

Universal-RCNN: Universal Object Detector via Transferable Graph R-CNN

Hang Xu,^{1*} Linpu Fang,^{2*} Xiaodan Liang,^{3†}
 Wenxiong Kang,² Zhenguo Li¹
¹Huawei Noah’s Ark Lab
²South China University of Technology
³Sun Yat-Sen University

Abstract

The dominant object detection approaches treat each dataset separately and fit towards a specific domain, which cannot adapt to other domains without extensive retraining. In this paper, we address the problem of designing a universal object detection model that exploits diverse category granularity from multiple domains and predict all kinds of categories in one system. Existing works treat this problem by integrating multiple detection branches upon one shared backbone network. However, this paradigm overlooks the crucial semantic correlations between multiple domains, such as categories hierarchy, visual similarity, and linguistic relationship. To address these drawbacks, we present a novel universal object detector called Universal-RCNN that incorporates graph transfer learning for propagating relevant semantic information across multiple datasets to reach semantic coherency. Specifically, we first generate a global semantic pool by integrating all high-level semantic representation of all the categories. Then an Intra-Domain Reasoning Module learns and propagates the sparse graph representation within one dataset guided by a spatial-aware GCN. Finally, an Inter-Domain Transfer Module is proposed to exploit diverse transfer dependencies across all domains and enhance the regional feature representation by attending and transferring semantic contexts globally. Extensive experiments demonstrate that the proposed method significantly outperforms multiple-branch models and achieves the state-of-the-art results on multiple object detection benchmarks (mAP: 49.1% on COCO).

Introduction

Object detection is a fundamental vision task that recognizes object’s location and category in an image. Modern object detectors widely benefits autonomous vehicles, surveillance camera, mobile phone, to name a few. According to the individual application, datasets annotated with different categories were created to train a highly-specific and distinct detector e.g. PASCAL VOC (Everingham et al. 2010) (20 categories), MSCOCO (Lin et al. 2014) (80 categories), Visual Genome (VG) (Krishna et al. 2016) (more than 33K

*Both authors contributed equally to this work.

†Corresponding Author: xliang328@gmail.com

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

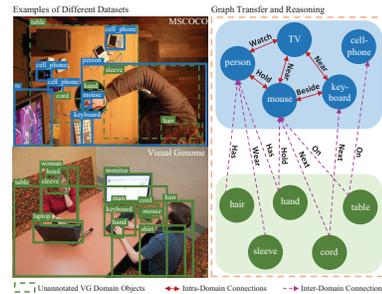


Figure 1: To identify objects in the upper image from COCO, we could use graph transfer and reasoning across COCO and Visual Genome (VG) (bottom): 1) Objects like “TV”, and “keyboard” in the upper image could help to identify “mouse” and “person”, which requires reasoning and this inspires our Intra-Domain Reasoning Module (blue). 2) Although “table”, “hand”, and “cord” are not annotated in the COCO image but present in VG, recognition of them could help to recognize “mouse”, “keyboard”, which motivates us the Inter-Domain Transfer Module (green to blue).

categories) and BDD (Yu et al. 2018) (10 categories). Those highly-tuned networks have sacrificed the generalization capability and only fit towards each dataset domain. It is impossible to directly adapt the model trained on one dataset to another related task, and thus requires new data annotations and additional computation to train each specific model. In contrast, human is capable of identifying all kinds of objects precisely under complex circumstances and reaches a holistic recognition across different domains. This may be due to the remarkable reasoning and transferring ability about the relationship, visual similarity and categories hierarchy across scenes. This inspires us to explore how to endow the current detection system the ability of incorporating and transferring relevant semantic information across multiple domains in an effective way, in order to mimic human recognition procedure.

With different categories and granularity across domains, transferring relevant information and reasoning among objects can help to make a correct identification. An example

of the necessity of transferring relevant information between multiple datasets and domains can be found in Figure 1.

The most widely-used and straightforward solutions to universal detection would be to consider it as a multi-task learning problem, and integrate multiple branches upon one shared backbone network (He et al.; Gong et al.; Li et al.; Liang et al.; Dai, He, and Sun 2017; 2017; 2016; 2015; 2016). However, they overlook the crucial semantic correlations between multiple domains and regions, such as categories hierarchy, visual similarity, and linguistic relationship since feature-level information sharing can only help to extract a more robust feature. Recently, some works explore visual reasoning and try to combine different information or interactions between objects in one image (Hu et al.; Chen and Gupta; Wang et al.; Liu et al. 2018a; 2017; 2018; 2018b). For example, Jiang et al. recently try to incorporate semantic relation reasoning in large-scale detection by different kinds of knowledge forms. Hu et al. introduced the Relation Networks which use an adapted attention module to allow interaction between the object’s visual features. However, they did not consider inter-domain relationship and their fully-connected relation is inefficient and noisy by incorporating redundant and distracted relationships from irrelevant objects and backgrounds.

With the advancement of geometric deep learning (Bronstein et al.; Monti et al.; Velickovic et al. 2017; 2017; 2017), using graph seems to be the most appropriate way to model relation and interaction with its flexible structure. In this paper, we present a novel universal detection system Universal-RCNN that incorporates graph learning for propagating and transferring relevant semantic information across multiple datasets to reach semantic coherency. The proposed framework first generates a global semantic pool for all domains to integrate all high-level semantic representation of categories by distilling the weights of the object classifiers. Then an Intra-Domain Reasoning Module learns a sparse regional graph to encode the regional interaction and propagates the high-level semantic graph representation from the global pool within one dataset guided by a spatial-aware Graph Convolutional Neural Network (GCN). Furthermore, an Inter-Domain Transfer Module is proposed to exploit diverse transfer dependencies across all domains and enhance the regional feature representation by attentively transferring related semantic contexts from the semantic pool in one domain to another, which bridges the gap between domains, and effectively utilize the annotations of multiple datasets. In this work, we exploit various graph transfer dependencies such as attribute/relationship knowledge, and visual/linguistic similarity. Our Universal-RCNN thus enables adaptive global reasoning and transfers over regions in multiple domains. The regional feature is greatly enhanced by abundant relevant contextual Intra/Inter-domain information and the performance on each domain is then boosted by sharing and distilling essential characteristics across domains.

