

Deep Object Co-Segmentation via Spatial-Semantic Network Modulation

Kaihua Zhang,¹ Jin Chen,¹ Bo Liu,^{2*} Qingshan Liu¹

¹B-DAT and CICAET, Nanjing University of Information Science and Technology, Nanjing, China

²JD Finance America Corporation
{zhkhua, kfliubo}@gmail.com

Abstract

Object co-segmentation is to segment the shared objects in multiple relevant images, which has numerous applications in computer vision. This paper presents a spatial and semantic modulated deep network framework for object co-segmentation. A backbone network is adopted to extract multi-resolution image features. With the multi-resolution features of the relevant images as input, we design a spatial modulator to learn a mask for each image. The spatial modulator captures the correlations of image feature descriptors via unsupervised learning. The learned mask can roughly localize the shared foreground object while suppressing the background. For the semantic modulator, we model it as a supervised image classification task. We propose a hierarchical second-order pooling module to transform the image features for classification use. The outputs of the two modulators manipulate the multi-resolution features by a shift-and-scale operation so that the features focus on segmenting co-object regions. The proposed model is trained end-to-end without any intricate post-processing. Extensive experiments on four image co-segmentation benchmark datasets demonstrate the superior accuracy of the proposed method compared to state-of-the-art methods. The codes are available at <http://kaihuazhang.net/>.

1 Introduction

As a special case of image object segmentation, object co-segmentation refers to the task of jointly discovering and segmenting the objects shared in a group of images. It has been widely used to support various computer vision applications, such as interactive image segmentation (Kamranian et al. 2018), 3D reconstruction (Mustafa and Hilton 2017) and object co-localization (Wei et al. 2019; Han et al. 2018), to name a few.

Image features that characterize the co-objects in the image group are vital for a co-segmentation task. Conventional approaches use the hand-crafted cues such as color histograms, Gabor filter outputs and SIFT descriptors as feature representations (Yuan, Lu, and Shivakumara 2014; Dai et al. 2013; Lee et al. 2015). Those hand-crafted



Figure 1: Object co-segmentation examples by our approach. (a) Horse group; (b) Horse group co-segmentation maps.

features cannot well handle the challenging cases in co-segmentation such as background clutter and large-scale appearance variations of the co-objects in images. In recent years, deep-learning-based co-segmentation methods have attracted much attention. For example, (Li et al. 2018; Chen et al. 2018) leverage a Siamese network architecture for object co-segmentation and an attention mechanism is used to enhance the co-object feature representations. These methods have shown superior performance compared to the traditional methods (Yuan, Lu, and Shivakumara 2014; Dai et al. 2013; Lee et al. 2015), which inspire us to explore a deep-learning-based solution to object co-segmentation.

One critical property of object co-segmentation is that the co-objects in images belong to the same semantic category. Those co-objects usually occupy part of each image. One illustrative example is shown in Figure 1. It is desirable that the deep convolutional network layers, being used as a feature extractor, are targeted on modelling the co-objects. To this end, we propose a spatial-semantic modulated network structure to model this property. The two modulators are achieved by the designed group-wise mask learning branch and co-category classification branch, respectively. We summarize the technical contributions of this work as follows:

- We propose a spatial-semantic modulated deep network for object co-segmentation. Image features extracted by a backbone network are used to learn a spatial modulator and a semantic modulator. The outputs of the modulators

*Corresponding author

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

guide the image features up-sampling to generate the co-segmentation results. The network parameter learning is formulated into a multi-task learning task, and the whole network is trained in an end-to-end manner.

- For the spatial modulation branch, an unsupervised learning method is proposed to learn a mask for each image. With the fused multi-resolution image features as input, we formulate the mask learning as an integer programming problem. Its continuous relaxation has a closed-form solution. The learned parameter indicates whether the corresponding image pixel corresponds to foreground or background.
- In the semantic modulation branch, we design a hierarchical second-order pooling (HSP) operator to transform the convolutional features for object classification. Spatial pooling (SP) is shown to be able to capture the high-order feature statistical dependency (Gao et al. 2019). The proposed HSP module has a stack of two SP layers. They are dedicated to capturing the long-range channel-wise dependency of the holistic feature representation. The output of the HSP layer is fed into a fully-connected layer for object classification and used as the semantic modulator.

We conduct extensive evaluations on four object co-segmentation benchmark datasets (Faktor and Irani 2013; Rubinstein et al. 2013), including the sub-set of MSRC, Internet, the sub-set of iCoseg and PASCAL-VOC datasets. The proposed model achieves a significantly higher accuracy than state-of-the-art methods. Especially, on the most challenging PASCAL-VOC dataset, our method outperforms the second best-performing state-of-the-art approach (Li et al. 2018) by 6% in terms of average Jaccard index \mathcal{J} .

The rest of this work is organized as follows: In § 2, we introduce the related works of our study. § 3 describes the proposed framework and its main components. Afterwards, comprehensive experimental evaluations are presented in § 4. Finally, we conclude this work in § 5.

2 Related Work

2.1 Object Co-segmentation

A more comprehensive literature review about image co-segmentation can be found in (Zhu et al. 2016). Existing object co-segmentation methods can be roughly grouped into four categories including graph-based model, saliency-based model, joint processing model and deep learning model. Conventional approaches such as (Yuan, Lu, and Shivakumara 2014; Collins et al. 2012; Lee et al. 2015) assume the pixels or superpixels in the co-objects can be grouped together and then they formulate co-segmentation as a clustering task to search for the co-objects. Saliency-detection-based methods assume regions of interest in the images are usually the co-objects to be segmented. They conduct image co-segmentation through detecting the regions that attract human attention most. Representative models include (Tsai et al. 2018; Zhang et al. 2019; Lu, Xu, and Liu 2019). The work in (Dai et al. 2013; Jerripothula et al. 2017) employs a coupled framework for co-skeletonization and co-segmentation tasks so that they are well informed

by each other, and benefit each other synergistically. The idea of joint processing can exploit the inherent interdependencies of two tasks to achieve better results jointly. Recently, (Li et al. 2018; Chen et al. 2018) respectively propose an end-to-end deep-learning-based method for object co-segmentation using a Siamese encoder-decoder architecture and a semantic-aware attention mechanism.

