

Unsupervised and semi-supervised co-salient object detection via segmentation frequency statistics

Souradeep Chakraborty^{1,2†}, Shujon Naha^{2,3}, Muhammet Bastan², Amit Kumar K C², Dimitris Samaras¹

¹Stony Brook University ²Visual Search & AR, Amazon ³Indiana University

{souchakrabor, samaras}@cs.stonybrook.edu, {mbastan, amitkrkc}@amazon.com, snaha@iu.edu

Abstract

In this paper, we address the detection of co-occurring salient objects (CoSOD) in an image group using frequency statistics in an unsupervised manner, which further enable us to develop a semi-supervised method. While previous works have mostly focused on fully supervised CoSOD, less attention has been allocated to detecting co-salient objects when limited segmentation annotations are available for training. Our simple yet effective unsupervised method US-CoSOD combines the object co-occurrence frequency statistics of unsupervised single-image semantic segmentations with salient foreground detections using self-supervised feature learning. For the first time, we show that a large unlabeled dataset e.g. ImageNet-1k can be effectively leveraged to significantly improve unsupervised CoSOD performance. Our unsupervised model is a great pre-training initialization for our semi-supervised model SS-CoSOD, especially when very limited labeled data is available for training. To avoid propagating erroneous signals from predictions on unlabeled data, we propose a confidence estimation module to guide our semi-supervised training. Extensive experiments on three CoSOD benchmark datasets show that both of our unsupervised and semi-supervised models outperform the corresponding state-of-the-art models by a significant margin (e.g., on the Cosal2015 dataset, our US-CoSOD model has an 8.8% F-measure gain over a SOTA unsupervised co-segmentation model and our SS-CoSOD model has an 11.81% F-measure gain over a SOTA semi-supervised CoSOD model).

1. Introduction

Co-salient object detection (CoSOD) focuses on detecting co-existing salient objects in an image group, whereas salient object detection (SOD) detects the same salient objects in single images [6, 36, 37, 43, 47, 52, 63]. CoSOD leverages the extra knowledge that the group images share a common object by finding semantic similarities across the




Figure 1. Visualization of co-saliency detections on an image group *train* from the Cosal2015 dataset [65]. Row 1: Original image, Row 2: DCFM [62] predictions (trained using all labeled data), Row 3: Our unsupervised model (US-CoSOD), Row 4: Our semi-supervised model (SS-CoSOD) trained using 1/4 labeled data with unsupervised pre-training, which has comparable performance to the fully supervised DCFM trained with all labels.

image regions in the group. Thus CoSOD models can localize the salient objects more accurately compared to the single image based SOD models [13, 16] in such image groups. Both SOD and CoSOD are joint segmentation and detection tasks as shown in the existing literature [13, 15, 62] and thus require segmentation labels. However, collecting segmentation annotations is time-consuming as well as expensive.

This annotation requirement is a drawback for a majority of the existing CoSOD models [13, 15, 62, 67, 71] that are fully supervised. To relieve the labeling burden, some works [21, 22, 39] focused on unsupervised co-segmentation and co-saliency detection. Semi-supervised learning in CoSOD [72] aims to learn an effective model from a training dataset using only a small set of labeled images along with a larger unlabeled set. Such models have an immense value in several real-world industrial applications such as in e-commerce (e.g. automatic product detection from customer review and query images of the product without the need for manual product annotations), content-based image retrieval, satellite imaging, bio-medical imaging, etc.

[†]Part of this work was done during an internship at Amazon.

However, the prediction performance of these models is significantly worse compared to the existing fully supervised models due to their inefficient use of the unlabeled data. In this paper, we first solve unsupervised CoSOD using a large unlabeled dataset and next use the unsupervised model as a pre-training initialization for our semi-supervised pipeline.

In this work, we take advantage of the recent progress in self-supervised semantic segmentation [18] as well as self-supervised self-attention [3] to develop a simple yet effective unsupervised algorithm for CoSOD (US-CoSOD). As part of our unsupervised approach, we first obtain the segmentation masks of the co-occurring objects in an image group using STEGO, an off-the-shelf self-supervised semantic segmentation model [18]. Next, we select the most common and salient segmentation mask (with guidance from the self-attention maps obtained from DINO [3], a self-supervised feature learning method) as the pseudo segmentation label for training an off-the-shelf CoSOD model in a supervised manner. We show a significant improvement in prediction performance using our methods. However, standard training datasets are relatively small. In our paper, we introduce a more up-to-date evaluation task for unsupervised CoSOD on a set of 150K unlabeled images from the ImageNet-1k dataset [10,31] (which only contains class labels without any segmentation annotation).

Next, we show that our unsupervised model forms a strong pre-training initialization for a CoSOD model trained in a semi-supervised manner. For this, we propose a confidence aware student-teacher architecture based semi-supervised model, SS-CoSOD. Here, we leverage the fact that in an input image group for CoSOD, we can mix the labeled and unlabeled images to effectively propagate knowledge from the labeled images to the unlabeled images in the image group via cross-region correspondences. We also introduce a confidence estimation module to block erroneous knowledge flow from inaccurate predictions on difficult unlabeled images. Similar to US-CoSOD, we leverage the large unlabeled ImageNet-1k [10,31] dataset to significantly improve semi-supervised CoSOD performance.

In Fig. 1, we compare our unsupervised and semi-supervised models with DCFM [62], a state-of-the-art fully supervised CoSOD model. Our US-CoSOD produces segmentations comparable with DCFM and our SS-CoSOD model further improves the segmentation predictions. Our main contributions are:

- We propose a simple yet effective unsupervised approach for CoSOD that effectively leverages single-image semantic segmentations and self-attention maps generated using self-supervision to generate pseudo-labels for supervised training of a CoSOD model.
- For the first time, we show that CoSOD can be significantly improved using large unlabeled datasets,

e.g. ImageNet-1k [31]. This approach helps us achieve state-of-the-art results for unsupervised CoSOD.

- We propose a novel approach for semi-supervised CoSOD by effectively propagating knowledge from a limited labeled set to a much larger unlabeled set via confidence estimation and cross-region correspondence between the labeled and the unlabeled sets.

2. Related Work

Co-salient object detection: Graphical models are used to model pixel relationships in an image group [23,26–28,60,68], followed by mining co-salient objects with consistent features. Some works used additional object saliency information to first mine the salient objects and then implement CoSOD [29,69,70]. Other works compute the shared attributes among input images [13,15,17,34,38,51,67,71,73] and supplement semantic information with classification information. The surveys [9,14,64] provide more information on CoSOD. DCFM [62] mines co-salient features with democracy while reducing background interference. We use DCFM as the backbone network in our study.

Unsupervised segmentation: Several unsupervised semantic segmentation approaches use self-supervised feature learning techniques [8,25,40,54,56]. Recently, STEGO [18] showed that semantically correlated dense features from unsupervised feature learning frameworks can help distill unsupervised features into high-quality semantic labels. We use this model as a component in our unsupervised pipeline. In [61], the semantic categories obtained using self-supervised learning are mapped to pixel-level features via class activation maps, which serve as pseudo labels for training. Some papers solve unsupervised co-segmentation [2,4,21,24,39] and CoSOD [22,65]. Li et al. [39] proposed an unsupervised co-segmentation model by ranking image complexities using saliency maps. Hsu et al. [21] developed an unsupervised co-attention based model for object co-segmentation. The same authors presented an unsupervised graphical model for CoSOD [22] that jointly solves single-image saliency and object co-occurrence. Our US-CoSOD outperforms all of these unsupervised models.