Extensive experiments demonstrate the effectiveness of the proposed method and achieve the state-of-the-art results on multiple object detection benchmarks. We observe a consistent gain over multiple-branch models. In particular, Universal-RCNN achieves around 16% of mAP improve-

ment on MSCOCO, 26% on VG, and 36% on ADE. The Universal-RCNN obtains 49.1% mAP on COCO *test-dev* with single-model result.

Related Work

Object Detection. Object detection is a core problem in computer vision. Most of the previous progress focus on developing new structures such as better feature fusion (Lin et al.; Liu et al.; Zhu, He, and Savvides; Xu et al. 2017; 2018a; 2019; 2019a) and better receptive field to improve feature representation (Luo et al.; Wang et al.; Li et al. 2017; 2018; 2018). However, their trained model cannot be applied directly to another related task without heavy fine-tuning.

Transfer Learning. Transfer learning (Zamir et al.; Peng et al.; Cui et al.; Xu et al. 2018; 2018; 2018; 2019b) tries to bridge the gap between different domains or tasks to reuse the information and mitigate the burden of manual labeling. Early work (Hoffman et al. 2014) tries to learn to transfer from the ImageNet’s classification network into object detection network with fewer categories. In Misra, Gupta, and Hebert, the classifier weights are constructed by regression or a small neural network from few-shot examples or different concepts.

Multi-task Learning. Multi-task learning aims at developing RCNN can sufficiently mining the semantic correlations systems that can provide multiple outputs simultaneously for an input (Tsai, Huang, and Salakhutdinov; He et al.; Kirillov et al.; Xiong et al.; Kendall, Gal, and Cipolla 2017; 2017; 2018; 2019; 2018). For example, Mask-RCNN solves the instance segmentation problem by considering two branches: one bounding box and classification head and another dense image prediction head. Xiao et al. introduced a multi-task network to handle heterogeneous annotations for unified perceptual scene parsing. However, these approaches simply create several branches separately for different tasks.

The Proposed Approach

Overview

In this paper, we introduce a unified detection framework to unify all kinds of detection annotations from different domains and tackle different detection tasks in one system, that is, detecting and predicting all kinds of categories defined from multiple domains using a single network while improving the detection performance of the individual domain. This framework can be implemented on any modern dominant detection system to further improve their performance by enhancing its original image features via graph transfer learning. An overview of our model can be found in Figure 2. Specifically, we first learn and propagate compact high-level semantic representation within one domain via Intra-Domain Graph Reasoning module, and then transfer and fuse the semantic information across multiple domains to enhance the feature representation via Inter-Graph Transfer module driven by learned transfer graph or human defined hierarchical categorical structures.

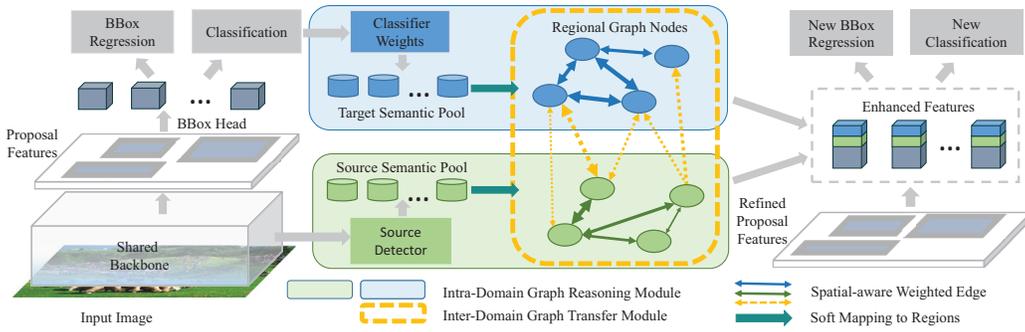


Figure 2: An overview of our Universal-RCNN. Built on the classification layer of base detection network, global semantic pools for each domain integrate all high-level semantic representation for each category by the weights of previous classifiers. Then an Intra-Domain Reasoning Module learns a sparse regional graph to encode the regional interaction and propagates the high-level semantic graph representation from the semantic pool within one domain. Furthermore, Inter-Domain Transfer Module exploits diverse transfer dependencies across domains and enhances the regional feature representation by attentively transferring related semantic contexts from source semantic pool to the target domain, which bridges the gap between domains and utilizes the annotations of multiple datasets. Finally, the outputs of modules concatenated with region proposal features are fed into the bounding-box regression layer and classification layer to obtain better detection results. Note that this figure only illustrates how the model works within a specific pair of domains. In the whole model, we have multiple pairs of the source/target domain for the inter-domain transfer modules.

Intra-Domain Reasoning Module

Given the extracted proposal visual features from the backbone, we introduce Intra-Domain Reasoning module to enhance local features for the domain T , by leveraging graph reasoning with key semantic and spatial relationships. More specifically, we first create a global semantic pool to integrate high-level semantic representation for each category by collecting the weights of the original classification layer. Then a region-to-region undirected graph $G : G_{T \rightarrow T} = \langle \mathcal{N}, \mathcal{E} \rangle$ for domain T is defined, where each node in \mathcal{N} corresponds to a region proposals and each edge $e_{i,j} \in \mathcal{E}$ encodes relationship between two nodes. An interpretable sparse adjacency matrix is learned from the visual feature which only retains the most relevant connections for recognition of the objects. Then semantic representations over the global semantic pool are mapped to each region and propagated with a learned structure of $G_{T \rightarrow T}$. Finally, the proposal features are enhanced and concatenated with the original features to obtain better detection results.

