

Can You Spot the Chameleon?

Adversarially Camouflaging Images from Co-Salient Object Detection

Ruijun Gao^{1*}, Qing Guo^{1,6*}, Felix Juefei-Xu², Hongkai Yu³, Huazhu Fu⁴, Wei Feng^{1†},
Yang Liu^{5,6}, Song Wang⁷

¹College of Intelligence and Computing, Tianjin University, China, ²Alibaba Group, USA,
³Cleveland State University, USA, ⁴IHPC, A*STAR, Singapore, ⁵Zhejiang Sci-Tech University, China,
⁶Nanyang Technology University, Singapore, ⁷University of South Carolina, USA

Abstract

Co-salient object detection (CoSOD) has recently achieved significant progress and played a key role in retrieval-related tasks. However, it inevitably poses an entirely new safety and security issue, i.e., highly personal and sensitive content can potentially be extracting by powerful CoSOD methods. In this paper, we address this problem from the perspective of adversarial attacks and identify a novel task: **adversarial co-saliency attack**. Specially, given an image selected from a group of images containing some common and salient objects, we aim to generate an adversarial version that can mislead CoSOD methods to predict incorrect co-salient regions. Note that, compared with general white-box adversarial attacks for classification, this new task faces two additional challenges: (1) low success rate due to the diverse appearance of images in the group; (2) low transferability across CoSOD methods due to the considerable difference between CoSOD pipelines. To address these challenges, we propose the very first black-box joint adversarial exposure and noise attack (Jadena), where we jointly and locally tune the exposure and additive perturbations of the image according to a newly designed high-feature-level contrast-sensitive loss function. Our method, without any information on the state-of-the-art CoSOD methods, leads to significant performance degradation on various co-saliency detection datasets and makes the co-salient objects undetectable. This can have strong practical benefits in properly securing the large number of personal photos currently shared on the Internet. Moreover, our method is potential to be utilized as a metric for evaluating the robustness of CoSOD methods.

*Ruijun Gao and Qing Guo are co-first authors and contribute equally to this work.

†Wei Feng is the corresponding author: wfeng@ieee.org

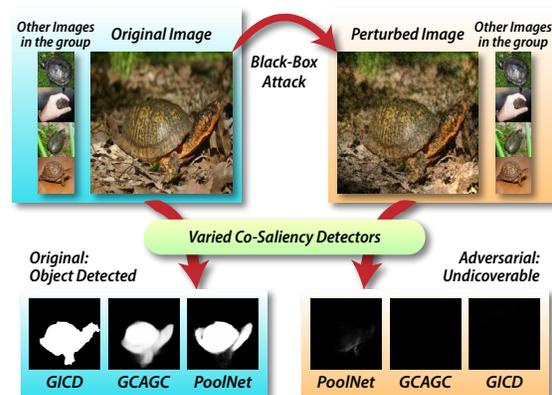


Figure 1. Overview of the novel problem, **adversarial co-saliency attack**, and our solution. We expect the perturbed image to be undiscoverable in an even dynamically growing group of images across multiple CoSOD methods (e.g., GICD [52], GCAGC [50], PoolNet [31]), which is much more challenging and practical in real-world scenarios. Note that our attack is black-box and can be performed without references provided in the group.

1. Introduction

Co-saliency is typically defined as the common and salient (usually in foreground) visual stimulus found across a given group of images. Co-salient object detection (CoSOD) therefore aims at detecting and highlighting the common and salient foreground region (object) in the given group of image [47]. Different from the traditional single-image saliency detection problem, the key to solving the co-saliency problem is to discover the correspondence (based on various cues) of the common and similar salient regions across multiple images in a group. With unknown semantic categories of the co-salient objects, the co-saliency task is a rather challenging one, and a good co-saliency detection algorithm should consider both the intra-image saliency cues and the inter-image common saliency cues [14].

At the moment, co-saliency detection is still an emerging research topic, with many recently proposed methods to

solve this challenging problem. We will discuss several of the current CoSOD algorithms, which range from low-level features to high-level semantic features, as well as deep learning-based models, in Section 2. Co-saliency detection plays a key role in many practical applications of computer vision and multimedia, including object co-segmentation [54], foreground discovery in video sequences [5], weakly supervised localization [48], image retrieval [14], multi-camera surveillance [32], etc.