Semi-supervised segmentation: While consistency-based methods enforce the predictions of unlabeled samples to be consistent under different perturbations [5,33,46], pseudo-label based methods [30,32,35,42,45,58], incorporate unlabeled data into training with high-quality pseudo labels. Some of these methods [30,32,45,58] apply pixel-level error correction mechanisms on the generated pseudo-labels (*e.g.* using an auxiliary decoder or employing a flaw detector or a discriminator) in order to avoid propagating label noise. Instead, we directly estimate the prediction error probability on an unlabeled image at the global level using a confidence estimation module, trained to estimate pre-

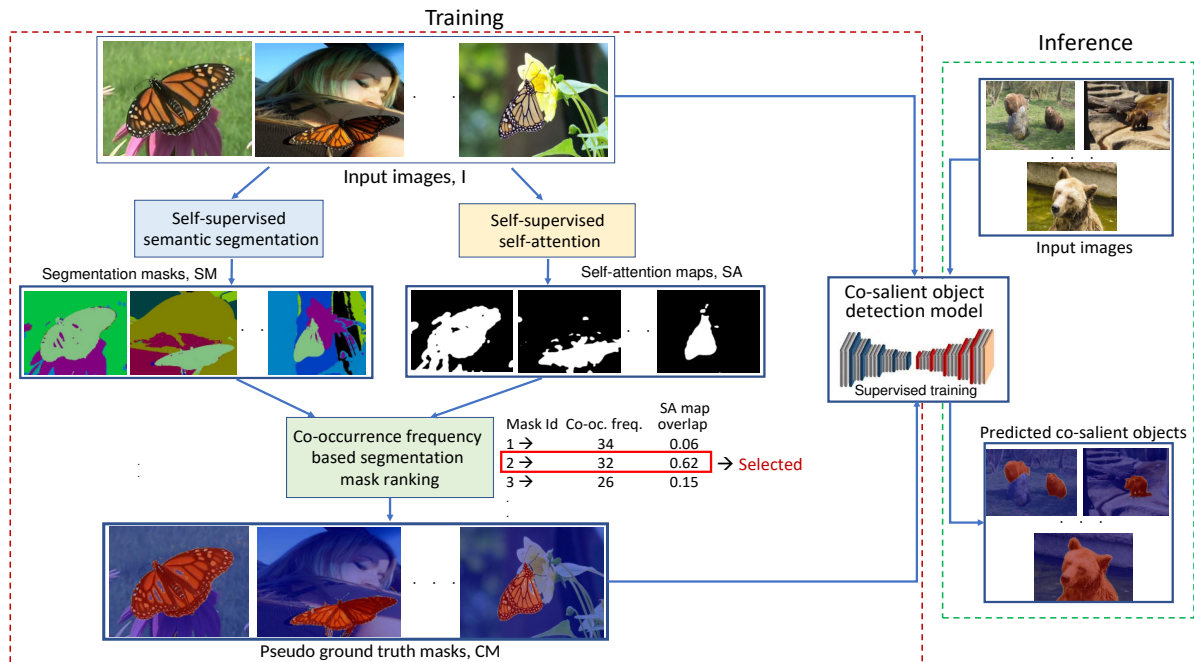


Figure 2. The proposed unsupervised CoSOD model, US-CoSOD. We first obtain unsupervised semantic segmentation maps from STEGO [18] and self-attention maps from DINO [3]. The unlabeled categories (from STEGO) in the image group are sorted based on their co-occurrence frequencies and the final category is chosen based on the overlap between the STEGO and DINO masks per image. The segmentation mask of the selected category is considered the pseudo ground truth for training an off-the-shelf supervised CoSOD model.

diction confidence (in terms of the segmentation accuracy) based on the labeled set to block error propagation. Wang et al. [59] proposed the first semi-supervised co-segmentation model by optimizing an energy function consisting of inter- and intra-image distances for an image group. While Zheng et al. [72] proposed the first semi-supervised CoSOD framework based on graph structure optimization, their model accuracy is low due to the use of hand-crafted features and their model has not been evaluated under sufficiently low labeled data. Some semi-supervised SOD models have also been proposed [42, 44, 66]. SAL [44] used active learning to gradually expand a small labeled set to include samples on which predictions are inaccurate. GWSCoSal [48] introduced a weakly supervised learning induced CoSOD model using group class activation maps.

Existing unsupervised and semi-supervised approaches suffer from limited performance because they use: (1) hand-crafted features, and (2) smaller unlabeled datasets. Our study fills this gap by introducing an unsupervised and a semi-supervised CoSOD method, both of which uses large-scale unlabeled data to significantly improve performance.

3. Methodology

Given a group of N images $I = \{I_1, I_2, \dots, I_n\}$ containing co-occurring salient objects of a certain class, CoSOD aims to detect them simultaneously and output their co-salient object segmentation masks. For unsuper-

vised CoSOD, the goal is to predict the co-salient segmentations $\{\hat{y}_i\}_{i=1}^n$ without using any labeled data. For semi-supervised CoSOD, given a labeled set $\mathcal{D}_l = \{(x_i^l, y_i^l)\}_{i=1}^{N_l}$ and a much larger unlabeled set $\mathcal{D}_u = \{(x_i^u)\}_{i=1}^{N_u}$, we aim to train a CoSOD model by efficiently utilizing both the limited labeled data and a large amount of unlabeled data.

3.1. Unsupervised co-salient object detection

Here, we describe our unsupervised CoSOD model (US-CoSOD) that effectively leverages the frequency statistics of self-supervised single-image segmentations.

Fig. 2 depicts our unsupervised pipeline for CoSOD. We first compute the pseudo co-saliency masks based on the single image segmentation masks and the self-attention masks, which are then used to train a fully supervised CoSOD model. Trained with the self-distillation loss [20], the attention maps associated with the class token from the last layer of DINO [3] have been shown to highlight salient foreground regions [3, 57, 61]. Motivated by this observation, we consider the averaged attention map (across all attention heads) from DINO as the foreground object segmentation. Also, to detect the co-occurring objects, we use the semantic segmentation masks from a recent self-supervised single-image semantic segmentation model, STEGO [18]. This model shows that feature correspondences across images form strong signals for unsupervised semantic segmentation. These correspondences are used to create pixel-wise

Algorithm 1 Pseudo co-saliency mask generation

Input: Image group $I = \{I_1, I_2, \dots, I_n\}$ **Output:** CoSOD masks $CM = \{CM_1, CM_2, \dots, CM_n\}$