Learning Graph $G_{T \rightarrow T}$ We seek to learn the $\mathcal{E}_T \in \mathbb{R}^{N_r \times N_r}$ in $G_{T \rightarrow T}$ thus the proposal node neighborhoods can be determined. We aim to produce a graphical representation of the relationship (e.g. attribute similarities and interactions) between proposal regions which is relevant to the object detection. Given the regional visual features $\mathbf{f} = \{f_i\}_{i=1}^{N_r}$, $f_i \in \mathbb{R}^D$ of D dimension extracted from the backbone network, we first transform \mathbf{f} to a latent space \mathbf{Z} by linear transformation denoted by $z_i = \phi(\mathbf{f})$, $i = 1, 2, \dots, N_r$, where $z_i \in \mathbb{R}^L$, L is the dimension of the latent space and $\phi(\cdot)$ is a linear function. Let $\mathbf{Z} \in \mathbb{R}^{N_r \times L}$ be the collection of normalized $\{z_i\}_{i=1}^{N_r}$, $z_i \in \mathbb{R}^L$, the adjacency matrix for $G_{T \rightarrow T}$ with self loops can be calculated as $\mathcal{E} = \mathbf{Z}\mathbf{Z}^T$, so that $e_{i,j} = \frac{z_i z_j^T}{\|z_i z_j^T\|}$, where $\|\cdot\|$ is the L2-norm.

To determine the neighborhoods of each node, using a fully connected \mathcal{E} directly will establish relationship between backgrounds(negative) samples which will lead to redundant edges and greater computation cost. In this paper, we consider a sparse graph. For each region proposal i , we only retain the top t largest values of each row of \mathcal{E} , that is: $\text{Neighbour}(\text{Node } i) = \text{Top-}t_{j=1, \dots, N_r}(e_{i,j})$. This sparsity constraint ensures a sparse graph structure focusing on the most relevant relationship for recognition of the objects.

Feature Enhanced via Intra-Domain Graph Reasoning.

Most recent existing works (Gong et al.; Chen et al.; Jiang et al. 2018; 2018; 2018) propagate visual features locally among regions in the image and the information is only from those categories appearing in the image. This limits the performance of graph reasoning because of poor or distracted feature representations. Instead, we create a global semantic pool \mathbf{P}_T to store high-level semantic representations for all categories. In some works (Wang, Ye, and Gupta; Gong et al.; Gidaris and Komodakis 2018; 2018; 2018) in zero/few-shot problem, they try to train a model to fit the weights of the classifier of an unseen/unfamiliar category and the weights of the classifier for each category can be regarded as containing high-level semantic information. Formally, let $\mathbf{P}_T \in \mathbb{R}^{C_T \times (D+1)}$ denote the weights and the bias of the previous classifier for all the C_T categories in the domain T . The global semantic pool \mathbf{P}_T is extracted from the previous classification layer in the bbox head of the detection network and can be updated in each iteration during training.

Since our graph G is a region-to-region graph, we first need to map the category semantic embedding $p_T \in \mathbf{P}_T$ to the regional representations of nodes $x_i \in \mathbf{X}$. In this paper, we use a *soft-mapping* which computes the mapping weights

$m_{p_T \rightarrow x_i} \in \mathbf{M}_T$ as $m_{p_T \rightarrow x_i} = \frac{\exp(s_{ij})}{\sum_j \exp(s_{ij})}$, where s_{ij} is the classification score for the region i towards category j from the previous classification layer of the detector. Thus the regional representations of the nodes $\mathbf{X} \in \mathbb{R}^{N_r \times (D+1)}$ can be computed as a matrix multiplication: $\mathbf{X} = \mathbf{M}_T \mathbf{P}_T$, where $\mathbf{M}_T \in \mathbb{R}^{N_r \times C_T}$ is the soft-mapping matrix.

Given the node representation $x_j \in \mathbf{X}$ and the learned graph $\mathbf{G}_{T \rightarrow T}$, it is natural to use a GCN for modeling the relation and interaction. We thus define a patch operator for the GCN as follows:

$$f'_k(i) = \sum_{j \in \text{Neighbour}(i)} w_k(g_{ij}) x_j e_{ij}, \quad (1)$$

where $\text{Neighbour}(i)$ denotes the neighborhood of node i and e_{ij} is the normalized adjacency element of \mathcal{E} . To capture the pairwise spatial information between proposals, we further add K spatial weight terms $w_k(g_{ij})$ which are calculated by a nonlinear transformation function $w_k(\cdot)$ encoding the spatial information of regions. g_{ij} is a four dimensional relative geometry feature between proposal i and proposal j : $(\log \frac{|x_i - x_j|}{w_i}, \log \frac{|x_i - x_j|}{h_i}, \log \frac{w_i}{w_j}, \log \frac{h_i}{h_j})$, where w_i and h_i denotes the width and height of the region. We consider K set of $w_k(\cdot)$ to encode different kinds of spatial interactions. Flowchart of the graph reasoning can be found in Figure 3.

Then $f'_k(i)$ for each node goes through a linear transformation $\mathbf{L} \in \mathbb{R}^{E \times (D+1)}$ and is concatenated together: $\mathbf{f}' = [L(f'_k(i))]$, and KE is the dimension of the output enhanced feature for each region. Finally, the \mathbf{f}' for each region is concatenated to the original region features \mathbf{f} to improve both classification and localization. Note that the \mathbf{f}' is a distilled information across the categories with connected edges such as similar attributes or relations. Thus, sharing the common features between categories can help to improve the feature representation by adding and discovering adaptive contexts from the global semantic pool.