However, the extensive and powerful CoSOD methods inevitably pose a new safety and security problem, *i.e.*, high-profile and personal-sensitive content may be subject to extraction and discovery by these methods. In order to prevent content-sensitive images from being discovered by co-saliency detection, in this work, we address the problem from the perspective of adversarial attacks and identify a novel task, *i.e.*, **adversarial co-saliency attack**. Specially, given an image selected from a group containing some common and salient objects, the aim of the task is to generate an adversarial version of the image that can mislead CoSOD methods to predict incorrect co-salient regions, thus evading co-saliency detection, and making the images undiscoverable by the co-saliency detection algorithms. Our method is able to conceal personal content containing sensitive information from the CoSOD methods. Moreover, our method can potentially be used as a metric to evaluate the robustness of the CoSOD methods and can be coupled with adversarial training to improve the performance as well.

It is worth noting that, compared with general white-box adversarial attacks for classification, this new task faces two additional challenges: (1) low success rate due to the diverse appearance of images in the group of images; (2) low transferability across CoSOD methods due to the considerable difference between CoSOD pipelines. To overcome these challenges, we propose the very first black-box joint adversarial exposure and noise attack (Jadena), where we jointly and locally tune the exposure and additive perturbations of the image according to a newly designed high-feature-level contrast-sensitive loss function. Our method, without needing any information on the state-of-the-art CoSOD methods, leads to significant performance degradation on various co-saliency detection datasets and makes the co-salient objects undetectable, as shown in Fig. 1. This has strong practical value nowadays where large-scale personal multimedia contents are shared on the public domain Internet and should be properly and securely preserved and protected from malicious extraction. We have released our code in this GitHub repo: <https://github.com/tsingqquo/jadena>.

2. Related work

Co-saliency detection. As an important subarea of saliency detection [1, 7, 28, 42, 44], the goal of co-saliency detec-

tion problem is to detect common salient objects in multiple images [12, 25, 27, 45, 50, 52]. Co-saliency detection considers the saliency cues within images and also the common cues among images [14, 17]. Understanding correspondences of co-salient objects among multiple images is the key to co-saliency detection, which can be addressed by optimization-based methods [3, 29], machine learning-based models [8, 49], and deep neural networks [45, 50, 52]. In addition, data from other modalities, such as depth images, can be fused to improve the co-saliency detection on color-based images [9]. Co-saliency detection has wide applications in computer vision, including co-segmentation [23] and co-localization [37], which segments and localizes similar objects in a group of images.

Adversarial attack. With the booming of deep learning, the model safety has attracted more and more attention. Therefore many works have been proposed for adversarial attacks against deep neural networks. Several classical methods, such as the fast gradient signed method (FGSM) [18, 21], momentum iterative fast gradient sign method (MI-FGSM) [10], basic iterative method (BIM) [26], C&W method [4], show promising results. Most current attacks learn additive noise under a white-box setting, in which the model parameters are given. For example, learnable noise will be added to a clean image as a perturbation to decrease the task performance. Some other studies have also try the adversarial natural corruptions including motion blur [19, 20], vignetting [38], rain streaks [46], face morphing [43], and haze [16], and the watermark [24]. The proposed method is different from these methods because it is a non-additive noise attack based on a black-box setup.

Adversarial attack for saliency detection. There is limited research on adversarial attacks and defenses for saliency detection. Tran *et al.* [39] have shown that adding adversarial attacks, *e.g.*, using FGSM [18], can result in a large performance drop for saliency detection, so they proposed a method to clean data affected by an adversarial attack. Different from attacks in image space, Che *et al.* [6] have proposed a sparser adversarial perturbation in feature space against deep saliency models, which needs some model information. To the best of our knowledge, this paper is the first work on the adversarial attacks for co-saliency detection.

3. Methodology

3.1. Problem formulation

Co-salient object detection. Consider a dynamic group of images $\mathcal{I} = \{\mathbf{I}_i\}$ containing $|\mathcal{I}|$ images that have common and salient objects. ‘Dynamic’ means the group of images can be changed or extended, since personal images are updated every day in real-world applications. A CoSOD

method denoted as $D(\cdot)$ takes the group of images as input and outputs the common and salient regions in all images

$$\mathbf{S} = \{\mathbf{S}_i\}_{i=1}^{|\mathcal{I}|} = D(\mathcal{I}), \quad (1)$$

where $\mathbf{S}_i = D(\mathcal{I})[i]$ is a binary map of the i -th image to represent its common and salient regions with ‘1’ and others with ‘0’. Currently, most CoSOD methods have different pipelines and can achieve effective co-saliency detection.