2.2 Network Modulation

Modulation module has been proved to be an effective way to manipulate network parameter learning. The modulator can be modelled as parameters or output of the auxiliary branch that are used to guide the main branch parameter learning. In the segmentation method (Dai, He, and Sun 2015), an image mask is used as a modulator for background removal. In the Mask R-CNN model (He et al. 2017), a classification branch is used to guide the segmentation branch learning. Feature-wise linear modulation is a widely-used scheme, which has been applied to object detection (Lin et al. 2017) and graph neural networks learning (Brockschmidt 2019). In visual reasoning problem, network modulation is used to encode the language information (De Vries et al. 2017; Perez et al. 2018). The attention module in the image caption model (Chen et al. 2017) can be viewed as a modulator. (Yang et al. 2018) proposes to model the visual and spatial information by a modulator for video object segmentation. In (Flores et al. 2019), a saliency detection branch is added to an existing CNN architecture as a modulator for fine-grained object recognition. A cross-modulation mechanism is proposed in (Prol, Dumoulin, and Herranz 2018) for few-shot learning.

3 Proposed Approach

3.1 Problem Formulation

Figure 2 presents an overview of our model. Given a group of N images $\mathcal{I} = \{I^n\}_{n=1}^N$ containing co-objects of a specific category, our objective is to learn a feed-forward network f that produces a set of object co-segmentation masks $\mathcal{M} = \{M^n\}_{n=1}^N$:

$$\mathcal{M} = f(\mathcal{I}; \theta), \quad (1)$$

where θ denotes the network parameters to be optimized. The network f is composed of three sub-networks: spatial modulation sub-net f_{spa} , semantic modulation sub-net f_{sem} and segmentation sub-net f_{seg} . The renowned SPP-Net (He et al. 2015) has shown that the convolutional feature maps (CFMs) for object recognition encode both spatial layouts of objects (by their positions) and the semantics (by strengths of their activations). Inspired by this model, we design f_{spa} and f_{sem} to encode the spatial and semantic information of the co-objects in \mathcal{I} , respectively. The two modulators guide the convolution layers learning in f_{seg} to focus on the co-objects in the images. Specifically, the sub-net f_{spa} is to learn a mask for each image to coarsely localize the co-object in it. Given the input CFMs $\{\varphi(I^n)\}_{n=1}^N$ produced by fusing all the output CFMs of our backbone network, the sub-net f_{spa} produces a set of spatial masks $\mathcal{S} = \{S^n \in \mathbb{R}^{w \times h}\}_{n=1}^N$ with width w and height h :

$$\mathcal{S} = f_{spa}(\{\varphi(I^n)\}_{n=1}^N; \theta_{spa}), \quad (2)$$

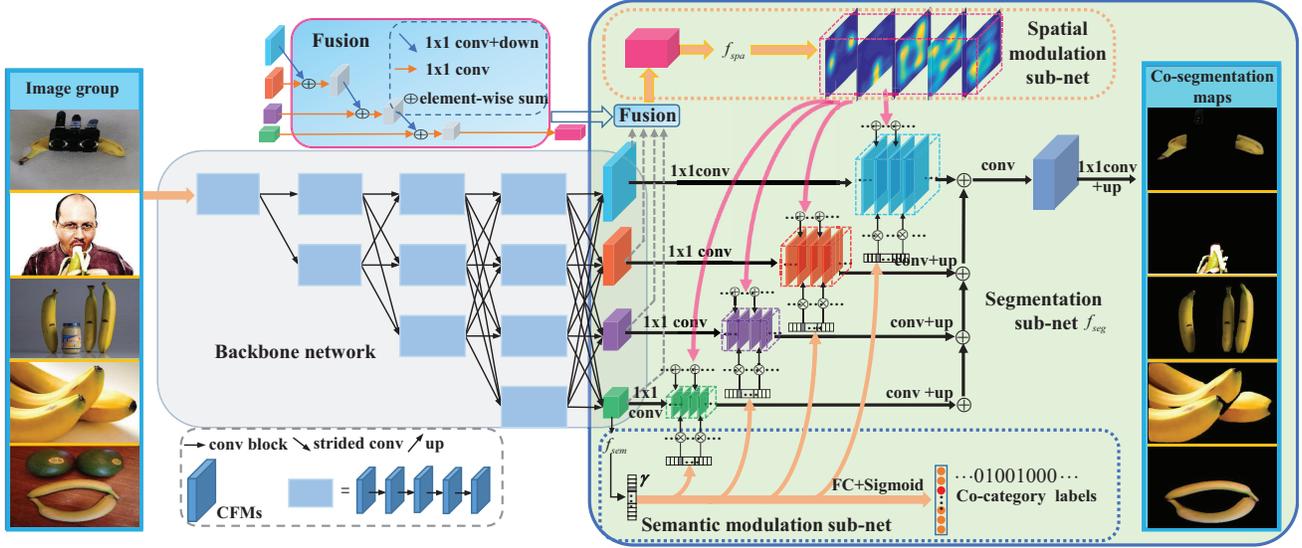


Figure 2: Overview of the proposed object co-segmentation framework. Firstly, a group of images are fed into the backbone network to yield a set of multi-resolution CFMs. Then, the CFMs are modulated by a group of spatial heatmaps and a feature channel selector vector. The former is generated by a clustering approach that can well capture the coarse localizations of the co-objects in the images. Under the supervision of co-category labels, the latter is obtained by learning a group-wise semantic representation that indicates the importance of the feature channels. Finally, the multi-resolution modulated CFMs are fused in a way similar to the feature pyramid network (FPN) (Lin et al. 2017) to produce the co-segmentation maps. ‘conv’, ‘FC’, ‘up’ and ‘down’ are short for convolutional, fully-connected, upsampling and downsampling layers, respectively.