Inter-Domain Transfer Module

To effectively distill relevant information from datasets with different annotations, we introduce Inter-Domain Transfer Module to bridge the gap between different domains. This module can be applied for multiple source domains as long as they have semantic correlation. More specifically, for each pair of domains, we implement an inter-domain transfer module from source domain \mathcal{S} to target domain \mathcal{T} . Given an image, the RPN head proposes a set of region proposals, some of which may contain objects defined in other source domains that have some relationships with objects defined in target domain \mathcal{T} . We will first extract region proposals semantic features of a certain source domain \mathcal{S} according to the global semantic \mathbf{P}_S pool of \mathcal{S} . Then these semantic features are transferred to the target domain \mathcal{T} by the GCN with the predefined transfer graph $\mathbf{G}_{S \rightarrow T}$ between the domain \mathcal{S} and \mathcal{T} . Finally, the proposal features are concatenated by the information of the source domain for better prediction.

Construct $\mathbf{G}_{S \rightarrow T}$ For different detection task with diverse categories, the attribute similarities, relationships or hierarchical correlations among them can be exploited. For

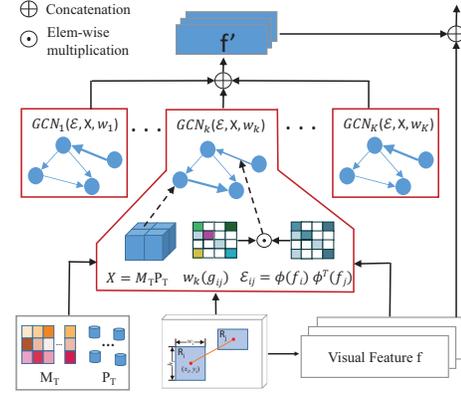


Figure 3: Flowchart of our graph reasoning and transfer in Equation (1). The global semantic pool \mathbf{P}_T is soft-mapped to the regional nodes by $\mathbf{X} = \mathbf{M}_T \mathbf{P}_T$. The graph edge \mathcal{E} and spatial weight terms $w_k(g_{ij})$ are calculated from \mathbf{f} and spatial information between proposal i and proposal j . Then graph conv $GCN_k(\cdot)$ is performed on the nodes \mathbf{X} according to the graph edge \mathcal{E} and w_k . Final enhanced feature \mathbf{f}' is a concatenation of all the GCN outputs. For inter-domain transfer module, the idea is similar except the input graph edge: $\mathcal{E}_{S \rightarrow T}$ and the nodes: $\mathbf{Y} = \mathbf{M}_S \mathbf{P}_S$ which are transferred from source semantic pool.

example, person label in a COCO dataset contains head, arm, and leg in VG domain, and the person label can also be composed of more fine-grained categories (e.g., man, woman, boy and girl) in ADE dataset. Also, there may further exist some location and co-occurrence relationship such as “road & car”, “street & truck” and “handbag & arm”. Thus, to transfer information from source domain to target domain, we define a transfer graph $\mathbf{G}_{S \rightarrow T}$, where $t_{c_i c_j} \in \mathbf{G}_{S \rightarrow T}$ and c_i, c_j denote the categories in source domain and target domain respectively.

In this paper we further try different methods to construct graph to explore various graph transfer dependencies between different label sets. Thus we consider and compare four schemes for $\mathbf{G}_{S \rightarrow T}$ including learning the graph from feature, handcrafted attribute/relationship, and word embedding similarity.

Handcrafted Attribute. Following Jiang et al., we consider a similar way to construct a handcrafted attribute $\mathbf{G}_{S \rightarrow T}$. Let us consider K attributes such as colors, size, materials, and status, we obtain a $(C_S + C_T) \times K$ frequency distribution table for each class-attribute pair. Then the pairwise Jensen–Shannon (JS) divergence between probability distributions P_{c_i} and P_{c_j} of two categories i and j from source/target domain can be measured as the edge weights of two classes: $t_{c_i c_j} = JS(P_{c_i} || P_{c_j})$, where JS divergence measures the similarity between two distributions.

Handcrafted Relationship. We may consider the pairwise relationship between classes, such as location relationship (e.g. *along, on*), the “subject-verb-object” relationship (e.g. *eat, wear*) or co-occurrence relationship. We can calculate frequent statistics f_{ij} from the occurrence among all

%	Method	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L	AR ₁	AR ₁₀	AR ₁₀₀	AR _S	AR _M	AR _L
VGMSCOCO	FPN (Lin et al. 2017)	38.6	60.4	41.8	22.3	43.2	49.8	31.8	50.5	53.2	33.9	58.0	67.1
	Multi Branches	39.9	61.5	43.4	23.5	44.9	51.4	32.7	52.4	55.2	36.3	60.3	68.3
	Universal-RCNN	44.7 ^{+6.1}	65.1 ^{+4.7}	49.1 ^{+7.3}	26.6 ^{+4.3}	49.1 ^{+5.9}	58.6 ^{+8.8}	35.6 ^{+3.8}	56.6 ^{+6.1}	59.5 ^{+6.3}	39.5 ^{+5.5}	64.3 ^{+6.3}	75.1 ^{+8.0}
VG	FPN (Lin et al. 2017)	7.0	12.3	7.0	3.9	7.6	10.2	14.0	19.1	19.3	10.5	19.4	22.5
	Multi Branches	7.1	12.3	7.2	4.1	7.6	10.3	13.8	19.0	19.2	10.9	19.0	22.2
	Universal-RCNN	8.8 ^{+1.8}	14.2 ^{+1.9}	9.3 ^{+2.3}	5.0 ^{+1.1}	9.4 ^{+1.8}	13.3 ^{+3.1}	17.5 ^{+3.5}	23.7 ^{+4.6}	24.0 ^{+4.7}	13.4 ^{+2.9}	23.3 ^{+3.9}	28.8 ^{+6.3}
ADE	FPN (Lin et al. 2017)	11.3	19.6	11.7	6.9	11.8	17.6	12.8	19.5	20.1	12.6	21.4	28.1
	Multi Branches	12.2	19.9	12.8	7.2	12.0	19.6	13.3	20.2	20.8	13.3	22.2	29.4
	Universal-RCNN	15.4 ^{+4.1}	24.2 ^{+4.6}	16.7 ^{+5.0}	9.4 ^{+2.5}	15.5 ^{+3.7}	24.3 ^{+6.7}	17.4 ^{+4.6}	26.1 ^{+6.6}	26.7 ^{+6.6}	17.3 ^{+4.7}	28.3 ^{+6.9}	37.1 ^{+9.0}

