







DOCK: Detecting Objects by Transferring Common-Sense Knowledge

Krishna Kumar Singh^{1,3} , Santosh Divvala^{2,3} , Ali Farhadi^{2,3} ,
and Yong Jae Lee¹ 

¹ University of California, Davis, USA
krsingh@ucdavis.edu

² University of Washington, Seattle, USA

³ Allen Institute for AI, Seattle, USA
<https://dock-project.github.io>

Abstract. We present a scalable approach for Detecting Objects by transferring Common-sense Knowledge (DOCK) from source to target categories. In our setting, the training data for the source categories have bounding box annotations, while those for the target categories only have image-level annotations. Current state-of-the-art approaches focus on image-level visual or semantic similarity to adapt a detector trained on the source categories to the new target categories. In contrast, our key idea is to (i) use similarity not at the image-level, but rather at the region-level, and (ii) leverage richer common-sense (based on attribute, spatial, etc.) to guide the algorithm towards learning the correct detections. We acquire such common-sense cues automatically from readily-available knowledge bases without any extra human effort. On the challenging MS COCO dataset, we find that common-sense knowledge can substantially improve detection performance over existing transfer-learning baselines.

1 Introduction

Object detection has witnessed phenomenal progress in recent years, where *fully-supervised* detectors have produced amazing results. However, getting large volumes of bounding box annotations has become an Achilles heel of this setting. To address this scalability concern, transfer-learning methods that transform knowledge from source categories (with bounding boxes) to *similar* target classes (with only image labels) have evolved as a promising alternative [17, 32, 37, 39].

While recent works [17, 18, 39] have demonstrated the exciting potential of transfer-learning on *object-centric* datasets like ImageNet, it has not yet been thoroughly explored on more complex *scene-centric* datasets like MS COCO. Why is it so?

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-01261-8_30) contains supplementary material, which is available to authorized users.

We hypothesize three key challenges:

(i) Existing transfer learning methods rely only on similarity knowledge between the source and target categories to compute the transformations that need to be transferred. Unfortunately, using similarity alone is often insufficient. For example, can you guess the orange-colored region proposal in the masked image using only the provided similarity cues (Fig. 1)? (ii) Existing methods depend on having a robust image-level object classifier for transferring knowledge, which for object-centric datasets like ImageNet is easy to obtain, but is challenging for scene-centric datasets like MS COCO (where multiple and potentially small objects like ‘toothbrush’ exist in each image). If the image classifier does not perform well, then transforming it into a detector will not perform well either. (iii) Finally, if instances of the target classes frequently co-occur with the source classes, then the target class regions can end up being undesirably learned as ‘background’ while training the detector for the source classes.

In this paper, we overcome the above limitations by proposing a new approach for Detecting Objects by transferring Common-sense Knowledge (DOCK). To overcome the first limitation, our key idea is to leverage multiple sources of *common-sense* knowledge. Specifically, we encode: (1) *similarity*, (2) *spatial*, (3) *attribute*, and (4) *scene*. For example, if ‘spoon’ is one of the target objects and ‘fork’, ‘table’, and ‘kitchen’ are among the source categories, we can learn to better detect the ‘spoon’ by leveraging the fact that it is usually ‘*similar* to fork’, ‘*on* a table’, ‘*is* metallic’, and ‘*seen* in kitchens’. Figure 2 shows another scenario, which builds upon the Fig. 1 example.

In this way, even if a target class does not have a visually/semantically *similar* class among the source classes, the other common-sense can help in obtaining a better detector. All the common-sense knowledge we use is freely-acquired from readily-available external knowledge bases [19, 23, 27, 29, 38, 44]. Further, our approach learns all the required common-sense models using source-class bounding box annotations only and does not require *any* bounding box annotations for the target categories.

To address the latter limitations, our idea is to directly model objects at the region-level rather than at the image-level. To this end, any detection framework that learns using region proposals with image-level labels for the target object categories is applicable. In this paper, we use an object proposal classification and ranking detection framework based upon [5] for its simplicity and competitive performance. It learns an object detector by minimizing an image




Fig. 1. Guess the object? (Color figure online)



Fig. 2. Using the multiple common-sense cues, can you now guess the object corresponding to the orange box? (For answer, see [1]) (Color figure online)

classification loss based on object proposals’ class probabilities. We inject common-sense into this framework by modulating the object proposals’ class probabilities with our proposed common-sense prior probabilities. The proposed priors give higher preference to regions that are more likely (under common-sense) to belong to the object-of-interest. Interestingly, since the common-sense cues are encoded only as *priors*, our algorithm can choose to ignore them when they are not applicable. This is particularly helpful to alleviate the concern when frequently co-occurring target classes are incorrectly learned as ‘background’.

We evaluate our approach on the challenging MS COCO dataset [22]. We find that transferring common-sense knowledge substantially improves object detection performance for the target classes that lack bounding box annotations, compared to other contemporary transfer-learning methods [17, 18, 39]. We also perform ablation analyses to inspect the contribution of our proposed idea of encoding common sense. Finally, we explore the potential of our proposed framework in the context of webly-supervised object detection.