Task definition. In this paper, we focus on an entirely new problem: given an image \mathbf{I}_k from the dynamic group of images \mathcal{I} with $k \in [1, \dots, |\mathcal{I}|]$, we aim to make it undiscoverable to CoSOD methods. To solve this problem, we formulate it as an adversarial attack, applying adversarial perturbations to image \mathbf{I}_k to fool a CoSOD method $D(\cdot)$ into predicting an incorrect salient map, *i.e.*, \mathbf{S}_k . Specifically, we can add perturbations to the original image \mathbf{I}_k by $\mathbf{I}_k^a = P_\theta(\mathbf{I}_k)$, where \mathbf{I}_k^a is the adversarial example aiming to fool the CoSOD method and $P_\theta(\mathbf{I}_k)$ is an image transformation (*e.g.*, $P_\theta(\mathbf{I}_k) = \mathbf{I}_k + \theta$ for the popular additive-perturbation-based attack). The key to the adversarial attack is how to calculate the parameter θ .

3.2. Adversarial noise attack and challenges

We first consider a straightforward way to address the task with the popular additive-perturbation-based attack. To be specific, we define

$$\mathbf{I}_k^a = P_\theta(\mathbf{I}_k) = \mathbf{I}_k + \theta, \quad (2)$$

where θ is a pixel-wise perturbation of the same size as \mathbf{I}_k . Then, we can define the following objective function for calculating θ :

$$\begin{aligned} \arg \max_{\theta} J(D(\{\mathbf{I}_k + \theta\} \cup \{\mathbf{I}_i | i \neq k\})[k], \mathbf{S}_k), \\ \text{subject to } \|\theta\|_p \leq \epsilon, \end{aligned} \quad (3)$$

where \mathbf{S}_k is the CoSOD result of \mathbf{I}_k based on the original group of images and $J(\cdot)$ is a loss function measuring the distance between the predictive results and \mathbf{S}_k . The constraint function, *i.e.*, $\|\theta\|_p \leq \epsilon$, is used to control the distortion degree determined by θ .

We can optimize the above function with the standard optimizations for additive-perturbation-based attacks, *i.e.*, using one-step signed gradient-based, iterative-based, and momentum-based optimization methods. However, we find such a general adversarial noise attack cannot address the new task very well since it faces two *challenges*: (1) It is difficult for the white-box adversarial noise attack to reach a high success rate since the dynamic group of images \mathcal{I} makes the optimized θ ineffective when new images are added. (2) Because there are several diverse pipelines for CoSOD, the optimized θ has low transferability. That is, the optimized θ for one CoSOD method fails to fool another.

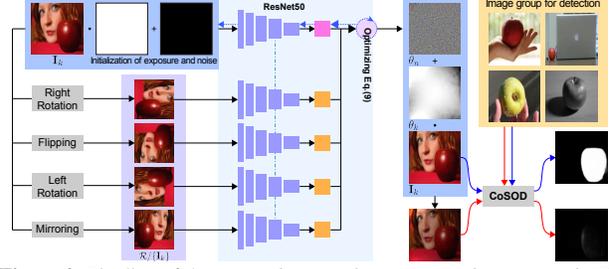


Figure 2. Pipeline of the joint adversarial exposure and noise attack. The clean image is augmented to generate references and the gradient back-propagates along the blue dashed lines.

3.3. Joint adversarial exposure and noise attack

In this section, we introduce a new black-box attack method that can address the two challenges faced by white-box adversarial noise attacks. In general, the saliency of an object is related to its contrast to other regions of the input image, while the contrast is usually affected by the image exposure. To achieve an effective attack, rather than only considering the noise, as in Sec. 3.2, we propose to jointly consider exposure and noise factors during the attack.