where θ_{spa} denotes the corresponding network parameters to be optimized. Although the coarse spatial layout information of the co-objects in all images can be embedded into \mathcal{S} in (2), the useful high-level semantic information that are essential to differentiate co-objects from distractors fails to be transferred into \mathcal{S} . To address this issue, we further propose f_{sem} as a complement. The sub-net f_{sem} learns a channel selector vector $\gamma \in \mathbb{R}^d$ with d channels. The entries of γ indicate the importance of feature channels, that is

$$\gamma = f_{sem}(\{\phi(I^n)\}_{n=1}^N; \theta_{sem}), \quad (3)$$

where ϕ denotes the output CFMs with the lowest resolution generated by our backbone network, and θ_{sem} is the corresponding sub-net parameters to be learned. γ is optimized using the co-category labels as supervision. Finally, we use the spatial and the semantic modulators as guidance to segment the co-object regions in each image I^n :

$$M^n = f_{seg}(I^n, S^n, \gamma; \theta_{seg}), \quad (4)$$

where θ_{seg} is the parameters of the segmentation sub-net. To be specific, we transfer the spatial and semantic guidance $\{\mathcal{S}, \gamma\}$ into f_{seg} using a simple shift-and-scale operation on the input CFMs of f_{seg} : for each image $I^n \in \mathcal{I}$, its modulated feature maps are formulated as

$$Y_c^n = \gamma_c X_c^n + S^n, c = 1, \dots, d, \quad (5)$$

where $X_c^n, Y_c^n \in \mathbb{R}^{w \times h}$ are the input and output CFMs in the c_{th} channel, γ_c is the c_{th} element of γ .

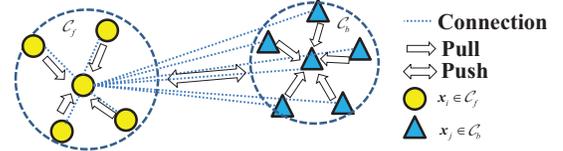


Figure 3: Schematic illustration of the proposed clustering objective for spatial modulator. The objective considers the mutual effects between any two samples, which pulls the samples of the same cluster together while pushing away the samples of different clusters.

3.2 Spatial Modulator

In the sub-net f_{spa} (2), the i -th channel feature $x_i^n \in \mathbb{R}^d$ of the input $\varphi(I^n) \in \mathbb{R}^{w \times h \times d}$ represents a corresponding local region in I^n . For expression clarity, we represent all the channel feature representations of \mathcal{I} as $\mathcal{X} = \{x_i \in \mathbb{R}^d\}_{i=1}^{whN}$. The sub-net f_{spa} aims at partitioning the data points in \mathcal{X} into two classes C_f, C_b of foreground and background. However, if training f_{spa} using a supervised learning method with a fixed set of categories, it cannot generalize well to unseen categories. To this end, we propose a simple yet effective clustering approach to partitioning \mathcal{X} into two clusters C_f, C_b without knowing object semantic categories. Our unsupervised method can highlight category-agnostic co-object regions in images and hence can better generalize to unseen categories. As shown by Figure 3, this can be achieved by maximizing all the distances between the fore-

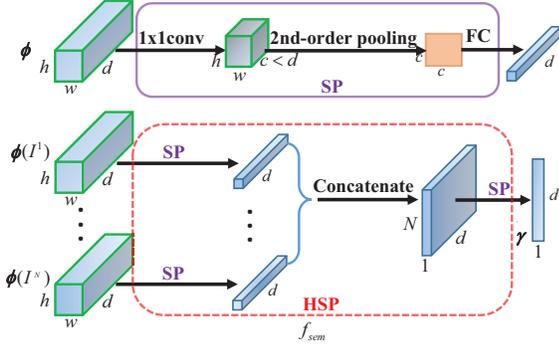


Figure 4: Illustration of the SP and the HSP. The sub-net f_{sem} (3) is composed of the HSP module.

ground and the background samples while minimizing all the distances between the foreground samples and between the background ones respectively. To this end, we define the clustering objective as follows:

$$\min\{\ell_{spa} = -2 \sum_{i \in \mathcal{C}_f, j \in \mathcal{C}_b} d_{ij} + \sum_{i, j \in \mathcal{C}_f} d_{ij} + \sum_{i, j \in \mathcal{C}_b} d_{ij}\}, \quad (6)$$

where $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|_2^2$ is the squared Euclidean distance between samples i and j . Since we use normalized channel features satisfying $\|\mathbf{x}_i\|_2^2 = 1$, d_{ij} can be reformulated as

$$d_{ij} = 2 - 2\mathbf{x}_i^\top \mathbf{x}_j. \quad (7)$$

Using a cluster indicator vector $\mathbf{s} = [s_1, \dots, s_{whN}]^\top$ subject to $\|\mathbf{s}\|_2^2 = 1$, where $s_i = 1/\sqrt{whN}$ if $i \in \mathcal{C}_f$ and $s_i = -1/\sqrt{whN}$ if $i \in \mathcal{C}_b$, the loss function ℓ_{spa} in (6) can be reformulated as

$$\ell_{spa}(\mathbf{s}) = whN\mathbf{s}^\top \mathbf{D}\mathbf{s}, \quad (8)$$

where the (i, j) -th entry of $\mathbf{D} = d_{ij}$. Putting (7) into (8) and removing the trivial constant whN , ℓ_{spa} can be reformulated as