- 1: Obtain self-attention (SA) maps, $SA = \{SA_i\}_{i=1}^n$ from DINO [3].
 - 2: Apply Otsu thresholding on the SA maps to obtain binary segmentation maps, $DM = \{DM_i\}_{i=1}^n$.
 - 3: Obtain the unsupervised single-image semantic segmentation maps, SM_i^c for each image i and discovered unlabeled category c from STEGO [18].
 - 4: Compute the frequency, f^c of each semantic unlabeled category c from STEGO in the image group, I .
 - 5: **for** $i = 1$ to n **do**
 - 6: $C_i = \{c^1, c^2, \dots, c^m\}$ is the set of discovered unlabeled categories in the image I_i .
 - 7: Sort the categories in C_i by their frequency f^c in the descending order and select the top- K frequent unlabeled categories, $U = \{u^1, u^2, \dots, u^K\}$.
 - 8: **for** $j = 1$ to K **do**
 - 9: For category u^j , compute overlap score:
 $O_i^{u^j} = Ar(SM_i^{u^j} \cap DM_i) / Ar(I_i)$, the overlapped area between the STEGO mask $SM_i^{u^j}$ for category u^j and the DINO SA map DM_i divided by the total image area.
 - 10: **end for**
Co-salient object mask, $CM_i = SM_i^{c_i^{coso}}$ is the STEGO mask of the class, c_i^{coso} that maximizes the overlap score O_i i.e. $c_i^{coso} = \arg \max_c O_i^c$
 - 11: **end for**
 - 12: **return** CM
-

embeddings, which yield high quality semantic segmentation maps upon clustering. We consider these co-occurring semantic clusters across the image group as unlabeled categories and leverage them to find the co-occurring objects.

We detail our unsupervised pseudo co-saliency mask generation in Algorithm 1. First, we average the self-attention maps from the n_h DINO attention heads to obtain the averaged self-attention map SA_i for an image I_i as: $SA_i = \frac{1}{n_h} \sum_{j=1}^{n_h} AM_i^j$, where AM_i^j is the attention map from the DINO attention head j for the image I_i . Map SA_i is normalized by min-max normalization. We then find the co-occurrence frequency f^c of the discovered categories across all images in the group I . Next, for each image I_i , we compute the top- K frequent STEGO categories and finally select a single unlabeled category c_i^{coso} per image based on the overlap score O_i^c between the STEGO and the DINO masks. We then consider the STEGO mask $SM_i^{c_i^{coso}}$ corresponding to the category c_i^{coso} as the pseudo ground truth mask for I_i . This filtering step ensures that the selected segmentation corresponds to the most common yet salient

object in the group, therefore preventing co-occurring backgrounds from being considered as the pseudo masks.

Thus we obtain the pseudo co-salient object masks CM_{train} for all groups I_{train} in our training set and train a CoSOD model [62] in a supervised manner using I_{train} as input and the corresponding pseudo segmentation masks CM_{train} as the training labels. The training loss, \mathcal{L}^{unsp} is:

$$\mathcal{L}^{unsp} = \frac{1}{|B|} \sum_{i=1}^{|B|} l_{iou}(f^{unsp}(x_i, \theta), CM_i) + \lambda_{sc} l_{sc} \quad (1)$$

where, l_{iou} is the IoU loss [49, 70] between the predicted segmentation map, $f^{unsp}(x_i, \theta)$ and the ground truth segmentation CM_i , B is the training batch, and x_i is the input image. l_{sc} is the self-contrastive loss as outlined in [62].

At inference, we use the trained f^{unsp} model to detect co-salient objects in the test image groups, I_{test} . Note that the self-supervised component models (STEGO and DINO) are only used during the training of our US-CoSOD model (to generate the pseudo-labels) and not during inference.

3.2. Semi-supervised co-salient object detection

While the unsupervised model does not require any labeled data for training, such a model is often not at par with supervised models in terms of prediction accuracy. Therefore, we develop a semi-supervised approach for CoSOD (SS-CoSOD), which can effectively leverage a large amount of unlabeled data with effective prediction confidence estimation. Fig. 3 depicts our semi-supervised pipeline.

In stage 1, we supervisedly pre-train a CoSOD model f_{PT} on the labeled set and then train our Confidence Estimation Network (CEN) on the same set. In stage 2, we employ a typical self-training framework with two models of the same architecture, namely student (model f_S) and teacher (model f_T) networks that are initialized as f_{PT} in stage 1. The student model's weights θ_s are updated via backpropagation while the teacher model's weights θ_t are updated with the exponential moving average (EMA) scheme [53], i.e., $\theta_t = \lambda_d \theta_{t-1} + (1 - \lambda_d) \theta_s$, where λ_d is the EMA decay factor (set to 0.95). At each training step, we sample B_l labeled images and B_u unlabeled images (maximum value of $|B_l|$ and $|B_u|$ being 16, following [62]). Next, we combine each batch of labeled images, B_l and unlabeled images, B_u into a single volume B_{l+u} before passing the combined volume through the student network. The training label for this volume is the combination of the ground truth labeled mask with the teacher model prediction as: $y^{(l+u)} = [y^l, f_T(x_j^u, \theta_T)]$. This step is different from other semi-supervised approaches, e.g., U2PL [58] where the labeled and the unlabeled sets are passed through the student model in two separate passes. We leverage the benefit of learning cross-pixel similarities across all images in the im-

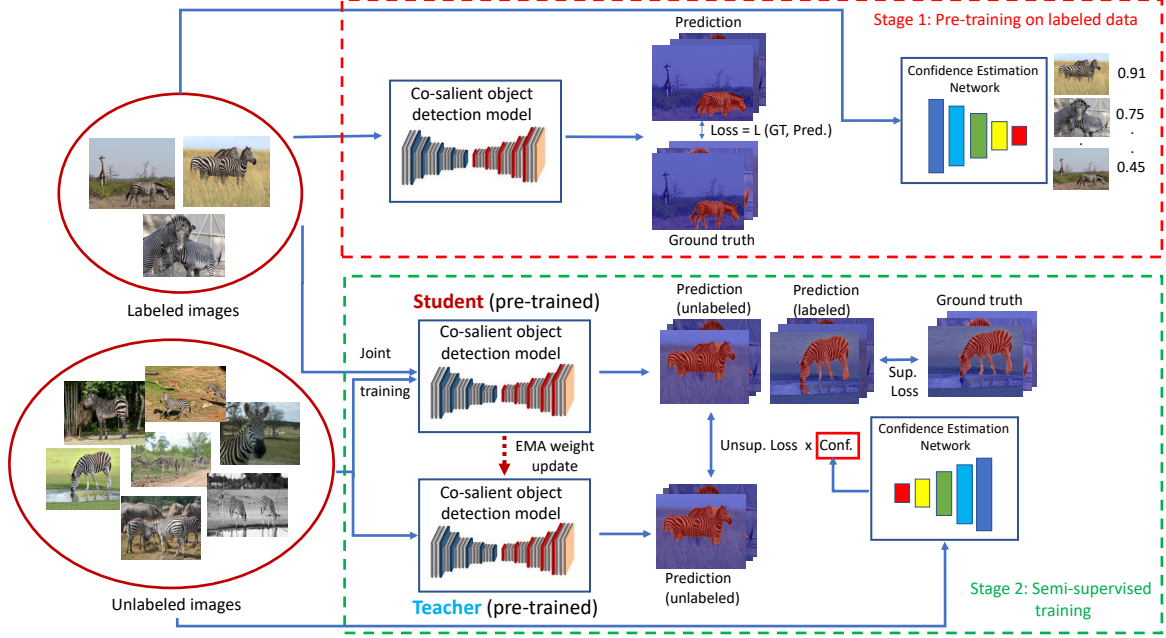


Figure 3. Proposed SS-CoSOD model for semi-supervised co-salient object detection. In the first stage, we pre-train a CoSOD model on the labeled set while also training our Confidence Estimation Network (CEN) on the same set. In stage 2, we employ a student-teacher model (initialized with the model from stage 1) for semi-supervised learning. The labeled and unlabeled data are jointly passed through the student during training. We weight the unsupervised loss by the confidence score estimated by the CEN module.

age group to effectively propagate object co-saliency information from the labeled images to the unlabeled images.