2 Related Work

Transfer learning for scalable object detection. Existing transfer learning approaches can be roughly divided into two groups: one that learns using bounding box annotations for both source and target categories [3, 11, 21, 36, 40, 43], and another that learns using bounding box annotations for source categories but only image-level annotations for target categories [17, 18, 32, 37, 39]. In this paper, we are interested in the latter setting, which is harder but likely to be more scalable. In particular, the recent state-of-the-art deep learning approaches of [17, 18, 39] adapt an image classifier into an object detector by learning a feature transformation between classifiers and detectors on the source classes, and transfer that transformation to related target classes based on visual or semantic similarity for which only classifiers exist. While our approach also leverages pre-trained object detectors for encoding visual and semantic similarity, we explore additional common-sense cues such as spatial and attribute knowledge. Moreover, both [17, 39] use similarity information at the image-level (which works well on only object-centric datasets like ImageNet), while our approach uses similarity at the region-level. All these contributions together help us achieve a significant performance boost on the scene-centric MS COCO dataset.

Use of context. Our use of common-sense is related to previous works on context which leverage additional information beyond an object’s visual appearance. Context has been used for various vision tasks including object detection [4, 8, 9, 15, 31, 33], semantic segmentation [28], and object discovery [10, 20]. As context by definition is something that frequently co-occurs with the object-of-interest, without bounding box annotations the contextual regions can easily be confused with the object-of-interest (e.g., a piece of ‘road’ context with a ‘car’). Our approach tries to address this issue by using external common-sense knowledge to model context for a target object in terms of its spatial relationship with

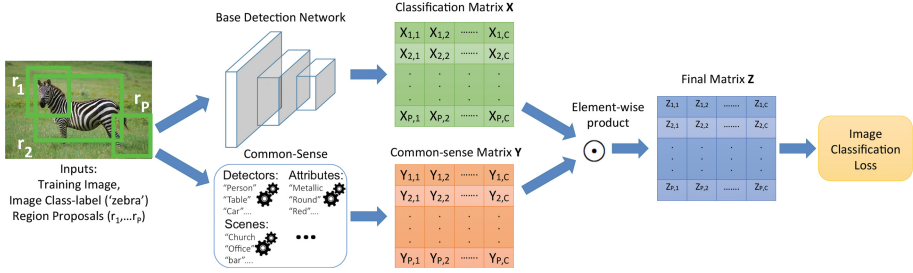


Fig. 3. Proposed framework for transferring common-sense knowledge for object detection. The base detection network computes a classification matrix $X_{P \times C}$ without any bounding box annotations (Sect. 3.1). We introduce a common-sense matrix $Y_{P \times C}$ that modulates the probabilities of region proposals belonging to various classes based on common-sense knowledge (Sect. 3.2). The common-sense matrix is computed using readily-available knowledge base resources (Sect. 3.3).

previously-learned source objects. This idea is related to [20], which makes use of already known objects to discover new categories from unlabeled images.

Using external knowledge for vision tasks. Our field has witnessed the rise of several interesting knowledge bases, including ConceptNet [23], BabelNet [29], WordNet [27], WebChild [38], Visual Genome [19], ImSitu [44], etc. While resources like WebChild [38] and BabelNet [29] are created automatically by crawling the web, others are generated with crowd-sourced effort. The key advantage of these resources is that they contain freely-available rich knowledge.

Such external knowledge bases have been used in several vision tasks including image classification [7, 25], VQA [41, 47], visual relationship detection [24, 30], and modeling object affordances [46]. However, there has been very limited work on using external knowledge for object detection [13], especially in the transfer learning setting in which bounding box annotations for the target classes are lacking. Tang et al. [39] use word2vec semantic similarity between classes to perform domain transfer between a classifier and a detector. In contrast, we go beyond using semantic similarity and explore spatial, scene, and attribute cues.

3 Proposed Approach

In this section, we first briefly describe the base detection network used in our framework and then explain our proposed approach for injecting common-sense knowledge into it. Finally, we describe our process for automatically gathering the different types of common-sense knowledge from existing resources.

3.1 Base Detection Network

Our idea of transferring common-sense knowledge to improve object detection is generic and could be incorporated into any detection approach that learns

from image-level labels. In our work, we use an object proposal classification and ranking framework based on [5] for its simplicity and end-to-end nature.

The initial layers of the network consist of convolution layers followed by spatial pyramid pooling layers to pool features corresponding to image region proposals (r_i). After pooling, the network has two data streams: the *recognition* stream assigns a classification score for each region proposal by applying a softmax over the classes to produce a $P \times C$ recognition matrix X_r , whereas the *detection* stream assigns probability of a region proposal to be selected for a specific class by applying a softmax over the proposals to produce a $P \times C$ detection matrix X_d . The final probability for each proposal to belong to different classes is computed by taking their element-wise dot product $X = X_r \odot X_d$. The network takes P proposals of a training image as input and outputs the probability for each of them to belong to C classes. This is shown as a $P \times C$ classification matrix X in Fig. 3. Note that the network learns to detect objects while being trained for the image classification task. The image-level class probabilities are obtained by summing the probabilities of each class (c_i) over the proposals:

$$Prob(c_i) = \sum_{n=1}^P X_{r_n, c_i}, i \in (1, C),$$

where X_{r_n, c_i} is the probability of proposal r_n belonging to class c_i . A binary cross-entropy loss is applied over the probabilities to learn the detection models.

3.2 Transferring Common-Sense

In order to transfer common-sense knowledge from the source categories with both image and bounding box annotations to the target categories with only image-level annotations, we augment the above base detection network with a novel *common-sense* matrix Y of size $P \times C$ (analogous to the classification matrix $X_{P \times C}$). Each element of Y_{r_n, c_i} can be thought of as representing a ‘prior’ probability of a proposal r_n belonging to class c_i according to common-sense knowledge (see Fig. 3). We will maintain a separate common-sense matrix for each type of common-sense (*similarity*, *attribute*, etc.) and later (Sect. 3.3) describe the details for acquiring and merging these matrices.