$$\ell_{spa}(\mathbf{s}) = -\mathbf{s}^\top \mathbf{G}\mathbf{s}, \quad (9)$$

where $\mathbf{G} = \mathbf{X}^\top \mathbf{X} - \mathbf{I}$ with $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_{whN}]$, \mathbf{I} denotes an all-ones matrix. Relaxing the elements in \mathbf{s} from binary indicator values to continuous values in $[-1, 1]$ subject to $\|\mathbf{s}\|_2^2 = 1$, the solution $\hat{\mathbf{s}} = \arg \min_{\mathbf{s}} \ell_{spa}(\mathbf{s})$ satisfies (Ding and He 2004)

$$\mathbf{G}\hat{\mathbf{s}} = \lambda_{max}\hat{\mathbf{s}}, \quad (10)$$

where λ_{max} denotes the maximum eigenvalue of \mathbf{G} , and its corresponding eigenvector is $\hat{\mathbf{s}} \in \mathbb{R}^{whN}$. The optimal solution $\hat{\mathbf{s}}$ is then reshaped to a set of N spatial masks $\{\hat{\mathbf{S}}^n \in \mathbb{R}^{w \times h}\}_{n=1}^N$ as the spatial guidance in (5).

3.3 Semantic Modulator

Figure 4 shows the diagram of the key modules in the sub-net f_{sem} (3), including the SP and the HSP. The SP exploits the high-order statistics of the holistic representation to enhance the non-linear representative capability of the

learned model (Gao et al. 2019), while the HSP can capture the long-range dependency along channel dimension of the group-wise feature tensors, paying more attention to important channels for classification task under the supervision of co-category labels.

SP: Given input feature tensor $\phi \in \mathbb{R}^{w \times h \times d}$, we firstly leverage a 1×1 convolution to reduce the number of channels from d to c to reduce the computational cost for the following operations. Then, we compute pairwise channel correlations of the $w \times h \times c$ tensor to yield a $c \times c$ covariance matrix. Each entry in the $c \times c$ covariance matrix measures the relevance between the feature maps in two channels, which leverages a quadratic operator to model high-order statistics of the holistic representation, hence enabling to enhance non-linear modeling capability. Afterwards, we use an FC layer to transform the $c \times c$ covariance matrix into a $1 \times 1 \times d$ tensor that indicates the feature channel importance.

HSP: For each image $I^n \in \mathcal{I}$, its feature tensor $\phi(I^n)$ is fed into an SP layer, outputting a $1 \times 1 \times d$ indicator tensor. Then, all the indicator tensors are concatenated vertically to yield a group-wise semantic representation, which is again fed into an SP layer to capture the long-range dependency along the channel dimension of the group-wise semantic representation, yielding an indicator vector γ that steers attention to the important channels that are essential for co-category classification.

Loss: The output γ of f_{sem} is followed by an FC layer and a sigmoid layer, yielding a classifier response:

$$\hat{\mathbf{y}} = \text{sigmoid}(\mathbf{W}\gamma + \mathbf{b}), \quad (11)$$

where $\mathbf{W} \in \mathbb{R}^{L \times d}$ and $\mathbf{b} \in \mathbb{R}^L$ are the parameters of the FC layer, L denotes the number of the co-category in the training set.

The widely used cross-entropy loss function for classification is adopted to learn the indicator γ in (11):

$$\ell_{sem} = -\frac{1}{L} \sum_{l=1}^L y_l \log \hat{y}_l - (1 - y_l) \log(1 - \hat{y}_l), \quad (12)$$

where \hat{y}_l is the l -th entry of $\hat{\mathbf{y}}$ that is the prediction value for the l -th co-category and $y_l \in \{0, 1\}$ is the ground-truth label.

3.4 Segmentation Sub-net

Given the input group-wise CFMs $\{\mathbf{X}^n\}_{n=1}^N$ of the images \mathcal{I} , the sub-net f_{seg} are modulated by the outputs $\{\mathcal{S}, \gamma\}$ of f_{spa} and f_{sem} , yielding a group of modulated representations $\{\mathbf{Y}^n\}_{n=1}^N$ using (5). Each \mathbf{Y}^n is composed of a group of multi-resolution representations $\{\mathbf{R}_i^n\}_{i=1}^4$. Similar to the FPN (Lin et al. 2017), we fuse $\{\mathbf{R}_i^n\}_{i=1}^4$ from coarse to fine: with the coarser-resolution feature maps, we use a 1×1 convolution layer to make the channel number equal to the corresponding top-down ones, following by an up-sampling layer to make their spatial resolutions the same. Then, the upsampled maps are merged with the corresponding top-down ones via element-wise addition. The process is repeated until the finest resolution maps are generated as $\mathbf{R}^n = \mathbf{R}_1^n \oplus \dots \oplus \mathbf{R}_4^n$. Finally, the maps \mathbf{R}^n are fed into a

convolutional layer, following by a 1×1 convolutional layer and an upsampling layer to generate the corresponding segmentation mask \widehat{M}^n .

Denoting the ground-truth binary co-segmentation masks in the training image group as $\mathcal{M}_{gt} = \{M_{gt}^n\}_{n=1}^N$, the loss function for the segmentation task is formulated as a weighted cross-entropy loss for pixel-wise classification:

$$\begin{aligned} \ell_{seg} = & -\frac{1}{NP} \sum_{n=1}^N \sum_{i=1}^P \{ \delta^n M_{gt}^n(i) \log \widehat{M}^n(i) \\ & - (1 - \delta^n)(1 - M_{gt}^n(i)) \log(1 - \widehat{M}^n(i)) \}, \end{aligned} \quad (13)$$

where P is the number of the pixels in each training image, i denotes the pixel index, δ^n is the ratio between all positive pixels and all pixels in image I^n , which balances the positive and negative samples.