A supervised loss, \mathcal{L}_s is computed for the labeled set and an unsupervised loss, \mathcal{L}_u is computed for the unlabeled samples. For every labeled image, our goal is to minimize the supervised IoU loss as:

$$\mathcal{L}_s = \frac{1}{|B_l|} \sum_{\substack{(x_i^l, y_i^l) \in B_l, \\ x_j^u \in B_u}} l_{iou}(f_S^l([x_i^l, x_j^u], \theta_S), y_i^l) + \lambda_{sc} l_{sc} \quad (2)$$

where l_{iou} is IoU loss and l_{sc} is self-contrastive (SC) loss. λ_{sc} is set as 0.1 (following [62]). The SC Loss, l_{sc} is computed as: $l_{sc} = -\log(\cos_c + \epsilon) - \log(1 - \cos_b + \epsilon)$ with $\cos_c = \cos(\text{proto}^{(l+u)}, \text{proto}_c^{(l+u)})$ and $\cos_b = \cos(\text{proto}^{(l+u)}, \text{proto}_b^{(l+u)})$, where proto is the prototype generated by the original inputs, proto_c is the co-salient prototype generated by the foreground regions, and proto_b is the background prototype generated from the background regions in the image. \cos is the cosine-similarity function and ϵ is a small positive constant to avoid overflow.

We pass the unlabeled batch, B_u , through the teacher network and compute the unsupervised loss between the predictions of the teacher and the student networks. The unsupervised loss is weighted using the confidence scores predicted by the CEN module for each unlabeled image as:

$$\mathcal{L}_u = \frac{1}{|B_u|} \sum_{\substack{x_i^u \in B_l, \\ x_j^u \in B_u}} \tilde{g}(x_j^u, \theta_C^l) l_{iou}(f_S^u([x_i^l, x_j^u], \theta_S), f_T(x_j^u, \theta_T)) \quad (3)$$

where $\tilde{g}(x_j^u, \theta_C^l) = \frac{g(x_j^u, \theta_C^l)}{\sum_{j=1}^{|B_u|} g(x_j^u, \theta_C^l)}$ is the confidence weight estimated by CEN, g for the unlabeled sample x_j^u , parameterized by θ_C^l that is learned from the labeled batch B_l . We observed that the normalized confidence weight is crucial for model convergence. Our objective is to minimize the overall loss, $\mathcal{L} = \mathcal{L}_s + \lambda_u \mathcal{L}_u$, where \mathcal{L}_s and \mathcal{L}_u represent supervised loss on the labeled set and unsupervised loss on the unlabeled set respectively. λ_u is set as 1 [62].

Confidence Estimation Network (CEN): The CEN model is trained to estimate the reliability score of the model prediction. To train this model, we use the labeled image set S_l as the input and the corresponding segmentation F-measure [1] scores of the pre-trained model f_{PT} predictions in stage 1 as the ground truth. We use a ResNet50 backbone [19] trained using the DINO method because of its ability to well segment the discriminative image regions. The model is trained by fine-tuning the pre-trained ResNet backbone and an $fc(2048, 1)$ layer using the MSE loss as:

$$\mathcal{L}_C = \frac{1}{|S_l|} \sum_{(x_i^l, y_i^l) \in S_l} (\text{CEN}(x_i^l, \theta_C^l) - F_\beta(f_{PT}(x_i^l, \theta_S), y_i^l))^2 \quad (4)$$

Table 1. Performance comparison of unsupervised models for CoSOD. Our US-CoSOD-ImgNet150 model achieves the best result.

Method	CoCA				Cosal2015				CoSOD3k			
	MAE↓	$F_\beta^{max} \uparrow$	$E_\phi^{max} \uparrow$	$S_\alpha \uparrow$	MAE↓	$F_\beta^{max} \uparrow$	$E_\phi^{max} \uparrow$	$S_\alpha \uparrow$	MAE↓	$F_\beta^{max} \uparrow$	$E_\phi^{max} \uparrow$	$S_\alpha \uparrow$
DINO (DI) [3] (ICCV 2021)	0.214	0.372	0.572	0.540	0.154	0.659	0.753	0.688	0.146	0.624	0.749	0.679
STEGO (ST) [18] (ICLR 2022)	0.235	0.353	0.555	0.523	0.164	0.618	0.717	0.676	0.204	0.543	0.660	0.615
TokenCut [56] (CVPR 2022)	0.167	0.467	0.704	0.627	0.139	0.805	0.857	0.793	0.151	0.720	0.811	0.744
DVFDVD [2] (ECCVW 2022)	0.223	0.422	0.592	0.581	0.092	0.777	0.842	0.809	0.104	0.722	0.819	0.773
SegSwap [50] (CVPRW 2022)	0.165	0.422	0.666	0.567	0.178	0.618	0.720	0.632	0.177	0.560	0.705	0.608
Ours (DI+ST)	0.165	0.461	0.676	0.610	0.112	0.760	0.823	0.767	0.124	0.684	0.793	0.724
Ours (US-CoSOD-COCO9213)	0.140	0.498	0.702	0.641	0.090	0.792	0.852	0.806	0.095	0.735	0.832	0.772
Ours (US-CoSOD-ImgNet150)	0.116	0.546	0.743	0.672	0.070	0.845	0.886	0.840	0.076	0.779	0.861	0.801
Ours (US-CoSOD-ImgNet450)	0.127	0.543	0.726	0.666	0.071	0.844	0.884	0.842	0.079	0.775	0.854	0.800

where f_S is the student model parameterized by θ_S , $F_\beta(p, q)$ is the F_β -metric computed between two segmentation maps p and q , and S_l is the labeled set. After training, CEN provides a reliability estimate of model prediction solely based on the image contents (Fig. 4 in supplement).

4. Experimental Results

4.1. Setup

Datasets and evaluation metrics: We used labeled data from COCO9213 [55], a subset of the COCO dataset [41] that contains 9,213 images selected from 65 groups, to train our semi-supervised model. We additionally constructed a dataset of 150K unlabeled images by selecting 150 images per class from the training subset of ImageNet-1K [10, 31] to train our unsupervised and semi-supervised models. We evaluate our methods on three popular CoSOD benchmarks: CoCA [71], Cosal2015 [65] and CoSOD3k [14]. CoCA and CoSOD3k are challenging real-world co-saliency evaluation datasets, containing multiple co-salient objects in some images, large appearance and scale variations, and complex backgrounds. Cosal2015 is a widely used dataset for CoSOD evaluation. Our evaluation metrics include the Mean Absolute Error (MAE↓) [7], maximum F-measure ($F_\beta^{max} \uparrow$) [1], maximum E-measure ($E_\phi^{max} \uparrow$) [12], and S-measure ($S_\alpha \uparrow$) [11].