Table 1: Main results on MSCOCO(*minival*), VG, and ADE. “Universal-RCNN” is our full model trained on all the three domains via graph transfer and reasoning. The backbones of all the models are ResNet-101.

source-target categories pairs from additional linguistic information or through counting from images. The symmetric transformation and row normalization are performed on edge weights: $t_{c_i c_j} = \frac{f_{ij}}{\sqrt{d_{ii} d_{jj}}}$, where $d_{ii} = \sum_j f_{ij}$.

Word Embedding Similarity. We further explore the linguistic knowledge to construct the $G_{S \rightarrow T}$ besides the visual information. We use the word2vec model Almazán et al. to map the semantic word of categories to a word embedding vector. Then we compute the similarity between the names of the categories in source domain and target domain: $t_{c_i c_j} = \frac{\exp(\cos(w_i, w_j))}{\sum_j \exp(\cos(w_i, w_j))}$, where $\cos(w_i, w_j)$ is the cosine similarity between the word embedding vectors of the i th source category and j th target category.

Learning the Graph from Features. This scheme is almost the same as the one used in the Intra-Domain Reasoning Module. The visual feature is transformed to the latent space z . Edge $\mathcal{E}_{S \rightarrow T}$ between the source and target domain can be calculated by $\frac{z_i z_j^T}{\|z_i z_j^T\|}$.

Feature Enhanced via Inter-Domain Transfer After creating a global semantic pool \mathbf{P}_S for all the C_S categories, it is then natural to transfer high level semantic information from the source domain $\mathbf{P}_S \in \mathbb{R}^{C_S \times (D+1)}$ to the target domain by the transfer graph $G_{S \rightarrow T} \in \mathbb{R}^{C_S \times C_T}$. Since the nodes are proposal regions, a $N_r \times N_r$ adjacent matrix $\mathcal{E}_{S \rightarrow T}$ between regions can be obtained by a matrix multiplication: $\mathcal{E}_{S \rightarrow T} = \mathbf{M}_S G_{S \rightarrow T} \mathbf{M}_T^T$ where $\mathbf{M}_S \in \mathbb{R}^{N_r \times C_S}$ which is created similar to the soft-mapping by the classification score for the region towards the source categories in Section . Then we can also obtain regional representations of nodes $y_i \in \mathbf{Y}$ from the relevant information from source domain by $\mathbf{Y} = \mathbf{M}_S \mathbf{P}_S$.

Finally, given the node representation $y_i \in \mathbf{Y}$ and the adjacent matrix $\mathcal{E}_{S \rightarrow T} \in \mathbb{R}^{N_r \times N_r}$, we can also use weighted GCN as to propagate the semantic representations from the source domain to the target region nodes. Finally, the output of the GCN \mathbf{f}'_s is concatenated to the original visual features \mathbf{f} for better classification and bounding box regression. Note that \mathbf{f}'_s serves as supplementary information since the relevant prediction in the source domain is transferred to each region to help better recognize the item.

Eval	Methods	Train with	mAP	Eval	Methods	Train with	mAP
MSCOCO	FPN	-	38.6	VG	FPN	-	7.0
	Multi-Branches	VG	39.8		Multi-Branches	COCO	7.0
	Fine-tuning	VG	39.2		Fine-tuning	COCO	7.4
	Overlap Labels	VG	38.7		Universal-RCNN	COCO	8.2
	Pseudo Labels	VG	38.7		FPN	-	7.0
	Universal-RCNN	VG	43.5		Multi-Branches	ADE	7.0
	Multi-Branches	ADE	38.8		Fine-tuning	ADE	7.3
	Fine-tuning	ADE	38.6		Universal-RCNN	ADE	8.0
ADE	Universal-RCNN	ADE	41.5	ADE	FPN	-	11.3
	Multi-Branches	VG	12.3		Multi-Branches	COCO	11.4
	Fine-tuning	VG	12.8		Fine-tuning	COCO	11.9
	Universal-RCNN	VG	14.6		Universal-RCNN	COCO	12.9

Table 2: Results of mAP with models trained based on two domains. “Universal-RCNN” is our full model trained on two domains via graph transfer and reasoning. The backbones are ResNet-101.

Experiments

Intra-Domain Reasoning	Inter-Domain Transfer				AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
	Attribute	Relation	Embed	Learn						
✓					38.6	60.4	41.8	22.3	43.2	49.8
✓	✓				41.4	60.9	46.0	22.9	46.1	55.2
✓		✓			43.0	62.8	47.3	25.0	47.6	56.3
✓			✓		43.2	63.0	47.3	25.4	47.5	56.8
✓				✓	42.7	62.7	47.4	24.9	47.4	56.3
✓				✓	41.9	60.9	46.1	23.5	46.2	55.8
✓				✓	43.5	63.5	47.7	25.8	47.8	57.0

Table 3: Ablation study on MSCOCO(*minival*) with models trained based on VG and MSCOCO domains. We consider four schemes to construct the transfer dependencies including handcrafted attribute(“Attribute”), relationship(“Relation”), word embedding similarity(“Embed”) and learning the graph from features(“Learn”). The backbones of all the models are ResNet-101.

Datasets and Evaluations. We evaluate the performance of our Universal-RCNN on three object detection domains with different annotations of categories: MSCOCO (Lin et al. 2014), Visual Genome(VG) (Krishna et al. 2016), and ADE (Zhou et al. 2017). MSCOCO is a common object detection dataset with 80 object classes, which contains 118K training images, 5K validation images (denoted as *minival*) and 20K unannotated testing images (denoted as *test-dev*) as common practice. VG and ADE are two large-scale object detection benchmarks with thousands of object classes. For VG, we use the synsets (Russakovsky et al.