Assuming we have access to this common-sense matrix Y , we utilize this information by taking an element-wise dot product of it with the classification matrix (X) to create a resultant matrix $Z_{P \times C}$:

$$Prob(c_i) = \sum_{n=1}^P Y_{r_n, c_i} * X_{r_n, c_i} = \sum_{n=1}^P Z_{r_n, c_i}, i \in (1, C),$$

which now will be used for obtaining the image-level class probabilities over which a binary cross-entropy loss is applied.

For example, in Fig. 3, the *attribute* common-sense matrix (which would encode the common-sense that a ‘zebra’ is *striped*) would have a low prior probability ($Y_{r_P, zebra}$) for the proposal r_P to be a ‘zebra’. And, the *class-similarity*

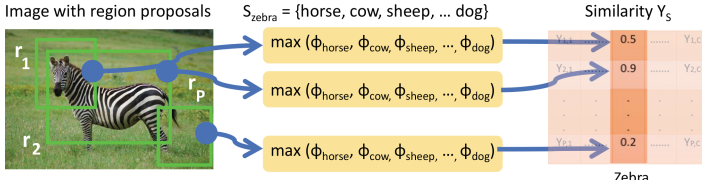


Fig. 4. Similarity Common-sense: For computing the $Y_s(., zebra)$ values, all the proposals of the input image are scored by the detectors of zebra’s semantically similar classes S_{zebra} . Observe that proposal r_2 , which contains the zebra’s full-body gets the highest Y_s value.

common-sense (which would encode the common-sense that a ‘zebra’ is similar to a ‘horse’) would have a low value for the zebra head proposal ($Y_{r_1, zebra}$) compared to the zebra full-body proposal ($Y_{r_2, zebra}$).

Intuitively, the common-sense matrix Y influences the values of the classification matrix X during training and over time the common-sense priorities are transferred from Y to X . To drive home this intuition, we take the example in Fig. 3. The $Prob_{zebra}$ should be high for this training image as it contains a ‘zebra’, i.e., $\sum_{n=1}^P Z_{r_n, zebra} = \sum_{n=1}^P Y_{r_n, zebra} * X_{r_n, zebra}$ should be high. This can be easily achieved if both $Y_{r_n, zebra}$ and $X_{r_n, zebra}$ are high. In this case $Y_{r_2, zebra}$ is high according to common-sense which in turn encourages the network to have high value for $X_{r_2, zebra}$. At the same time, due to a low value for $Y_{r_1, zebra}$ and $Y_{r_p, zebra}$, the network is discouraged to have a high value for $X_{r_1, zebra}$ and $X_{r_p, zebra}$. Therefore, during training itself, the network learns to incorporate the common-sense prior information of the Y matrix into the X matrix.

3.3 Acquiring Common-Sense

Now that we have seen how we transfer common-sense information (i.e., matrix Y) into the base detection framework, we next explain our approach for automatically gathering this matrix using existing knowledge resources.

Class Similarity Common-sense. Our goal here is to leverage the semantic similarity of a new target class to previously-learned source classes. For example, as a ‘zebra’ is semantically similar to a ‘horse’, the proposals that are scored higher by a ‘horse’ detector should be more probable of being a ‘zebra’. More generally, for any class c_i , the proposals looking similar to its semantically similar classes should receive higher prior for c_i .

To construct the class-similarity common-sense matrix Y_s , we tap into the readily-available set of pre-trained detectors (ϕ) for the source object classes (C_{source}) in the PASCAL VOC [12] knowledge base. Let c_i be one of the new target object classes for which we are trying to learn a detector with just image-level labels. To find the set of semantically similar source classes (i.e., $S_{c_i} \subset C_{source}$) to c_i , we represent all the classes (c_i as well as C_{source}) using their word2vec textual feature representation [26] and then compute the cosine-similarity between

them. We choose all classes from C_{source} with cosine-similarity above a threshold (0.35) as S_{c_i} .

We use the detectors of the classes in S_{c_i} to compute the values in Y_s as:

$$Y_s(r_n, c_i) = \max_{c_j \in S_{c_i}} \phi_{c_j}(r_n), n \in (1, P).$$

Specifically, we set the value $Y_s(r_n, c_i)$ for a proposal r_n to be of class c_i as being equal to the maximum detection probability of the classes similar to c_i . Figure 4 illustrates how the class-similarity common-sense Y_s is assigned in case of the target ‘zebra’ object class, where S_{zebra} consists of the source object classes {‘horse’, ‘cow’, ‘sheep’, ‘dog’, ‘cat’, ‘bird’}. Observe that a correct proposal (i.e., r_2 containing the ‘zebra’ full-body) gets the highest similarity common-sense probability as it is scored higher by the object detectors in S_{zebra} .




Fig. 5. Attribute Common-sense: For computing $Y_a(., clock)$ values, we apply the attribute classifiers for the common attributes of ‘clock’ on the proposals. Observe that r_1 containing a ‘clock’ with *white* & *round* attributes gets the highest Y_a value.

Attribute Common-sense. Attributes are mid-level semantic visual concepts (e.g., *furry*, *red*, *round*, etc.) that are shareable across object categories [14]. For example, an ‘apple’ is usually *red*, a ‘clock’ is typically *round*, etc. Therefore region proposals that possess the characteristic attributes of a specific class should be more probable to belong to that class.