Implementation details: We used DCFM [62] as the backbone model for all experiments. While the student network in SS-CoSOD consists of the Democratic Prototype Generation Module (DPG) and the Self-Contrastive Learning Module (SCL) from DCFM, we removed SCL in the teacher network because this network is not updated via backpropagation [62]. We used the Adam optimizer for training. The total training time is 5 hours for US-CoSOD and 8 hours for SS-CoSOD using ImageNet-1K. The inference time is 84.4 fps. More details in the supplementary.

4.2. Quantitative evaluation

In Table 1, we quantitatively evaluate the predictions obtained using seven different unsupervised baseline models: DINO self-attention mask (DI) [3], the most frequently co-occurring semantic unlabeled category mask

from STEGO (ST) [18], the pseudo co-saliency mask (DI+ST), predictions from TokenCut [56], DVFDVD [2], SegSwap [50], and predicted masks from our US-CoSOD model. We evaluated SegSwap [50] (that only predicts pairwise co-segmentations) by averaging the predicted co-segmentations over all image pairs in an image group i.e. between one image and every other image in the group. Table 1 shows that our US-CoSOD model trained on the 150K images from ImageNet (150 images per class) achieves the best performance. However, increasing the number of images to 450 per class slightly reduced the model accuracy. This could be attributed to the fact that adding more unlabeled images to the training set may lead to erroneous training due to the noisy pseudo ground truth masks generated by DI+ST (using which US-CoSOD is trained).

In Table 2, we compare the performance of different versions of our SS-CoSOD models using different proportions of labeled data. The “US-CoSOD” prefix indicates that the model is initialized with the pre-trained US-CoSOD model. “DJ” indicates that the model is trained with labeled and unlabeled images passed ‘disjointly’ to the student model without any cross-region interaction between the two sets. The SS-CoSOD model is trained semi-supervisedly using only images from the COCO9213 dataset. Finally, the “SS-CoSOD with ImgNet” version utilizes the extra 150 images per 1K ImageNet classes during semi-supervised training. In order to include the unlabeled ImageNet images in our semi-supervised setting, we use CEN to infer the confidence weights for the unlabeled images and then include the samples with predicted score ≥ 0.9 in the labeled set and the remaining samples in the unlabeled set.

In Table 2, we observe a consistent improvement in SS-CoSOD performance when the reliability scores from CEN are used to modulate the unsupervised loss. Joint labeled-unlabeled student training further leads to a small but consistent improvement. We further observe a significant improvement in performance when the ImageNet-1K dataset is used for semi-supervised learning, e.g. on the CoCA dataset, we obtain a reduction of 20% for MAE and gains of 4.79% for maximum F-measure, 2.64% for maximum E-measure, and 2.76% for S-measure compared to SS-CoSOD

Table 2. Performance comparison of the different versions of our unsupervised and semi-supervised models. In column 1, we indicate the fraction of labeled data for training, followed by the actual number of images. See supplementary for the results with 1/8 labeled data.

Labeled data	Method	CoCA				Cosal2015				CoSOD3k			
		MAE↓	$F_{\beta}^{max} \uparrow$	$E_{\phi}^{max} \uparrow$	$S_{\alpha} \uparrow$	MAE↓	$F_{\beta}^{max} \uparrow$	$E_{\phi}^{max} \uparrow$	$S_{\alpha} \uparrow$	MAE↓	$F_{\beta}^{max} \uparrow$	$E_{\phi}^{max} \uparrow$	$S_{\alpha} \uparrow$
1/16 (576)	US-CoSOD	0.108	0.557	0.754	0.683	0.068	0.854	0.888	0.846	0.076	0.783	0.857	0.801
	SS-CoSOD-DJ (w/o CEN)	0.107	0.485	0.728	0.635	0.094	0.771	0.834	0.771	0.089	0.709	0.817	0.742
	SS-CoSOD-DJ (w/ CEN)	0.115	0.488	0.730	0.639	0.086	0.782	0.847	0.787	0.086	0.717	0.828	0.755
	SS-CoSOD	0.113	0.492	0.733	0.641	0.085	0.788	0.850	0.792	0.084	0.721	0.830	0.758
	US-CoSOD+SS-CoSOD	0.111	0.554	0.751	0.681	0.066	0.855	0.890	0.849	0.075	0.783	0.858	0.803
	SS-CoSOD with ImgNet	0.098	0.562	0.757	0.684	0.072	0.837	0.880	0.828	0.068	0.784	0.865	0.800
1/4 (2303)	US-CoSOD	0.109	0.569	0.758	0.685	0.069	0.855	0.888	0.844	0.077	0.783	0.854	0.797
	SS-CoSOD-DJ (w/o CEN)	0.097	0.552	0.763	0.678	0.076	0.828	0.874	0.818	0.075	0.776	0.859	0.790
	SS-CoSOD-DJ (w/ CEN)	0.096	0.560	0.764	0.685	0.069	0.839	0.885	0.831	0.069	0.784	0.867	0.802
	SS-CoSOD	0.097	0.562	0.765	0.686	0.068	0.841	0.886	0.833	0.068	0.785	0.868	0.803
	US-CoSOD+SS-CoSOD	0.107	0.566	0.757	0.686	0.066	0.858	0.891	0.848	0.073	0.787	0.859	0.803
	SS-CoSOD with ImgNet	0.091	0.581	0.772	0.698	0.066	0.851	0.891	0.841	0.064	0.799	0.875	0.812
1/2 (4607)	US-CoSOD	0.105	0.569	0.760	0.688	0.068	0.856	0.889	0.843	0.074	0.793	0.862	0.804
	SS-CoSOD-DJ (w/o CEN)	0.092	0.572	0.771	0.694	0.068	0.846	0.885	0.834	0.071	0.791	0.865	0.802
	SS-CoSOD-DJ (w/ CEN)	0.090	0.578	0.772	0.699	0.062	0.851	0.892	0.843	0.067	0.795	0.870	0.810
	SS-CoSOD	0.088	0.582	0.773	0.700	0.062	0.854	0.892	0.843	0.066	0.797	0.872	0.809
	US-CoSOD+SS-CoSOD	0.110	0.563	0.755	0.686	0.064	0.858	0.894	0.850	0.072	0.794	0.866	0.810
	SS-CoSOD with ImgNet	0.088	0.590	0.775	0.705	0.062	0.861	0.896	0.850	0.063	0.804	0.876	0.817
Full (9213)	US-CoSOD	0.102	0.573	0.764	0.692	0.068	0.860	0.890	0.845	0.077	0.791	0.856	0.799
	SS-CoSOD with ImgNet	0.091	0.591	0.778	0.707	0.061	0.865	0.901	0.852	0.062	0.809	0.882	0.821

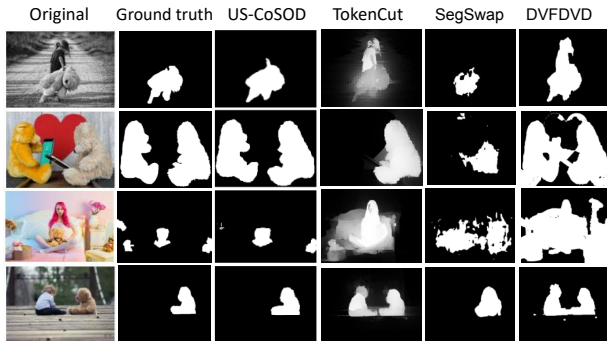


Figure 4. Qualitative comparisons of our US-CoSOD model with other baselines on the *teddy bear* image group from CoCA. US-CoSOD produces the most accurate segmentation masks.

trained using the COCO9213 dataset on the 1/2 data split.