3.5 Loss Function

The three sub-nets f_{spa} , f_{sem} and f_{seg} are trained jointly via optimizing the following multi-task loss function

$$\ell = \ell_{spa} + \ell_{sem} + \ell_{seg}, \quad (14)$$

where ℓ_{spa} , ℓ_{sem} and ℓ_{seg} are defined by (9), (12) and (13), respectively.

4 Experiments

4.1 Implementation Details

We leverage the HRNet (Sun et al. 2019) pre-trained on ImageNet (Deng et al. 2009) as the backbone network to extract the multi-resolution semantic features. Moreover, we also report the results of using the VGG16 backbone network (Simonyan and Zisserman 2014), which still demonstrate competing performance over state-of-the-art methods. Except for using the pretrained backbone network parameters as initialization, all other parameters are trained from scratch. We follow the same settings as (Wei et al. 2017; Wang et al. 2019): the input image group \mathcal{I} consists of $N = 5$ images that are randomly selected from a group of images with co-object category, and a mini-batch of $4 \times \mathcal{I}$ is fed into the model simultaneously during training. All images in \mathcal{I} are resized to 224×224 as input, and then the predicted co-segmentation maps are resized to the original image sizes as outputs. We leverage the Adam algorithm (Kingma and Ba 2014) to optimize the whole network in an end-to-end manner, among which the exponential decay rates for estimating the first and the second moments are set to 0.9 and 0.999, respectively. The learning rate starts from $1e-4$ and reduces by a half every 25,000 steps until the model converges at about 200,000 steps. Our model is implemented in PyTorch and a Nvidia RTX 2080Ti GPU is adopted for acceleration.

We adopt the COCO-SEG dataset released by (Wang et al. 2019) to train our model. The dataset contains 200,000 images belonging to $L = 78$ groups, among which each image has a manually labeled binary mask with co-category label information. The training process takes about 40 hours.

Table 1: Quantitative comparison results on the sub-set of MSRC. The bold numbers indicate the best results.

MSRC Dataset	Ave. \mathcal{P} (%)	Ave. \mathcal{J} (%)
(Vicente, Rother, and Kolmogorov 2011)	90.2	70.6
(Rubinstein et al. 2013)	92.2	74.7
(Wang, Huang, and Guibas 2013)	92.2	-
(Faktor and Irani 2013)	92.0	77.0
(Mukherjee, Lall, and Lattupally 2018)	84.0	67.0
(Li et al. 2018)	92.4	79.9
(Chen et al. 2018)	95.2	77.7
Ours-VGG16	94.3	79.4
Ours-HRNet	95.2	81.9

Table 2: Quantitative comparison results on the Internet. The bold numbers indicate the best results.

Internet Dataset	Airplane		Car		Horse	
	\mathcal{P} (%)	\mathcal{J} (%)	\mathcal{P} (%)	\mathcal{J} (%)	\mathcal{P} (%)	\mathcal{J} (%)
(Joulin et al. 2012)	47.5	11.7	59.2	35.2	64.2	29.5
(Rubinstein et al. 2013)	88.0	55.8	85.4	64.4	82.8	51.6
(Chen et al. 2014)	90.2	40.3	87.6	64.9	86.2	33.4
(Jerripothula et al. 2016)	90.5	61.0	88.0	71.0	88.3	60.0
(Quan et al. 2016)	91.0	56.3	88.5	66.8	89.3	58.1
(Sun and Ponce 2016)	88.6	36.3	87.0	73.4	87.6	54.7
(Tao et al. 2017)	79.8	42.8	84.8	66.4	85.7	55.3
(Yuan et al. 2017)	92.6	66.0	90.4	72.0	90.2	65.0
(Li et al. 2018)	94.1	65.4	93.9	82.8	92.4	69.4
(Chen et al. 2018)	-	65.9	-	76.9	-	69.1
(Chen et al. 2019)	94.1	65.0	94.0	82.0	92.2	63.0
Ours-VGG16	94.6	66.7	89.7	68.1	93.2	66.2
Ours-HRNet	94.8	69.6	91.6	82.5	94.4	70.2

4.2 Datasets and Evaluation Metrics

Datasets: We conduct extensive evaluations on four widely-used benchmark datasets (Faktor and Irani 2013; Rubinstein et al. 2013) including sub-set of MSRC, Internet, sub-set of iCoseg, and PASCAL-VOC. Among them, the sub-set of MSRC includes 7 classes: bird, car, cat, cow, dog, plane, sheep, and each class contains 10 images. The Internet has 3 categories of airplane, car and horse. Each class has 100 images including some images with noisy labels. The sub-set of iCoseg contains 8 categories, and each has a different number of images. The PASCAL-VOC is the most challenging dataset with 1,037 images of 20 categories selected from the PASCAL-VOC 2010 dataset (Everingham et al. 2010).

Evaluation Metrics: We adopt two widely-used metrics to evaluate the co-segmentation results, including the *precision* \mathcal{P} and the *Jaccard index* \mathcal{J} . The precision \mathcal{P} measures the percentage of the correctly segmented pixels for both foreground and background, while the Jaccard index \mathcal{J} is defined as the intersection area of the predicted foreground objects and the ground truth divided by their union area.