2015) instead of the raw names of the categories due to inconsistent label annotations. We use 88K images for training and 5K images for testing with 1000 most frequent classes, following Chen et al.; Jiang et al.. For ADE, we consider 445 classes and use 20K images for training and 1K images for testing, following Chen et al.; Jiang et al.. Since ADE is a segmentation dataset, we convert segmentation masks to bounding boxes for all instances. For all the evaluation, we use the standard COCO metrics including mean Average Precision (mAP) and Average Recall (AR).

Implementation Details. The proposed Universal-RCNN is a single network and trained in an end-to-end style. All tasks for multiple domains share a backbone for intermediate image feature extraction while having separated heads and transfer modules for multi-domain object feature learning. We use the popular FPN as our baseline detector and implement Universal-RCNN based on it. ResNet-101 (He et al. 2016) pretrained on ImageNet is used as the shared backbone network. The hyper-parameters in training mostly follow Lin et al.. During both training and testing, we resize the input image such that the shorter side has 800 pixels. RPN and bbox-head are applied to all levels in the feature pyramid. The total number of proposed regions after NMS is $N_r = 512$. In the bbox-head, 2 shared FC layer is used for proposal visual feature extraction and the output is a 1024-d vector feed into the bbox regression and class-agnostic classification layer following FPN.

For the Intra-Domain graph reasoning and Inter-Domain transfer module, we use the previous 2 shared FC layer for re-extracting visual features \mathbf{f} of region proposals. In the graph learner module, we use a linear transformation layer of size 256 to learn the latent representation \mathbf{Z} ($L = 256$) and most $t = 32$ relevant nodes are retained. The global semantic pools are created for each domain by copying the weights of the classification layers of different domains. For the spatial-aware GCN, we use two weighted graph convolutional layers with dimensions of 256 and 128 respectively so that the output size of the module for each region is 128. Each GCN consists of $K = 8$ spatial weight terms forming a multi-head graph attention layer (Monti et al. 2017). All experiments are conducted on a single server with 8 Tesla V100 GPUs by using the Pytorch framework. For training, SGD with weight decay of 0.0001 and momentum of 0.9 is adopted to optimize all models. The batch size is set to be 16 with 2 images on each GPU. The initial learning rate is 0.02, reduce twice (x0.1) during the training process. We train 12 epochs for all models in an end-to-end manner. We follow Jiang et al. to prepare handcrafted Attribute/Relationship matrix with the help of the annotations in the VG dataset.

Comparison with state-of-the-art domain transfer baselines. To show the effectiveness of our method in transferring relevant information across multiple domains, we compare the overall performance of Universal-RCNN with several state-of-the-art domain transfer baselines including a) training an unified model on all datasets with separate RPN and bbox heads for each domain stacked on a shared backbone (**Multiple Branches**); b) pretraining on source dataset, fine-tuning on target dataset (**Fine-tuning**); c) training based on the overlap labels (**Overlap Labels**); d) us-

# training	Method	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
50%	FPN w Finetune	36.3	57.3	39.3	20.6	40.3	47.4
	Multi Branches	37.6	59.0	40.8	21.8	42.6	48.4
	Universal-RCNN	42.0	62.1	46.3	24.4	46.3	56.4
80%	FPN w Finetune	38.3	59.4	41.8	21.8	42.6	50.0
	Multi Branches	39.2	60.5	42.7	23.5	44.0	50.7
	Universal-RCNN	42.9	63.1	46.8	25.1	47.0	56.5
100%	FPN w Finetune	39.2	60.5	42.5	22.8	43.8	50.7
	Multi Branches	39.8	61.2	43.2	23.1	44.3	51.1
	Universal-RCNN	43.5	63.5	47.7	25.8	47.8	57.0

Table 4: Training with less data on MSCOCO(*minival*) with models trained based on VG and MSCOCO domains. We only use a portion of MSCOCO data to train the models. “FPN w Finetune” is the FPN model first trained on VG, then finetuned on MSCOCO. “FPN w Multi Branches” is the method FPN with two branches.

ing a fully trained model from target dataset and generate pseudo-labels on source images and train again (**Pseudo-labels**). The comparison results are reported on MSCOCO, VG and ADE in Table 1 and Table 2.

Table 1 shows the results of our Universal R-CNN that exploits the semantic correlations between all three datasets. We only compare our Universal R-CNN with Multiple Branches method under this setting since b), c), d) can only be used on two datasets. As can be seen, the Universal-RCNN significantly outperforms Multiple Branches method in terms of precision (1.7% to 4.8%) and recall (4.7% to 6.6%), which indicates the superiority of explicitly extracting the semantic correlations across domains within our Universal-RCNN than implicitly learning feature-level interactions between domains within multi-task learning.

Table 2 shows the experiment results of training on each pair of domains. It can be found that our method consistently surpasses the above-mentioned domain transfer methods by a large margin through graph reasoning and transfer. Both Multiple Branches method and Fine-tuning method cannot improve detection performance much, which demonstrates the implicit way to utilize semantic information across multiple domains is ineffective. Since VG domain nearly contains all class annotations of COCO domain (79 classes out of total 80 classes), we test Overlap Labels method and Pseudo Labels method on COCO, the results show that both of them only obtain negligible 0.1% AP improvement. The reason is that the classes of VG domain have many fine-grained objects while the classes of COCO domain only have some specific objects. Table 2 also shows that transferring information from VG to COCO or ADE can achieve better detection performance, this is because that VG annotations contains a huge number of classes and many of them overlap with that of COCO and ADE. Overall, our Universal-RCNN can sufficiently mining the semantic correlations across all three domains and significantly boost the detection performance on each target domain.