Comparison with existing CoSOD models: In Table 3, we compare the performance of existing unsupervised and semi-supervised CoSOD models with our model using the F_{β}^{max} -measure and the S_{α} -measure metrics. Both of our unsupervised and semi-supervised models outperform the corresponding state-of-the-art models by a significant margin *e.g.*, on the Cosal2015 dataset, our US-CoSOD model has an 8.8% F-measure gain over DVFDVD [2], an unsupervised SOTA co-segmentation model and our SS-CoSOD model (using ImageNet-1K) has an 11.81% F-measure gain over FASS [72], a semi-supervised CoSOD model.

Ablation studies:

Variation model: For certain image groups (*e.g. key, frisbee, etc.*), the co-saliency of the common objects can be lesser than that of other bigger objects solely due to the lesser area, which could impact performance. To test this hypothe-

sis, we investigate a variant model that normalizes the overlap score by STEGO mask area as: $O_i^j = \frac{Ar(SM_i^j \cap DM_i)}{Ar(SM_i^j)}$ (see Algorithm 1). US-CoSOD (without area normalization) has more accurate predictions compared to this variant on all test sets, *e.g.* US-CoSOD achieves F-measures 0.461, 0.760, and 0.684 against the variant model’s 0.410, 0.613, and 0.579 on CoCA, Cosal2015, and CoSOD3k respectively. See supplementary for more details.

Performance on challenging categories: Our US-CoSOD outperforms the pre-trained DINO and DINO+STEGO models by a significant margin on challenging categories (categories over which DINO scored lesser than the average DINO F-measure score over the test dataset). For instance, the average F-measures of DINO, DINO+STEGO, and US-CoSOD are 0.598, 0.654, and 0.738 on the Cosal2015 dataset respectively. See supplementary for more details.

CEN backbone: We compared the confidence estimation error (Mean Squared Error, MSE) of CEN using different backbone networks on the unlabeled set. DINO (ResNet50) yielded the least MSE across different data splits *e.g.* MSE using ResNet50, MobileNetV2, ViTB, and ViTS are 0.166, 0.171, 0.176, and 0.177 respectively on the 1/4 labeled data split. See supplementary for details. We attribute the lower accuracy of MobileNetV2 to its lower feature representation power. Also the transformer models DINO (ViTB, ViTS) fail to outperform the convolutional models (*e.g.* ResNet50) due to the lesser training data (in different data splits).

4.3. Qualitative evaluation

In Fig. 4, we qualitatively compare the CoSOD predictions from TokenCut [56], SegSwap [50], and DVFDVD [2] with our US-CoSOD on the *teddy bear* image group from




Figure 5. Qualitative comparisons of our model with different baselines on three image groups selected each from the CoCA, CoSOD3k, and Cosal2015 datasets. Our SS-CoSOD model (all labels) produces the most accurate segmentation mask compared to the other baselines.

CoCA. We see that US-CoSOD more accurately detects the teddy bear in the four images compared to the baselines.

In Fig. 5, we qualitatively compare the CoSOD predictions from different baselines with SS-CoSOD on three image groups, each from the CoCA, CoSOD3k, and Cosal2015 datasets. We see that our US-CoSOD generates reasonable masks and our SS-CoSOD approaches further improve these predictions while performing better than the fully supervised DCFM trained with limited labels. While DCFM (all labels) produces incomplete segmentations (columns 5, 9, 11, and 12) and overestimates the co-saliency (in columns 2, 3, 5, and 8) in certain image regions, our SS-CoSOD (all labels) model predictions suffer fewer inaccuracies, producing more accurate CoSOD masks.

Table 3. Comparison of our models with existing unsupervised and semi-supervised models for CoSOD on Cosal2015.

Method	Type	Label %	$F_{\beta}^{max} \uparrow$	$S_{\alpha} \uparrow$
CSSCF [24] (TMM 2016)	Unsup	-	0.682	0.671
CoDW [65] (IJCV 2016)	Unsup	-	0.705	0.647
UCCDGO [22] (ECCV 2018)	Unsup	-	0.758	0.751
DVFDVD [2] (ECCVW 2022)	Unsup	-	0.777	0.809
SegSwap [50] (CVPRW 2022)	Unsup	-	0.618	0.632
Ours (US-CoSOD)	Unsup	-	0.845	0.840
FASS [72] (ACMM 2018)	Semi-sup	50%	0.770	-
Ours (SS-CoSOD w/ ImgNet)	Semi-sup	25%	0.851	0.841
Ours (SS-CoSOD w/ ImgNet)	Semi-sup	50%	0.861	0.850

Fig. 6 shows that while multiple categories can co-exist in an image group (e.g. *flowers* and *butterflies*), our DI+ST model extracts the more co-salient butterflies, guided by the DINO SA maps. Also, our US-CoSOD model trained using

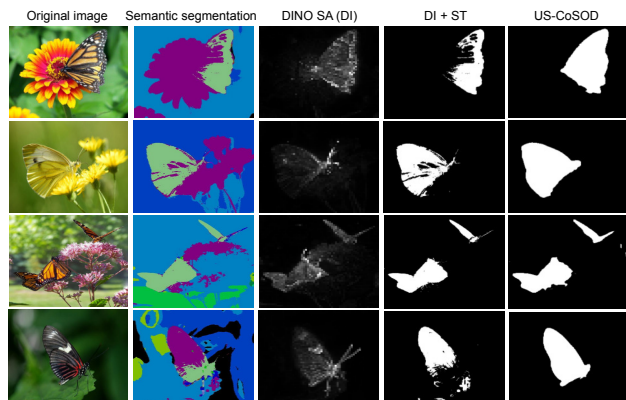


Figure 6. DINO self-attention maps and US-CoSOD predictions using DI+ST for training. US-CoSOD produces the best results.

the pseudo co-saliency masks from DI+ST improves the co-saliency masks (better structural consistency) from DI+ST.

5. Conclusion

We presented a novel unsupervised approach for CoSOD based on the frequency statistics of semantic segmentations, which forms a strong pre-training initialization for a semi-supervised CoSOD model. Our semi-supervised model employs a student-teacher approach with an effective confidence estimation module. We demonstrate that both our unsupervised and semi-supervised CoSOD models can significantly improve prediction performance over a fully-supervised model trained with limited labeled data. As future work, we aim to improve our unsupervised model by avoiding the use of any off-the-shelf component.