To build the attribute common-sense matrix Y_a , we leverage the pre-trained set of attribute classifiers (θ) from the ImageNet Attribute [35] knowledge base and the readily-available set of object-attribute relationships from the Visual Genome [19] knowledge base. Let c_i be one of the new target classes for which we are trying to learn a detector and A_{c_i} be its set of common attributes (determined by the frequency by which it is used to describe c_i). The classifiers of the attributes in A_{c_i} are used to compute the values of the matrix Y_a as

$$Y_a(r_n, c_i) = \max_{a_j \in A_{c_i}} \theta_{a_j}(r_n), n \in (1, P).$$

As the attributes in ImageNet [35] knowledge base (that we use in this work) have been additionally grouped into sets of *color* A^{col} , *shape* A^{shape} , and *texture* A^{text} attributes, we adopt this information by updating $Y_a(r_n, c_i)$ as: $Y_a(r_n, c_i) = \text{mean}(Y_a^{col}(r_n, c_i), Y_a^{shape}(r_n, c_i), Y_a^{text}(r_n, c_i))$, where Y^{col} , Y^{shape} , Y^{text} have been computed over the A^{col} , A^{shape} , A^{text} domains. In Fig. 5, for the ‘clock’




Fig. 6. Spatial common-sense distribution: We model the relative location distributions $\gamma_{c_j, rel}^L$ encoding the pixel-wise probability of all other source objects for a given object c_j (shown as a red box) under a given relation rel . (Color figure online)

class, the proposal r_1 containing its attributes, i.e., *white/black* and *round*, get the highest Y value, while the other proposals (r_2, r_P) get lower values as they lack the characteristic ‘clock’ attributes.

Spatial Common-sense. In our day-to-day experience objects often appear in characteristic spatial relations with other objects. For example, a ‘bowl’ is typically *on* a ‘table’, a ‘backpack’ is typically *behind* a ‘person’, etc. Therefore region proposals that have the characteristic spatial relation of a target class to other source classes should be more probable of belonging to that target class.

To obtain the spatial-relation common-sense matrix Y_{sp} , we utilize the information about relative locations and sizes of *source* object classes C_{source} in the Visual Genome [19] knowledge base that contains visually-grounded triplets {object1, relation, object2}.

For each class c_j in C_{source} , we model the relative location distributions $\gamma_{c_j, rel}^L$ encoding the pixel-wise probability of all other source objects under a given relation rel . For example, Fig. 6 shows the distribution of objects with respect to the ‘along’ relationship for the source ‘person’ object class. In a similar way, for each class c_j in C_{source} and a given relation rel , we also model the relative size distributions $\gamma_{c_j, rel}^S$. Note that these distributions need to be learned just once using *only the source classes* and then can be reused for any of the target classes (i.e., without bounding box annotations for C_{target} classes).

For a new target class c_i from C_{target} , we first gather its most common relation with respect to source classes c_j in C_{source} using the {object1, relation, object2} triplet information from [19]¹ and then compute the Y_{sp} matrix as:

$$Y_{sp}(r_n, c_i) = \max_{c_j \in C_{vis}} \frac{1}{2} (\gamma_{c_j, rel}^L(\mathbf{x}_{r_n}^{center}) + \gamma_{c_j, rel}^S(area_{r_n})),$$

where \mathbf{x}^{center} and $area$ denote the center coordinate and the size of a proposal, and C_{vis} are the subset of classes from C_{source} that are visible in the given image with their locations determined by running the pre-trained detectors ϕ .

In Fig. 7, the proposal r_P gets a higher Y_{sp} value for the ‘skateboard’ class as it is in sync with the ‘along’ relation under the $\gamma_{person, along}^L$ distribution.

¹ For any target class, we gather this information from existing knowledge bases [19, 38] by analyzing the rel that is commonly used to associate it with the source classes.




Fig. 7. Spatial Common-sense: The $Y_{sp}(\cdot, \text{skateboard})$ values are computed by measuring the spatial relationship of proposals with respect to the source ‘person’ class. Based on our knowledge base [19], ‘along’ is the most common relation between the ‘skateboard’ and ‘person’ class. Observe that the proposal r_P gets the highest Y_{sp} value.

Scene Common-sense. Some objects appear more in certain scenes than others; e.g., a ‘surfboard’ is more likely to be found on a *beach*. Hence, images depicting scenes associated with class c_i are more likely to contain instances of c_i .

To obtain the scene common-sense matrix Y_{sc} , we leverage the SceneUNderstanding (SUN) [42] and Places [45] knowledge bases. These databases not only contain information about commonly occurring scene labels for the different object classes but also provide access to pre-trained scene classifiers (β). Let SC_{c_i} denote the set of scene labels associated with a new target object class c_i , and I be the given image containing the proposals, then Y_{sc} can be computed as: $Y_{sc}(r_n, c_i) = \sum_{s_j \in SC_{c_i}} \beta_{s_j}(I), n \in (1, P)$. All proposals r_n in image I get high prior for class c_i if I depicts the scene that frequently contains instances of class c_i . Note that this scene common-sense knowledge would be helpful in case of noisy image-level labels (e.g., in the webly-supervised setting), and may not be relevant when we already have clean human-annotated labels that indicate the presence/absence of objects in images.

4 Experimental Results and Analysis

In this section, we report quantitative and qualitative analyses for validating the benefit of transferring common-sense knowledge for object detection. We also conduct ablation studies to dissect the various components of our approach.

Dataset. Recent state-of-the-art transfer learning approaches for object detection [17,39] have limited their analysis to ImageNet [34]. While this dataset has enabled substantial progress in this domain, we believe the time is ripe for advancing the transfer-learning paradigm to the next level, i.e., the more complex MS COCO [22] benchmark. The MS COCO dataset is challenging not only in terms of its diversity (non-iconic views, varying sizes, etc.) but also in number of classes, and thus poses a tough challenge for methods that do not have access to bounding box annotations for the target classes. We believe our idea of leveraging multiple common-sense would be more relevant to address this challenge.