Ablative Analysis. Table 3 shows the ablative analysis of our method trained based on VG and MSCOCO domains. Our Intra-domain reasoning is effective and acquires 2.8% improvements compared with the basic network on

%	Method	AP	AP ₅₀	AP ₇₅
MSCOCO	SIN (Liu et al. 2018b)	23.2	44.5	22.0
	Relation Network (Hu et al. 2018b)	38.9	60.5	43.3
	HKRM (Jiang et al. 2018)	37.8	58.0	41.3
	Ours w Intra-Domain	41.4	60.9	46.0
VG	HKRM (Jiang et al. 2018)	7.8	13.4	8.1
	Ours w Intra-Domain	7.9	13.7	8.3
ADE	HKRM (Jiang et al. 2018)	10.3	18.0	10.4
	Ours w Intra-Domain	14.0	23.1	14.9

Table 5: Comparison between intra-domain semantic correlation methods. The backbone are ResNet-101.

MSCOCO. For the Inter-domain transfer module, it can be found that adding the transfer module alone can boost the performance by 3.3% which demonstrates the importance of graph transfer across domains. Combining these two modules can further improve the performance by 1.6%, which demonstrates the effectiveness of the proposed Universal-RCNN in fully exploiting semantic correlations across multiple domains. We further compare four schemes to construct the transfer dependencies including handcrafted attribute, relationship, word embedding similarity and learning the graph from feature. Learning the graph from features performs better than other methods. Thus, we choose learning the graph from feature for the final model.

Training with Less Data. Since bounding box annotation is expensive, we are curious about how our Universal-RCNN performs with less training data. Table 4 shows the performance on MSCOCO with models trained based on VG and MSCOCO(less data). “FPN w Finetune” is the FPN model first trained on VG, then finetuned on MSCOCO with less data. The Universal-RCNN is better than baseline methods with a large margin in all scenarios. Moreover, it can be found that Universal-RCNN trained on half of the data (42% mAP) outperforms FPN trained on full data (39.2% mAP). This superior performance confirms the effectiveness of our method that seamlessly bridges the gap between domains and fully utilizes data annotations. Thus, our method can be quickly adapted to new dataset with less annotation by borrowing information from other domains.

Comparison with semantic correlation works within one domain. There have been many works that use graph based models for additional reasoning over bounding boxes, we compare our intra-domain graph reasoning method with those intra-domain semantic correlation works in this section. We train FPN with our intra-domain module using one dataset and report the results in Table 5. It can be found that our Intra-Domain graph reasoning module is superior to multiple competing methods. It should be noted that previous works (Hu et al.; Jiang et al.; Liu et al. 2018a; 2018; 2018b) use fully connected graphs to build object-object relationships, our method instead learn a sparse spatial-aware graph structure to perform graph inference, which can reduce lots of redundant relationships and improve the detection performance significantly. Furthermore, compare to the implicit knowledge module in Jiang et al., our method uses

%	Method	backbone	<i>test-val</i>					
			AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
MSCOCO	DetNet	DetNet-59	35.1	57.3	37.2	22.3	40.8	39.9
	DetNet w MB	DetNet-59	36.8	59.0	39.3	22.6	42.3	42.3
	DetNet w ours	DetNet-59	40.1	61.8	43.2	24.7	46.5	46.5
	HKRM	ResNet-101	40.1	61.2	44.8	23.8	43.7	51.2
	HKRM w MB	ResNet-101	41.6	62.9	45.5	24.3	45.2	52.8
	HKRM w ours	ResNet-101	43.7	63.5	47.3	24.8	46.5	55.4
	Mask-RCNN	ResNet-101	39.7	61.6	43.3	23.1	43.3	49.7
	Mask-RCNN w MB	ResNet-101	41.0	62.2	45.0	24.2	44.9	51.0
	Mask-RCNN w ours	ResNet-101	45.1	65.2	49.3	26.1	48.6	57.3
	FPN	ResNet-101	38.9	61.0	42.3	22.3	42.3	48.5
	FPN w MB	ResNet-101	40.1	61.7	43.6	23.2	43.8	49.7
	FPN w ours	ResNet-101	44.4	64.9	48.5	25.3	47.8	56.7
FPN w ours *	ResNext-101	49.1	69.3	54.2	30.4	52.7	61.7	

Table 6: Methods trained on multiple datasets(COCO, VG, ADE) and eval on COCO. “MB” is model with Multi Branches. * is the model adding multi-scale training and testing with Soft-NMS

global semantic pool to construct bbox features and use a GCN to aggregate information.

Generalization Capacity. To validate the generalization capability of the Universal-RCNN, we further implement our method upon more recent detection methods such as DetNet (Li et al. 2018), HKRM (Jiang et al. 2018) and Mask-RCNN (He et al. 2017) and compare. The results are reported on COCO in Table 6. Table 6 shows that Universal-RCNN consistently boost the performance by 4~6 points in terms of AP, which suggests that the Universal-RCNN is widely applicable across different detection baselines. We further add some bells and whistles to test the upper limit of the Universal-RCNN (Li et al. 2018). Specifically, we utilize ResNext-101 as the backbone and apply multi-scale training, multi-scale testing and Soft-NMS. Finally, the Universal-RCNN obtains 49.1 mAP on COCO *test-dev* with single-model result.

Conclusion

In this work, we proposed a new universal object detector (Universal-RCNN) to alleviate the categories discrepancy and fully utilize the data annotation. Our method can be easily plugged into any existing detection pipeline via Transferable Graph R-CNN with multiple domains for endowing its ability to global reasoning and transferring. Extensive experiments demonstrate the effectiveness of the proposed method and achieve the state-of-the-art results on multiple object detection benchmarks.