References

- [1] Radhakrishna Achanta, Sheila Hemami, Francisco Estrada, and Sabine Susstrunk. Frequency-tuned salient region detection. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1597–1604. IEEE, 2009. 5, 6
- [2] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *ECCVW What is Motion For?*, 2022. 2, 6, 7, 8
- [3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. 2, 3, 4, 6
- [4] Souradeep Chakraborty and Pabitra Mitra. A site entropy rate and degree centrality based algorithm for image co-segmentation. *Journal of Visual Communication and Image Representation*, 33:20–30, 2015. 2
- [5] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2613–2622, 2021. 2
- [6] Zuyao Chen, Qianqian Xu, Runmin Cong, and Qingming Huang. Global context-aware progressive aggregation network for salient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 10599–10606, 2020. 1
- [7] Ming-Ming Cheng, Jonathan Warrell, Wen-Yan Lin, Shuai Zheng, Vibhav Vineet, and Nigel Crook. Efficient salient region detection with soft image abstraction. In *Proceedings of the IEEE International Conference on Computer vision*, pages 1529–1536, 2013. 6
- [8] Jang Hyun Cho, Utkarsh Mall, Kavita Bala, and Bharath Hariharan. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16794–16804, 2021. 2
- [9] Runmin Cong, Jianjun Lei, Huazhu Fu, Ming-Ming Cheng, Weisi Lin, and Qingming Huang. Review of visual saliency detection with comprehensive information. *IEEE Transactions on circuits and Systems for Video Technology*, 29(10):2941–2959, 2018. 2
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2, 6
- [11] Deng-Ping Fan, Ming-Ming Cheng, Yun Liu, Tao Li, and Ali Borji. Structure-measure: A new way to evaluate foreground maps. In *Proceedings of the IEEE international conference on computer vision*, pages 4548–4557, 2017. 6
- [12] Deng-Ping Fan, Cheng Gong, Yang Cao, Bo Ren, Ming-Ming Cheng, and Ali Borji. Enhanced-alignment measure for binary foreground map evaluation. *arXiv preprint arXiv:1805.10421*, 2018. 6
- [13] Deng-Ping Fan, Tengpeng Li, Zheng Lin, Ge-Peng Ji, Dingwen Zhang, Ming-Ming Cheng, Huazhu Fu, and Jianbing Shen. Re-thinking co-salient object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 1, 2
- [14] Deng-Ping Fan, Zheng Lin, Ge-Peng Ji, Dingwen Zhang, Huazhu Fu, and Ming-Ming Cheng. Taking a deeper look at co-salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2919–2929, 2020. 2, 6
- [15] Qi Fan, Deng-Ping Fan, Huazhu Fu, Chi-Keung Tang, Ling Shao, and Yu-Wing Tai. Group collaborative learning for co-salient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12288–12298, 2021. 1, 2
- [16] Guangshuai Gao, Wenting Zhao, Qingjie Liu, and Yunhong Wang. Co-saliency detection with co-attention fully convolutional network. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(3):877–889, 2020. 1
- [17] Yanliang Ge, Qiao Zhang, Tian-Zhu Xiang, Cong Zhang, and Hongbo Bi. Tcnet: Co-salient object detection via parallel interaction of transformers and cnns. *IEEE Transactions on Circuits and Systems for Video Technology*, 2022. 2
- [18] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T Freeman. Unsupervised semantic segmentation by distilling feature correspondences. *arXiv preprint arXiv:2203.08414*, 2022. 2, 3, 4, 6
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [20] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 3
- [21] Kuang-Jui Hsu, Yen-Yu Lin, Yung-Yu Chuang, et al. Co-attention cnns for unsupervised object co-segmentation. In *IJCAI*, volume 1, page 2, 2018. 1, 2
- [22] Kuang-Jui Hsu, Chung-Chi Tsai, Yen-Yu Lin, Xiaoning Qian, and Yung-Yu Chuang. Unsupervised cnn-based co-saliency detection with graphical optimization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 485–501, 2018. 1, 2, 8
- [23] Rongyao Hu, Zhenyun Deng, and Xiaofeng Zhu. Multi-scale graph fusion for co-saliency detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7789–7796, 2021. 2
- [24] Koteswar Rao Jerripothula, Jianfei Cai, and Junsong Yuan. Image co-segmentation via saliency co-fusion. *IEEE Transactions on Multimedia*, 18(9):1896–1909, 2016. 2, 8
- [25] Xu Ji, Joao F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9865–9874, 2019. 2
- [26] Bo Jiang, Xingyue Jiang, Jin Tang, and Bin Luo. Co-saliency detection via a general optimization model and adaptive graph learning. *IEEE Transactions on Multimedia*, 23:3193–3202, 2020. 2

- [27] Bo Jiang, Xingyue Jiang, Jin Tang, Bin Luo, and Shilei Huang. Multiple graph convolutional networks for co-saliency detection. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, pages 332–337. IEEE, 2019. 2
- [28] Bo Jiang, Xingyue Jiang, Ajian Zhou, Jin Tang, and Bin Luo. A unified multiple graph learning and convolutional network model for co-saliency estimation. In *proceedings of the 27th ACM International Conference on Multimedia*, pages 1375–1382, 2019. 2
- [29] Wen-Da Jin, Jun Xu, Ming-Ming Cheng, Yi Zhang, and Wei Guo. Icnnet: Intra-saliency correlation network for co-saliency detection. *Advances in Neural Information Processing Systems*, 33:18749–18759, 2020. 2
- [30] Zhanhan Ke, Di Qiu, Kaican Li, Qiong Yan, and Rynson WH Lau. Guided collaborative training for pixel-wise semi-supervised learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pages 429–445. Springer, 2020. 2
- [31] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. 2, 6
- [32] Donghyeon Kwon and Suha Kwak. Semi-supervised semantic segmentation with error localization network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9957–9967, 2022. 2
- [33] Xin Lai, Zhuotao Tian, Li Jiang, Shu Liu, Hengshuang Zhao, Liwei Wang, and Jiaya Jia. Semi-supervised semantic segmentation with directional context-aware consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1205–1214, 2021. 2
- [34] Hieu Le, Chen-Ping Yu, Gregory Zelinsky, and Dimitris Samaras. Co-localization with category-consistent features and geodesic distance propagation. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1103–1112, 2017. 2
- [35] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4071–4080, 2021. 2
- [36] Aixuan Li, Jing Zhang, Yunqiu Lv, Bowen Liu, Tong Zhang, and Yuchao Dai. Uncertainty-aware joint salient object and camouflaged object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10071–10081, 2021. 1
- [37] Jia Li, Jinming Su, Changqun Xia, Mingcan Ma, and Yonghong Tian. Salient object detection with purificatory mechanism and structural similarity loss. *IEEE Transactions on Image Processing*, 30:6855–6868, 2021. 1
- [38] Long Li, Junwei Han, Ni Zhang, Nian Liu, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, and Fahad Shabbaz Khan. Discriminative co-saliency and background mining transformer for co-salient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7247–7256, 2023. 2
- [39] Lina Li, Zhi Liu, and Jian Zhang. Unsupervised image co-segmentation via guidance of simple images. *Neurocomputing*, 275:1650–1661, 2018. 1, 2
- [40] Yunfan Li, Peng Hu, Zitao Liu, Dezhong Peng, Joey Tianyi Zhou, and Xi Peng. Contrastive clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 8547–8555, 2021. 2
- [41] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [42] Jiawei Liu, Jing Zhang, and Nick Barnes. Semi-supervised salient object detection with effective confidence estimation. *arXiv preprint arXiv:2112.14019*, 2021. 2, 3
- [43] Yun Liu, Xin-Yu Zhang, Jia-Wang Bian, Le Zhang, and Ming-Ming Cheng. Samnet: Stereoscopically attentive multi-scale network for lightweight salient object detection. *IEEE Transactions on Image Processing*, 30:3804–3814, 2021. 1
- [44] Yunqiu Lv, Bowen Liu, Jing Zhang, Yuchao Dai, Aixuan Li, and Tong Zhang. Semi-supervised active salient object detection. *Pattern Recognition*, 123:108364, 2022. 3
- [45] Robert Mendel, Luis Antonio De Souza, David Rauber, Joao Paulo Papa, and Christoph Palm. Semi-supervised segmentation based on error-correcting supervision. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pages 141–157. Springer, 2020. 2
- [46] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high-and low-level consistency. *IEEE transactions on pattern analysis and machine intelligence*, 43(4):1369–1379, 2019. 2
- [47] Yongri Piao, Jian Wang, Miao Zhang, and Huchuan Lu. Mfnet: Multi-filter directive network for weakly supervised salient object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4136–4145, 2021. 1
- [48] Xiaoliang Qian, Yinfeng Zeng, Wei Wang, and Qiuwen Zhang. Co-saliency detection guided by group weakly supervised learning. *IEEE Transactions on Multimedia*, 2022. 3
- [49] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. Basnet: Boundary-aware salient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7479–7489, 2019. 4
- [50] Xi Shen, Alexei A Efros, Armand Joulin, and Mathieu Aubry. Learning co-segmentation by segment swapping for retrieval and discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5082–5092, 2022. 6, 7, 8
- [51] Yukun Su, Jingliang Deng, Ruizhou Sun, Guosheng Lin, Hanjing Su, and Qingyao Wu. A unified transformer framework for group-based segmentation: Co-segmentation, co-saliency detection and video salient object detection. *IEEE Transactions on Multimedia*, 2023. 2