4.3 Results

We quantitatively and qualitatively compare our algorithm with several state-of-the-art co-segmentation methods on the

Table 3: Quantitative comparison results on the sub-set of iCoseg. The bold numbers indicate the best results.

iCoseg Dataset	Ave. \mathcal{J} (%)	bear2	brownbear	cheetah	elephant	helicopter	hotballoon	panda1	panda2
(Rubinstein et al. 2013)	70.2	65.3	73.6	69.7	68.8	80.3	65.7	75.9	62.5
(Jerripothula et al. 2014)	73.8	70.1	66.2	75.4	73.5	76.6	76.3	80.6	71.8
(Faktor and Irani 2013)	78.2	72.0	92.0	67.0	67.0	82.0	88.0	70.0	55.0
(Jerripothula et al. 2016)	70.4	67.5	72.5	78.0	79.9	80.0	80.2	72.2	61.4
(Li et al. 2018)	84.2	88.3	92.0	68.8	84.6	79.0	91.7	82.6	86.7
(Chen et al. 2018)	86.0	88.3	91.5	71.3	84.4	76.5	94.0	91.8	90.3
Ours-VGG16	88.0	87.4	90.3	84.9	90.6	76.6	94.1	90.6	87.5
Ours-HRNet	89.2	91.1	89.6	88.6	90.9	76.4	94.2	90.4	87.5

Table 4: Quantitative comparison results on the PASCAL-VOC. The bold numbers indicate the best results.

PASCAL-VOC Dataset	\mathcal{P} (%)	\mathcal{J} (%)	A.P.	Bike	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	D.T.	Dog	Horse	M.B.	P.S.	P.P.	Sheep	Sofa	Train	TV
(Faktor and Irani 2013)	84.0	46	65	14	49	47	44	61	55	49	20	59	22	39	52	51	31	27	51	32	55	35
(Lee et al. 2015)	69.8	33	50	15	29	37	27	55	35	34	13	40	10	37	49	44	24	21	51	30	42	16
(Chang and Wang 2015)	82.4	29	48	9	32	32	21	34	42	35	13	50	6	22	37	39	19	17	41	21	41	18
(Quan et al. 2016)	89.0	52	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
(Hati et al. 2016)	72.5	25	44	13	26	31	28	33	26	29	14	24	11	27	23	22	18	17	33	27	26	25
(Jerripothula et al. 2016)	85.2	45	64	20	54	48	42	64	55	57	21	61	19	49	57	50	34	28	53	39	56	38
(Jerripothula et al. 2017)	80.1	40	53	14	47	43	42	62	50	49	20	56	13	38	50	45	29	26	40	37	51	37
(Wang et al. 2017)	84.3	52	75	26	53	59	51	70	59	70	35	63	26	56	63	59	35	28	67	52	52	48
(Li et al. 2018)	94.2	65	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
(Hsu et al. 2018)	91.0	60	77	27	70	61	58	79	76	79	29	75	28	63	66	65	37	42	75	67	68	51
Ours-VGG16	93.7	66	83	35	75	69	58	87	77	80	26	86	7	74	79	71	45	39	81	68	83	59
Ours-HRNet	94.9	71	82	37	74	70	67	88	82	77	36	87	15	75	82	72	58	46	82	77	84	69

four benchmark datasets.

Quantitative Results: Tables 1, 2, 3, 4 list the comparison results of our method with other state-of-the-arts on the sub-set of MSRC, Internet, sub-set of iCoseg and PASCAL-VOC. For fair comparisons, the reported results of the compared methods are directly obtained from their publications. We can observe that our algorithm outperforms the other state-of-the-arts in term of both metrics on most object categories in each dataset. Especially on the PASCAL-VOC, which has more challenging scenarios, the proposed algorithm achieves the best average \mathcal{P} and average \mathcal{J} with a score of 94.9% and 71%, respectively, significantly outperforming the others by a large margin. Moreover, on the sub-set of MSRC and sub-set of iCoseg, the average \mathcal{J} by our method has a score of 81.9% and 89.2%, outperforming the others by about 3%. Besides, on the Internet, our algorithm achieves the best performance on airplane and horse categories, as well as a competitive performance on car category in terms of both metrics average \mathcal{P} and average \mathcal{J} .

Qualitative Results: Figure 5 shows some qualitative results by comparing our method with SAAB (Chen et al. 2018) and DCOS (Li et al. 2018). Those images are chosen from all of the four datasets composed of co-objects with seen categories (inside the 78 categories of the COCO-SEG) and unseen categories (outside the categories of the COCO-SEG). For the seen categories shown by Figure 5(a), we can observe that SAAB and DCOS cannot discover the co-objects in the dog group accurately and two distractors (sheep) have been mis-classified as co-objects. However, the proposed approach does not suffer from this issue since it uses co-category labels as supervision to learn an effective semantic modulator that can well capture high-level

Table 5: Ablative experiments of the proposed model on the PASCAL-VOC. The bold numbers indicate the best results. The symbol ‘-f’ denotes removing the module f .

PASCAL-VOC	Avg. \mathcal{P} (%)	Avg. \mathcal{J} (%)
$f_{spa} \& f_{sem} \& f_{seg}$	94.9	71
- f_{spa}	94.5	69
- f_{sem}	85.0	38
- ($f_{spa} \& f_{sem}$)	82.0	27

semantic category information. Besides, as shown by Figure 5(a), (b), the proposed approach can discover the whole co-objects of seen and unseen categories well because its spatial modulator is learned by an unsupervised method that can not only help to locate the co-object regions of seen categories well, but also generalize well to unseen categories.