References

- Almazán, J.; Gordo, A.; Fornés, A.; and Valveny, E. 2014. Word spotting and recognition with embedded attributes. *IEEE transactions on pattern analysis and machine intelligence* 36(12):2552–2566.
- Bronstein, M. M.; Bruna, J.; LeCun, Y.; Szlam, A.; and Vandergheynst, P. 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* 34(4):18–42.
- Chen, X., and Gupta, A. 2017. Spatial memory for context reasoning in object detection. In *ICCV*.

- Chen, X.; Li, L.-J.; Fei-Fei, L.; and Gupta, A. 2018. Iterative visual reasoning beyond convolutions. In *CVPR*.
- Cui, Y.; Song, Y.; Sun, C.; Howard, A.; and Belongie, S. 2018. Large scale fine-grained categorization and domain-specific transfer learning. In *CVPR*.
- Dai, J.; He, K.; and Sun, J. 2016. Instance-aware semantic segmentation via multi-task network cascades. In *CVPR*.
- Everingham, M.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* 88(2):303–338.
- Gidaris, S., and Komodakis, N. 2018. Dynamic few-shot visual learning without forgetting. In *CVPR*.
- Gong, K.; Liang, X.; Zhang, D.; Shen, X.; and Lin, L. 2017. Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing. In *CVPR*.
- Gong, C.; He, D.; Tan, X.; Qin, T.; Wang, L.; and Liu, T.-Y. 2018. Frage: Frequency-agnostic word representation. In *NIPS*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*.
- Hoffman, J.; Guadarrama, S.; Tzeng, E. S.; Hu, R.; Donahue, J.; Girshick, R.; Darrell, T.; and Saenko, K. 2014. Lsd: Large scale detection through adaptation. In *NIPS*.
- Hu, H.; Gu, J.; Zhang, Z.; Dai, J.; and Wei, Y. 2018a. Relation networks for object detection. In *CVPR*.
- Hu, R.; Dollár, P.; He, K.; Darrell, T.; and Girshick, R. 2018b. Learning to segment every thing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4233–4241.
- Jiang, C.; Xu, H.; Liang, X.; and Lin, L. 2018. Hybrid knowledge routed modules for large-scale object detection. In *NIPS*.
- Kendall, A.; Gal, Y.; and Cipolla, R. 2018. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*.
- Kirillov, A.; He, K.; Girshick, R.; Rother, C.; and Dollár, P. 2018. Panoptic segmentation. *arXiv preprint arXiv:1801.00868*.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; Bernstein, M.; and Fei-Fei, L. 2016. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*.
- Li, X.; Zhao, L.; Wei, L.; Yang, M.-H.; Wu, F.; Zhuang, Y.; Ling, H.; and Wang, J. 2016. Deepsaliency: Multi-task deep neural network model for salient object detection. *IEEE Transactions on Image Processing* 25(8):3919–3930.
- Li, Z.; Peng, C.; Yu, G.; Zhang, X.; Deng, Y.; and Sun, J. 2018. Detnet: A backbone network for object detection. In *ECCV*.
- Liang, X.; Xu, C.; Shen, X.; Yang, J.; Liu, S.; Tang, J.; Lin, L.; and Yan, S. 2015. Human parsing with contextualized convolutional neural network. In *CVPR*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *CVPR*.
- Liu, S.; Qi, L.; Qin, H.; Shi, J.; and Jia, J. 2018a. Path aggregation network for instance segmentation. In *CVPR*.
- Liu, Y.; Wang, R.; Shan, S.; and Chen, X. 2018b. Structure inference net: Object detection using scene-level context and instance-level relationships. In *CVPR*.
- Luo, Z.; Mishra, A.; Achkar, A.; Eichel, J.; Li, S.; and Jodoin, P.-M. 2017. Non-local deep features for salient object detection. In *CVPR*.
- Misra, I.; Gupta, A.; and Hebert, M. 2017. From red wine to red tomato: Composition with context. In *CVPR*.
- Monti, F.; Boscaini, D.; Masci, J.; Rodola, E.; Svoboda, J.; and Bronstein, M. M. 2017. Geometric deep learning on graphs and manifolds using mixture model cnns. In *CVPR*.
- Peng, P.; Tian, Y.; Xiang, T.; Wang, Y.; Pontil, M.; and Huang, T. 2018. Joint semantic and latent attribute modelling for cross-class transfer learning. *IEEE transactions on pattern analysis and machine intelligence* 40(7):1625–1638.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115(3):211–252.
- Tsai, Y.-H. H.; Huang, L.-K.; and Salakhutdinov, R. 2017. Learning robust visual-semantic embeddings. In *ICCV*.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. In *ICLR*.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *CVPR*.
- Wang, X.; Ye, Y.; and Gupta, A. 2018. Zero-shot recognition via semantic embeddings and knowledge graphs. In *CVPR*.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified perceptual parsing for scene understanding. In *ECCV*.
- Xiong, Y.; Liao, R.; Zhao, H.; Hu, R.; Bai, M.; Yumer, E.; and Urtasun, R. 2019. Upsnet: A unified panoptic segmentation network. *arXiv preprint arXiv:1901.03784*.
- Xu, H.; Jiang, C.; Liang, X.; and Li, Z. 2019a. Spatial-aware graph relation network for large-scale object detection. In *CVPR*.
- Xu, H.; Jiang, C.; Liang, X.; Lin, L.; and Li, Z. 2019b. Reasoning-rnn: Unifying adaptive global reasoning into large-scale object detection. In *CVPR*.
- Yu, F.; Xian, W.; Chen, Y.; Liu, F.; Liao, M.; Madhavan, V.; and Darrell, T. 2018. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*.
- Zamir, A. R.; Sax, A.; Shen, W.; Guibas, L. J.; Malik, J.; and Savarese, S. 2018. Taskonomy: Disentangling task transfer learning. In *CVPR*.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *CVPR*.
- Zhu, C.; He, Y.; and Savvides, M. 2019. Feature selective anchor-free module for single-shot object detection. *arXiv preprint arXiv:1903.00621*.