Recall that, in our framework, we use the PASCAL VOC object detectors as one of our sources of common-sense knowledge (i.e., for similarity and spatial). Hence, we avoid using the 20 VOC classes (C_{source}) within the MS COCO dataset and focus our analysis on the remaining 60 object classes (C_{target}). We train our network with MS COCO 2015 training images, evaluate on the validation images, and use the standard performance metrics (mAP with 50% IoU threshold).

Implementation details. Our base network of [4] is initialized with VGG-CNN-F [6]. We train the network with a learning rate of 10^{-5} for the first 20 epochs and 10^{-6} for the next 10 epochs. During training, the images are randomly flipped horizontally and scaled to one of 5 scales: 800×608 , 656×496 , 544×400 , 960×720 , and 1152×864 . During testing, we average the corresponding 10 detection scores and apply NMS with an overlap threshold of 0.4 after filtering out proposals with probability less than 10^{-4} . We use MCG [2] object proposals. In order to combine the common-sense matrices (Y_a, Y_{sp}, Y_s), we take the average of the three matrices and obtain Y_{all} , which is used to train the base network. Although our approach involves a few task-specific insights to effectively leverage common-sense knowledge, we found our approach to be robust to a wide range of choices when constructing these common-sense matrices (e.g. number of similar known classes for similarity common-sense, number of common relations between known/unknown classes for spatial common-sense, etc). For similarity common-sense, we represent each class name using a 300-d word2vec representation (obtained using pre-trained model on Google news dataset [26]).

Table 1. Detection results on MS COCO validation set. Our proposed approach substantially improves over existing transfer-learning methods.

Methods	AP	AP^S	AP^M	AP^L
Classification Network ('No Adapt' [17])	3.4	0.9	2.9	6.1
LSDA [17]	4.6	1.2	5.1	7.8
LSDA+Semantic [39]	4.7	1.1	5.1	8.0
LSDA+MIL [18]	5.9	1.5	8.3	10.7
Fine-tuned Detection Network	10.8	1.2	8.9	18.6
Proposed Approach	14.4	2.0	12.8	24.9
Oracle: Full Detection Network [16]	25.2	5.8	26.0	41.6

4.1 Quantitative Results

Table 1 presents the results obtained using our approach and compares it to other relevant approaches. As an upper-bound, we also include the fully-supervised Fast-RCNN [16] detection result (using VGG-CNN-F and trained with bounding box annotations for the target classes) obtaining 25.2% mAP. While our

approach falls well short of this upper bound, it reveals to us the possibility of using common-sense for bridging the gap to fully-supervised methods.

Comparison to Transfer-learning. The most relevant state-of-the-art transfer learning methods are LSDA [17], [18] and [39]. However, as [17, 18, 39] neither report results on the challenging MS COCO dataset nor provide their training code, we re-implemented LSDA [17]². Running this baseline on the 60 MS COCO C_{target} classes yields 4.6% mAP, which is substantially inferior than our proposed approach (14.4%). We hypothesize the poor performance of the LSDA frameworks on the MS COCO dataset to the following reasons:

- (i) LSDA approaches [17, 18, 39] are inherently limited to exploiting only *similarity* common-sense. While similarity suffices in case of object-centric datasets, more richer common-sense (such as *attribute*, *spatial*, etc., that we use) needs to be leveraged when dealing with complex scene-centric datasets. Further when the size of C_{source} is small (e.g., set of 20 classes in our MS COCO experiments), similarity fails to work well in gathering meaningful neighbors between C_{target} and C_{source} classes. As a result, LSDA methods cannot transform the classifier successfully. Particularly, the class-invariant transformation of the weights from conv1 to fc7 layers in LSDA will not generalize well when the similarity overlap between the classes in C_{source} and C_{target} is poor. Our approach alleviates this concern by explicitly using detection probabilities of the C_{source} classes for each image (rather than doing a weight transformation), and also by using other common-sense knowledge (*attribute*, *spatial*, etc., which would weigh more in the absence of similar C_{source} classes).
- (ii) As MS COCO has images with multiple categories and small objects, the initial classifier network learned in the LSDA methods is poor (giving only 3.4% mAP, see Table 1 top-row) and therefore transforming it results in a poor detector. Our approach alleviates this concern by employing a proposal classification and ranking detection network (11.7% mAP, see Table 2, ‘Base network’).
- (iii) Finally, class co-occurrences in images poses a formidable challenge for the LSDA methods. For example, in case of the target class ‘spoon’, the top nearest neighbor (among the 20 C_{source} classes) is ‘dining-table’, which often co-occurs with it and thus transferring its detection knowledge will cause confusion to the detector. Our approach is robust to such cases of incorrect similarities as the similarity information is only used as a ‘prior’ during our base detection network training.

Transfer learning by fine-tuned detection network. We also explored an alternate transfer-learning strategy, where we initialize our base detection network (Sect. 3.1) with pre-trained Fast-RCNN weights from the source categories and

² Similar setting as ours: use image-level labels for all training images of MS COCO and VOC, use bounding box for only 20 VOC classes, and VGG-CNN-F as base network.

then fine-tune it on the target classes using only image-level labels. While this method produces a relatively higher mAP of 10.8% than LSDA, it is still lower than our approach. We believe this is due to the network weights getting overfit to the 20 C_{source} classes and subsequently failing to generalize well to the 60 C_{target} classes. In contrast, our approach does not have this overfitting issue as we initialize our base detection network with the weights from a more general network (trained on 1000 ImageNet classes) and then use the 20 C_{source} classes pre-trained detection models only for computing similarity common-sense.