4.4 Ablative Study

To further show our main contributions, we compare different variants of our model including those without spatial modulator ($-f_{spa}$), semantic modulator ($-f_{sem}$) and both modulators ($-(f_{spa} \& f_{sem})$), respectively. Table 5 lists the results of ablative experiments on the PASCAL-VOC. We can observe that without f_{spa} , the average \mathcal{P} score drops from 94.9% to 94.5% while the average \mathcal{J} score reduces by 2% from 71% to 69%, which verifies the effectiveness of the proposed module f_{spa} . Moreover, without f_{sem} , the performance suffers from a significant loss with a big drop of 9.9% and 33% for the average \mathcal{P} and \mathcal{J} scores, respectively, indicating the critical role of the semantic modulator as a guidance to learn an effective segmentation network for ac-

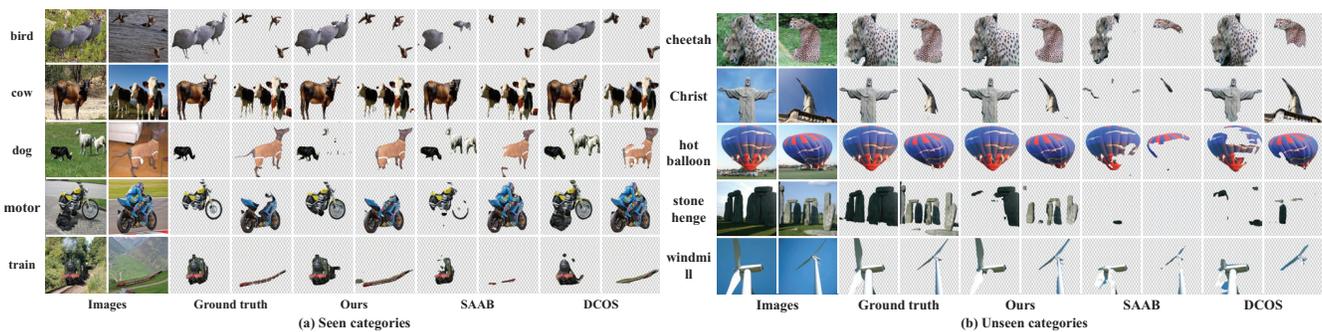


Figure 5: Some qualitative comparison results generated by the proposed method, SAAB (Chen et al. 2018) and DCOS (Li et al. 2018) for co-segmenting objects associated to the training categories and unseen categories, respectively.

curate co-segmentation. Besides, compared to that only removes f_{sem} , removing both modulators f_{spa} and f_{sem} further makes the performance of our model drop by 3% and 11% in terms of average \mathcal{P} and average \mathcal{J} , respectively. These experiments confidently validate that both modulators have a positive effect to boost the performance of our model.

5 Conclusions

In this paper, we have presented a spatial-semantic modulated deep network framework for object co-segmentation. Our model is composed of a spatial modulator, a semantic modulator and a segmentation sub-net. The spatial modulator is to learn a mask to coarsely localize the co-object regions in each image that captures the correlations of image feature descriptors with unsupervised learning. The semantic modulator is to learn a channel importance indicator under the supervision of co-category labels. We have proposed the HSP module to transform the input image features of the semantic modulator for classification use. The outputs of the two modulators manipulate the input feature maps of the segmentation sub-net by a simple shift-and-scale operation to adapt it to target on segmenting the co-object regions. Both quantitative and qualitative evaluations on four image co-segmentation benchmark datasets have demonstrated superiority of the proposed method to the state-of-the-arts.

Acknowledgments

This work is supported in part by National Major Project of China for New Generation of AI (No. 2018AAA0100400), in part by the Natural Science Foundation of China under Grant nos. 61876088, 61825601, in part by the Natural Science Foundation of Jiangsu Province under Grant no. BK20170040.

References

Brockschmidt, M. 2019. Gnn-film: Graph neural networks with feature-wise linear modulation. *arXiv preprint arXiv:1906.12192*.

Chang, H.-S., and Wang, Y.-C. F. 2015. Optimizing the decomposition for multiple foreground cosegmentation. *CVIU* 141:18–27.

Chen, X.; Shrivastava, A.; Gupta, A.; and et al. 2014. Enriching visual knowledge bases via object discovery and segmentation. In *CVPR*, 2027–2034.

Chen, L.; Zhang, H.; Xiao, J.; Nie, L.; Shao, J.; Liu, W.; and Chua, T.-S. 2017. Sca-cnn: Spatial and channel-wise attention in convolutional networks for image captioning. In *CVPR*, 5659–5667.

Chen, H.; Huang, Y.; Nakayama, H.; and et al. 2018. Semantic aware attention based deep object co-segmentation. In *ACCV*, 435–450.

Chen, Y.-C.; Lin, Y.-Y.; Yang, M.-H.; and Huang, J.-B. 2019. Show, match and segment: Joint learning of semantic matching and object co-segmentation. In *arXiv*.

Collins, M. D.; Xu, J.; Grady, L.; and Singh, V. 2012. Random walks based multi-image segmentation: Quasiconvexity results and gpu-based solutions. In *CVPR*, 1656–1663.

Dai, J.; Nian Wu, Y.; Zhou, J.; and Zhu, S.-C. 2013. Cosegmentation and cosketch by unsupervised learning. In *ICCV*, 1305–1312.

Dai, J.; He, K.; and Sun, J. 2015. Convolutional feature masking for joint object and stuff segmentation. In *CVPR*, 3992–4000.

De Vries, H.; Strub, F.; Mary, J.; Larochelle, H.; Pietquin, O.; and Courville, A. C. 2017. Modulating early visual processing by language. In *NIPS*, 6594–6604.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*, 248–255.

Ding, C., and He, X. 2004. K-means clustering via principal component analysis. In *ICML*, 29.

Everingham, M.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *IJCV* 88(2):303–338.

Faktor, A., and Irani, M. 2013. Co-segmentation by composition. In *ICCV*, 1297–1304.

Flores, C. F.; Gonzalez-Garcia, A.; van de Weijer, J.; and Raducanu, B. 2019. Saliency for fine-grained object recognition in domains with scarce training data. *PR* 94:62–73.