- [52] Lv Tang, Bo Li, Yijie Zhong, Shouhong Ding, and Mofei Song. Disentangled high quality salient object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3580–3590, 2021. [1](#)
- [53] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. [4](#)
- [54] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Scan: Learning to classify images without labels. In *European conference on computer vision*, pages 268–285. Springer, 2020. [2](#)
- [55] Chong Wang, Zheng-Jun Zha, Dong Liu, and Hongtao Xie. Robust deep co-saliency detection with group semantic. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8917–8924, 2019. [6](#)
- [56] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14543–14553, 2022. [2](#), [6](#), [7](#)
- [57] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L. Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *Conference on Computer Vision and Pattern Recognition*, 2022. [3](#)
- [58] Yuchao Wang, Haochen Wang, Yujun Shen, Jingjing Fei, Wei Li, Guoqiang Jin, Liwei Wu, Rui Zhao, and Xinyi Le. Semi-supervised semantic segmentation using unreliable pseudo-labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4248–4257, 2022. [2](#), [4](#)
- [59] Zhengxiang Wang and Rujie Liu. Semi-supervised learning for large scale image cosegmentation. In *2013 IEEE International Conference on Computer Vision*, pages 393–400, 2013. [3](#)
- [60] Lina Wei, Shanshan Zhao, Omar El Farouk Bourahla, Xi Li, Fei Wu, and Yueting Zhuang. Deep group-wise fully convolutional network for co-saliency detection with graph propagation. *IEEE Transactions on Image Processing*, 28(10):5052–5063, 2019. [2](#)
- [61] Zhaoyuan Yin, Pichao Wang, Fan Wang, Xianzhe Xu, Hanling Zhang, Hao Li, and Rong Jin. Transfgu: a top-down approach to fine-grained unsupervised semantic segmentation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIX*, pages 73–89. Springer, 2022. [2](#), [3](#)
- [62] Siyue Yu, Jimin Xiao, Bingfeng Zhang, and Eng Gee Lim. Democracy does matter: Comprehensive feature mining for co-salient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 979–988, 2022. [1](#), [2](#), [4](#), [5](#), [6](#)
- [63] Siyue Yu, Bingfeng Zhang, Jimin Xiao, and Eng Gee Lim. Structure-consistent weakly supervised salient object detection with local saliency coherence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3234–3242, 2021. [1](#)
- [64] Dingwen Zhang, Huazhu Fu, Junwei Han, Ali Borji, and Xuelong Li. A review of co-saliency detection algorithms: fundamentals, applications, and challenges. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 9(4):1–31, 2018. [2](#)
- [65] Dingwen Zhang, Junwei Han, Chao Li, Jingdong Wang, and Xuelong Li. Detection of co-salient objects by looking deep and wide. *International Journal of Computer Vision*, 120(2):215–232, 2016. [1](#), [2](#), [6](#), [8](#)
- [66] Dingwen Zhang, Haibin Tian, and Jungong Han. Few-cost salient object detection with adversarial-paced learning. *Advances in Neural Information Processing Systems*, 33:12236–12247, 2020. [3](#)
- [67] Kaihua Zhang, Mingliang Dong, Bo Liu, Xiao-Tong Yuan, and Qingshan Liu. Deepacg: Co-saliency detection via semantic-aware contrast gromov-wasserstein distance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13703–13712, 2021. [1](#), [2](#)
- [68] Kaihua Zhang, Tengteng Li, Shiwen Shen, Bo Liu, Jin Chen, and Qingshan Liu. Adaptive graph convolutional network with attention graph clustering for co-saliency detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9050–9059, 2020. [2](#)
- [69] Ni Zhang, Junwei Han, Nian Liu, and Ling Shao. Summarize and search: Learning consensus-aware dynamic convolution for co-saliency detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4167–4176, 2021. [2](#)
- [70] Qijian Zhang, Runmin Cong, Junhui Hou, Chongyi Li, and Yao Zhao. Coadnet: Collaborative aggregation-and-distribution networks for co-salient object detection. *Advances in neural information processing systems*, 33:6959–6970, 2020. [2](#), [4](#)
- [71] Zhao Zhang, Wenda Jin, Jun Xu, and Ming-Ming Cheng. Gradient-induced co-saliency detection. In *European Conference on Computer Vision*, pages 455–472. Springer, 2020. [1](#), [2](#), [6](#)
- [72] Xiaoju Zheng, Zheng-Jun Zha, and Liansheng Zhuang. A feature-adaptive semi-supervised framework for co-saliency detection. In *Proceedings of the 26th ACM International Conference on Multimedia, MM '18*, page 959–966, New York, NY, USA, 2018. Association for Computing Machinery. [1](#), [3](#), [7](#), [8](#)
- [73] Ziyue Zhu, Zhao Zhang, Zheng Lin, Xing Sun, and Ming-Ming Cheng. Co-salient object detection with co-representation purification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. [2](#)