Alternate approaches for leveraging common-sense. To analyze the significance of our proposed approach for leveraging common-sense, we also studied alternative strategies for leveraging common sense cues that use exactly the same external knowledge/bounding box information as our approach.

The first strategy uses common-sense as a contextual post-processing tool [9]. Specifically, for a test image, we compute the common-sense matrix Y , and then modulate its classification matrix X via element-wise multiplication to produce the final score matrix: $Z_{test} = X_{test} \cdot Y_{test}$. Table 2 (‘Post-process’) displays the result obtained using this post-processing strategy, which obtains 11.8% (and 14.1% when common-sense matrix Y is also used during training). Observe that the post-processing result is lower compared to our approach, which transfers common-sense during training only and not during testing. This indicates that X has already incorporated the common-sense knowledge during training, and thus using Y is redundant at test time. When Y is introduced during testing, it is difficult for common sense to fix any incorrect biases that the detector may have learned (e.g., focusing only on the most discriminative part of an object). It may even hurt as any mistakes in common-sense information Y cannot be circumvented when used directly at test time. In contrast, by transferring common sense during training, it can guide the algorithm to *learn* to focus on the correct image regions.

The second alternate strategy for leveraging common-sense analyzes the possibility of improving detection performance by simply having access to pre-trained object/spatial/attribute classifiers. In this case, we trivially append a 45-dimensional feature to the fc7 appearance feature of each proposal during training. The first 20 dimensions correspond to the detection probabilities of the

Table 2. Analysis on the MS COCO dataset: ‘Base network’ is the result obtained by using our base detection network using no common-sense. ‘Alternatives’ are the alternate strategies for leveraging common-sense knowledge. Their performance is lower than our model demonstrating the benefit of our proposed approach. ‘Ablations’ show the improvement obtained using each common-sense cue over the base network. Combining all common-sense cues produces the best result indicating their complementarity.

Method	Base network	Alternatives		Ablations				Ours
		Post-process	Feature	Attr	Spatial	Sim	Joint	+Sim(Bbox)
mAP	11.7	11.8/14.1	12.7	12.2	13.0	13.7	14.1	14.4

20 C_{source} classes and the remaining 25 dimensions correspond to the attribute probabilities of the classifiers (θ) pre-trained on the ImageNet attribute knowledge base. While this model yields a boost of 1.0% mAP compared to the base detection network (see Table 2 ‘Feature’), it is 1.7% lower than our proposed model. This reveals that merely concatenating external knowledge features to visual appearance features is insufficient.

4.2 Ablation and Qualitative Analysis

We also analyzed the importance of the various common-sense cues in our proposed approach. Table 2 ‘Ablations’ shows that each common sense gives a boost over the base network, which only relies on appearance cues. Among the individual cues, we see that similarity helps the most. (Scene cue was not explored in this scenario and its influence will be analyzed in the webly-supervised setting.) Combining attribute, spatial, and similarity common sense (‘Joint’) leads to a greater boost of 2.4% mAP, which shows their complementarity. Finally, we also borrow the bounding box regressors trained on the similar C_{source} classes and apply them to the C_{target} classes, which further boosts performance to 14.4%.

Taking a closer look at individual classes, we find that using the attribute common-sense for ‘oven’ (that is usually *white/black*) results in a boost of 7.3% AP. By using the spatial common-sense for ‘frisbee’ with respect to the source object ‘person’, we get an improvement of 10.5% in AP, whereas using the spatial relation for a ‘bowl’ with respect to the source object ‘table’ gives a boost of 2.2%. For ‘giraffe’ and ‘bed’, we use the common-sense that they are semantically similar to {‘sheep’, ‘horse’, ‘dog’, ‘cow’, ‘bird’, ‘cat’}, and ‘couch’ respectively, which leads to an improvement of 28.7% and 12.1% AP, respectively.

We next analyze the importance of using word2vec for similarity common-sense. For this, we replace word2vec similarity with visual similarity, which results in 12.1% compared to our 13.7%. Word2vec similarity is more robust than visual similarity (used in LSDA [17]), particularly for challenging dataset like MS COCO (with small objects and co-occurring objects from different classes). We also tried WordNet [27] treepath based similarity which also gives an inferior result of 13.1%.

Figure 8 shows some qualitative detections (for each of the common-sense) produced by our approach (green box) and compares them to competing base-lines. We can observe that using common-sense helps improve performance. For example, by using spatial common-sense, our approach gets rid of the co-occurring background for ‘frisbee’ and ‘surfboard’ (person and water, respectively). Interestingly, for ‘frisbee’, our approach uses the spatial common-sense that a ‘frisbee’ often spatially-overlaps with a ‘person’ in order to learn its appearance during training. It is then able to detect a ‘frisbee’ at test time even when a ‘person’ is not present (see ‘frisbee’ image in the second column). This indicates that while our approach leverages common-sense as a prior during training to learn about the object’s appearance, at the same time, our network is not dependent on common-sense during testing. Figure 9 shows some failure cases.