Gao, Z.; Xie, J.; Wang, Q.; and Li, P. 2019. Global second-

- order pooling convolutional networks. In *CVPR*, 3024–3033.
- Han, J.; Quan, R.; Zhang, D.; and Nie, F. 2018. Robust object co-segmentation using background prior. *T-IP* 27(4):1639–1651.
- Hati, A.; Chaudhuri, S.; Velmurugan, R.; and et al. 2016. Image co-segmentation using maximum common subgraph matching and region co-growing. In *ECCV*, 736–752.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Spatial pyramid pooling in deep convolutional networks for visual recognition. *T-PAMI* 37(9):1904–1916.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*, 2961–2969.
- Hsu, K.-J.; Lin, Y.-Y.; Chuang, Y.-Y.; and et al. 2018. Co-attention cnns for unsupervised object co-segmentation. In *IJCAI*, 748–756.
- Jerripothula, K. R.; Cai, J.; Meng, F.; and Yuan, J. 2014. Automatic image co-segmentation using geometric mean saliency. In *ICIP*, 3277–3281.
- Jerripothula, K. R.; Cai, J.; Yuan, J.; and et al. 2016. Image co-segmentation via saliency co-fusion. *T-MM* 18(9):1896–1909.
- Jerripothula, K. R.; Cai, J.; Lu, J.; and Yuan, J. 2017. Object co-skeletonization with co-segmentation. In *CVPR*, 3881–3889.
- Joulin, A.; Bach, F.; Ponce, J.; and et al. 2012. Multi-class cosegmentation. In *CVPR*, 542–549.
- Kamranian, Z.; Nilchi, A. R. N.; Monadjemi, A.; and Navab, N. 2018. Iterative algorithm for interactive co-segmentation using semantic information propagation. *AI* 48(12):5019–5036.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lee, C.; Jang, W.-D.; Sim, J.-Y.; and Kim, C.-S. 2015. Multiple random walkers and their application to image cosegmentation. In *CVPR*, 3837–3845.
- Li, W.; Jafari, O. H.; Rother, C.; and et al. 2018. Deep object co-segmentation. In *ACCV*, 638–653.
- Lin, T. Y.; Dollar, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *CVPR*, 2117–2125.
- Lu, Z.; Xu, H.; and Liu, G. 2019. A survey of object co-segmentation. *IEEE Access* 7:62875–62893.
- Mukherjee, P.; Lall, B.; and Lattupally, S. 2018. Object cosegmentation using deep siamese network. *arXiv preprint arXiv:1803.02555*.
- Mustafa, A., and Hilton, A. 2017. Semantically coherent co-segmentation and reconstruction of dynamic scenes. In *CVPR*, 422–431.
- Perez, E.; Strub, F.; De Vries, H.; Dumoulin, V.; and Courville, A. 2018. Film: Visual reasoning with a general conditioning layer. In *AAAI*.
- Prol, H.; Dumoulin, V.; and Herranz, L. 2018. Cross-modulation networks for few-shot learning. *arXiv preprint arXiv:1812.00273*.
- Quan, R.; Han, J.; Zhang, D.; and Nie, F. 2016. Object co-segmentation via graph optimized-flexible manifold ranking. In *CVPR*, 687–695.
- Rubinstein, M.; Joulin, A.; Kopf, J.; and Liu, C. 2013. Unsupervised joint object discovery and segmentation in internet images. In *CVPR*, 1939–1946.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sun, J., and Ponce, J. 2016. Learning dictionary of discriminative part detectors for image categorization and cosegmentation. *IJCV* 120(2):111–133.
- Sun, K.; Xiao, B.; Liu, D.; and Wang, J. 2019. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 5693–5703.
- Tao, Z.; Liu, H.; Fu, H.; and Fu, Y. 2017. Image cosegmentation via saliency-guided constrained clustering with cosine similarity. In *AAAI*.
- Tsai, C.-C.; Li, W.; Hsu, K.-J.; Qian, X.; and Lin, Y.-Y. 2018. Image co-saliency detection and co-segmentation via progressive joint optimization. *T-IP* 28(1):56–71.
- Vicente, S.; Rother, C.; and Kolmogorov, V. 2011. Object cosegmentation. In *CVPR*, 2217–2224.
- Wang, C.; Zhang, H.; Yang, L.; Cao, X.; and Xiong, H. 2017. Multiple semantic matching on augmented n -partite graph for object co-segmentation. *T-IP* 26(12):5825–5839.
- Wang, C.; Zha, Z.-J.; Liu, D.; and Xie, H. 2019. Robust deep co-saliency detection with group semantic. In *AAAI*.
- Wang, F.; Huang, Q.; and Guibas, L. J. 2013. Image co-segmentation via consistent functional maps. In *ICCV*, 849–856.
- Wei, L.; Zhao, S.; El Farouk Bourahla, O.; Li, X.; and Wu, F. 2017. Group-wise deep co-saliency detection. In *IJCAI*, 3041–3047.
- Wei, X. S.; Zhang, C. L.; Wu, J.; Shen, C.; and Zhou, Z. H. 2019. Unsupervised object discovery and co-localization by deep descriptor transforming. *PR* 88:113–126.
- Yang, L.; Wang, Y.; Xiong, X.; Yang, J.; and Katsaggelos, A. K. 2018. Efficient video object segmentation via network modulation. In *CVPR*, 6499–6507.
- Yuan, Z.; Lu, T.; Wu, Y.; and et al. 2017. Deep-dense conditional random fields for object co-segmentation. In *IJCAI*, 3371–3377.
- Yuan, Z.; Lu, T.; and Shivakumara, P. 2014. A novel topic-level random walk framework for scene image co-segmentation. In *ECCV*, 695–709.
- Zhang, K.; Li, T.; Liu, B.; and Liu, Q. 2019. Co-saliency detection via mask-guided fully convolutional networks with multi-scale label smoothing. In *CVPR*, 3095–3104.
- Zhu, H.; Meng, F.; Cai, J.; and Lu, S. 2016. Beyond pixels: A comprehensive survey from bottom-up to semantic image segmentation and cosegmentation. *JVCIR* 34:12–27.