computer vision and learning lab at Imperial College London. His research interests include action recognition, deep learning, weakly supervised learning, topic model, object localisation and attribute learning.



Yongxin Yang Yongxin Yang is a PhD student at Queen Mary University of London (QMUL). His research interests include machine learning (Transfer Learning, Domain Adaptation, Multi-Task Learning and Deep Learning) and computer vision.



Timothy M. Hospedales received the PhD degree in neuroinformatics from the University of Edinburgh in 2008. He is currently a Reader (associate professor) in Image and Vision Computing at The University of Edinburgh. His research interests include probabilistic modelling and machine learning applied to various problems in computer vision and beyond. He has published more than 60 papers in major international journals and conferences. He is a member of the IEEE.



Tao Xiang received the Ph.D. degree in electrical and computer engineering from the National University of Singapore in 2002. He is currently a reader (associate professor) in the School of Electronic Engineering and Computer Science, Queen Mary University of London. His research interests include computer vision, machine learning, and data mining. He has published over 120 papers in international journals and conferences.