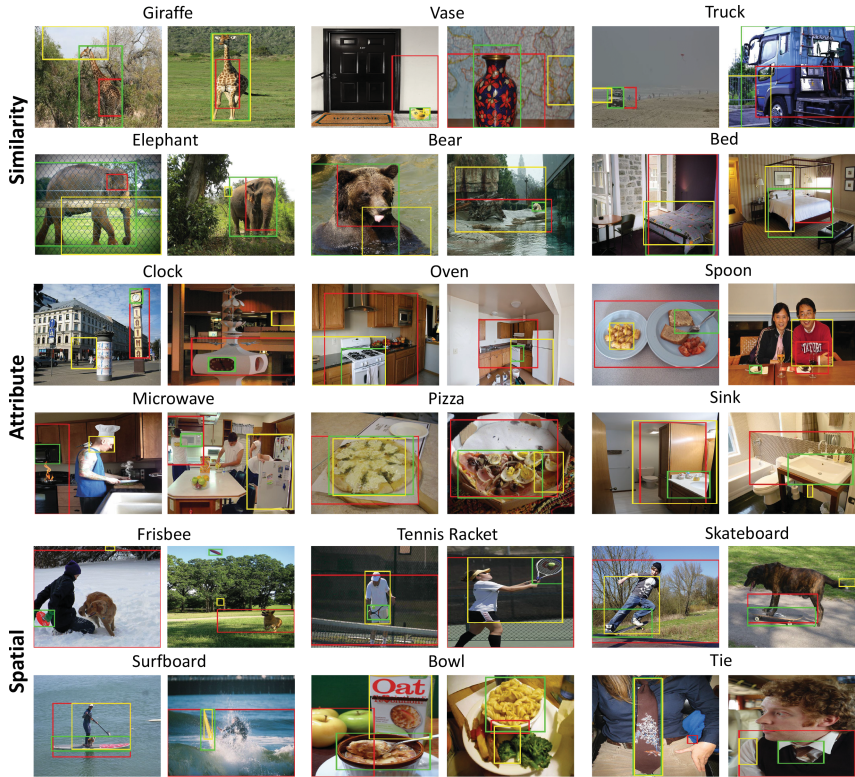


Fig. 8. Qualitative detection results on MS COCO (Ours: green boxes; Base network: red; LSDA+Semantic [39]: Yellow): Observe that our approach produces better detections than the base network for all three common-sense. For ‘giraffe’ and ‘elephant’, by using similarity common-sense (i.e., being similar to other animal categories in C_{source}), our approach detects the full body extent rather than localizing a discriminative body part. By using spatial and attribute common-sense, e.g., ‘clock’ being *round*, ‘spoon’ being *metallic*, and ‘microwave’ being *white/black*, we get better detections. (Color figure online)




Fig. 9. Example failures: Our approach fails when the object-of-interest is hardly-visible (‘handbag’) or when source objects with similar attribute (*metallic*) are cluttered together (‘spoon’). For ‘wine glass’, we falsely detect the ‘bottle’ because during training we provided the common-sense that wine-glass is semantically similar to a bottle.

Towards webly-supervised detection using common-sense supervision. What happens when we apply our method in cases when we do not even have explicit human-annotated image labels for the target class? This is exactly the setting studied in the webly-supervised domain where images retrieved from the web are used for training detection models.

We conducted a preliminary investigation wherein we ran our proposed approach on training images retrieved from the web (i.e., instead of the MS COCO training imageset). As images retrieved from web are potentially noisy, common-sense knowledge should be particularly useful in mitigating the noise. Our preliminary results indicate that our proposed idea is promising even in the webly setting (Base network [5]: 6.8%, vs. Ours: 8.3%).

Further, to analyze potential concerns about the generalizability of our acquired common-sense knowledge, we also tested these webly+commonsense models on the ImageNet 200 detection valset [34]. Even in this case, our approach yields interesting performance gains (Base network [5]: 6.2%, vs. Ours: 8.8%).

5 Conclusion

In this paper, we presented DOCK, a novel approach for transferring common-sense knowledge from a set of categories with bounding box annotations to a set of categories that only have image-level annotations for object detection. We explored how different common-sense cues based on similarity, attributes, spatial relations, and scene could be jointly used to guide the algorithm towards improved object localization. Our experiments showed that common-sense knowledge can improve detection performance on the challenging MS COCO dataset. We hope our work will spur further exciting research in this domain.

Acknowledgements. This work is in part supported by ONR N00014-13-1-0720, NSF IIS-1338054, NSF IIS-1748387, NSF IIS-1751206, NSF-1652052, NRI-1637479, ARO YIP W911NF-17-1-0410, Allen Distinguished Investigator Award, Allen Institute for AI, Microsoft Azure Research Award and the GPUs donated by NVIDIA. A part of this work was done while Krishna Kumar Singh was an intern at Allen Institute for AI.

References

1. Toothbrush. <http://cocodataset.org/#explore?id=315486>
2. Arbeláez, P., Pont-Tuset, J., Barron, J., Marques, F., Malik, J.: Multiscale combinatorial grouping. In: CVPR (2014)
3. Aytar, Y., Zisserman, A.: Tabula rasa: model transfer for object category detection. In: ICCV (2011)
4. Bilen, H., Pedersoli, M., Nambodiri, V., Tuytelaars, T., Gool, L.V.: Object classification with adaptable regions. In: CVPR (2014)
5. Bilen, H., Vedaldi, A.: Weakly supervised deep detection networks. In: CVPR (2016)

6. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: delving deep into convolutional nets. In: BMVC (2014)
7. Chowdhury, S.N., Tandon, N., Ferhatosmanoglu, H., Weikum, G.: VISIR: visual and semantic image label refinement. In: WSDM (2018)
8. Desai, C., Ramanan, D., Fowlkes, C.: Discriminative models for multi-class object layout. *IJCV* **95**(1), 1–12 (2011)
9. Divvala, S.K., Hoiem, D., Hays, J.H., Efros, A.A., Hebert, M.: An empirical study of context in object detection. In: CVPR (2009)
10. Doersch, C., Gupta, A., Efros, A.A.: Context as supervisory signal: discovering objects with predictable context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8691, pp. 362–377. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10578-9_24
11. Donahue, J., Hoffman, J., Rodner, E., Saenko, K., Darrell, T.: Semi-supervised domain adaptation with instance constraints. In: CVPR (2013)
12. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>
13. Fang, Y., Kuan, K., Lin, J., Tan, C., Chandrasekhar, V.: Object detection meets knowledge graphs. In: IJCAI (2017)
14. Farhadi, A., Endres, I., Hoiem, D., Forsyth, D.: Describing objects by their attributes. In: CVPR (2009)
15. Gidaris, S., Komodakis, N.: Object detection via a multi-region and semantic segmentation-aware CNN model. In: ICCV (2015)
16. Girshick, R.: Fast R-CNN. In: ICCV (2015)
17. Hoffman, J., et al.: LSDA: Large scale detection through adaptation. In: NIPS (2014)
18. Hoffman, J., Pathak, D., Darrell, T., Saenko, K.: Detector discovery in the wild: joint multiple instance and representation learning. In: CVPR (2015)
19. Krishna, R., et al.: Visual genome: connecting language and vision using crowd-sourced dense image annotations. *IJCV* **123**(1), 32–73 (2017)
20. Lee, Y.J., Grauman, K.: Object-graphs for context-aware visual category discovery. *TPAMI* **34**(2), 346–358 (2012)
21. Lim, J.J., Salakhutdinov, R.R., Torralba, A.: Transfer learning by borrowing examples for multiclass object detection. In: NIPS (2011)
22. Lin, T.-Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
23. Liu, H., Singh, P.: ConceptNet - a practical commonsense reasoning tool-kit. *BT Technol. J.* **22**(4), 211–226 (2004)
24. Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 852–869. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46448-0_51
25. Marino, K., Salakhutdinov, R., Gupta, A.: The more you know: using knowledge graphs for image classification. In: CVPR (2017)
26. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: NIPS (2013)
27. Miller, G.A.: WordNet: a lexical database for English. *Commun. ACM* **38**(11), 39–41 (1995)
28. Mottaghi, R., et al.: The role of context for object detection and semantic segmentation in the wild. In: CVPR (2014)

29. Navigli, R., Ponzetto, S.P.: BabelNet: the automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artif. Intell.* **193**, 217–250 (2012)
30. Plummer, B.A., Mallya, A., Cervantes, C.M., Hockenmaier, J., Lazebnik, S.: Phrase localization and visual relationship detection with comprehensive linguistic cues. In: *ICCV* (2017)
31. Rabinovich, A., Vedaldi, A., Galleguillos, C., Wiewiora, E., Belongie, S.: Objects in context. In: *ICCV* (2007)
32. Rochan, M., Wang, Y.: Weakly supervised localization of novel objects using appearance transfer. In: *CVPR* (2015)
33. Russakovsky, O., Lin, Y., Yu, K., Fei-Fei, L.: Object-centric spatial pooling for image classification. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) *ECCV 2012*. LNCS, pp. 1–15. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33709-3_1
34. Russakovsky, O., et al.: ImageNet large scale visual recognition challenge. *IJCV* **115**, 211–252 (2015)
35. Russakovsky, O., Fei-Fei, L.: Attribute learning in large-scale datasets. In: Kutulakos, K.N. (ed.) *ECCV 2010*. LNCS, vol. 6553, pp. 1–14. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35749-7_1
36. Salakhutdinov, R., Torralba, A., Tenenbaum, J.: Learning to share visual appearance for multiclass object detection. In: *CVPR* (2011)
37. Shi, Z., Siva, P., Xiang, T.: Transfer learning by ranking for weakly supervised object annotation. In: *BMVC* (2012)
38. Tandon, N., de Melo, G., Suchanek, F., Weikum, G.: WebChild: Harvesting and organizing commonsense knowledge from the web. In: *WSDM* (2014)
39. Tang, Y., Wang, J., Gao, B., Dellandréa, E., Gaizauskas, R., Chen, L.: Large scale semi-supervised object detection using visual and semantic knowledge transfer. In: *CVPR* (2016)
40. Wang, Y.X., Hebert, M.: Model recommendation: generating object detectors from few samples. In: *CVPR* (2015)
41. Wu, Q., Wang, P., Shen, C., Dick, A., van den Hengel, A.: Ask me anything: free-form visual question answering based on knowledge from external sources. In: *CVPR* (2016)
42. Xiao, J., Ehinger, K.A., Hays, J., Torralba, A., Oliva, A.: Sun database: exploring a large collection of scene categories. *IJCV* **119**, 3–22 (2016)
43. Xu, J., Ramos, S., Vázquez, D., López, A.M.: Domain adaptation of deformable part-based models. *PAMI* **36**, 2367–2380 (2014)
44. Yatskar, M., Zettlemoyer, L., Farhadi, A.: Situation recognition: visual semantic role labeling for image understanding. In: *CVPR* (2016)
45. Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., Oliva, A.: Learning deep features for scene recognition using places database. In: *NIPS* (2014)
46. Zhu, Y., Fathi, A., Fei-Fei, L.: Reasoning about object affordances in a knowledge base representation. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8690, pp. 408–424. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10605-2_27
47. Zhu, Y., Zhang, C., Ré, C., Fei-Fei, L.: Building a large-scale multimodal knowledge base system for answering visual queries. *arXiv preprint* (2015)