Formulation. We add the joint exposure and noise based perturbations to the image \mathbf{I}_k by

$$\mathbf{I}_k^a = P_\theta(\mathbf{I}_k) = \theta_e \mathbf{I}_k + \theta_n, \quad (4)$$

where $\theta = [\theta_e, \theta_n]$ and θ_e and θ_n denote the pixel-wise exposure and noise matrices, respectively. Note that the above exposure model has been widely used in various image enhancement methods [15, 51]. However, achieving an effective co-saliency attack is still challenging: (1) The exposure is smooth across the whole image domain and cannot be tuned pixel-wise, as this will result in a speckle-like appearance. (2) We still do not know how to optimize θ to achieve a black-box attack across different CoSOD methods.

For the first challenge, we construct an objective function to encourage the naturalness of the exposure. For the second, instead of optimizing θ based on a CoSOD method, we explore a full black-box approach, where we define the objective functions on the high-level features extracted from a pretrained classification model denoted as $\phi(\cdot)$.

Objective function for the naturalness of the adversarial exposure. We emphasize that exposure in the real world is usually smooth (*i.e.*, neighboring pixels tend to have very similar exposure) across the whole image. In contrast, the adversarial attacks require the exposure to be locally different for a high success rate. Besides, we want the exposure to modify the original input as little as possible. To meet these requirements, we propose to represent the exposure with the locally-variant multivariate polynomial model in

the logarithm domain:

$$\log(\theta_{e,p}) = \sum_{d=0}^D \sum_{l=0}^{D-d} a_{d,l} (x_p + u_p)^d (y_p + v_p)^l, \quad (5)$$

where x_p and y_p are the coordinates of the p -th pixel and u_p and v_p denote the variation offsets of the p -th pixel. $\{a_{d,l}\}$ denotes the parameters of the polynomial model and is concatenated as \mathbf{a} with D being the degree. Note that a polynomial model with fewer parameters $\{a_{d,l}\}$ leads to smoother exposure. With Eq. (5), we can represent the exposure as a function of \mathbf{a} and the offset map $\mathbf{U} = \{(u_p, v_p)\}$, *i.e.*, $\theta_e(\mathbf{a}, \mathbf{U})$. More specifically, to maintain the smoothness of exposure and make it capable of an adversarial attack, we define the following objective function:

$$J_{\text{smooth}}(\mathbf{a}, \mathbf{U}) = -\lambda_b \|\log(\theta_e(\mathbf{a}, \mathbf{U}))\|_2^2 - \lambda_s \|\nabla \mathbf{U}\|_2^2, \quad (6)$$

where the first term encourages the model to not perturb the original \mathbf{I}_k and the second term is a total variation regularization, which encourages smooth variation of the exposure.

Objective function against single saliency. We argue that the image \mathbf{I}_k should be perturbed by the exposure and noise to make its high-level features consistent across all image regions. As a result, the salient objects will not be invisible at a high level. To this end, we define the objective function against single saliency:

$$J_{\text{cons}}(\mathbf{a}, \mathbf{U}, \theta_n) = -\text{avg}(\{\text{std}(\phi_j(P_\theta(\mathbf{I}_k))) | j \in \mathcal{L}\}) \quad (7)$$

where $\phi_j(\cdot)$ is the j -th layer feature of the CNN $\phi(\cdot)$ and \mathcal{L} denotes the index set of employed layers. The function $\text{std}(\cdot)$ calculates the standard variation of each feature channel and outputs the average across all channels. The function $\text{avg}(\cdot)$ computes the average of all $\text{std}(\phi_j(P_\theta(\mathbf{I}_k)))$ across all required layers.

Objective function against co-saliency. Eq. (7) is used to make the object less salient in a single image and may be less effective for co-saliency detection since the non-salient object could be salient again in a group of images. We thus extend Eq. (7) to the feature level of a group of images, *i.e.*, \mathcal{R} with $P_\theta(\mathbf{I}_k) \in \mathcal{R}$, and have

$$J_{\text{co-cons}}(\mathbf{a}, \mathbf{U}, \theta_n) = -\text{avg}(\{\text{std}(\Phi_j(\mathcal{R})) | j \in \mathcal{L}\}), \quad (8)$$

with $\Phi_j(\mathcal{R}) = \text{spl}(\{\phi_j(\mathbf{I}_i) | \mathbf{I}_i \in \mathcal{R}\})$.

Note that we splice the j -th layer features of all images in \mathcal{R} with the function $\text{spl}(\cdot)$, in a channel-wise way. For example, if the function $\phi_j(\mathbf{I}_i)$ outputs a tensor with size $11 \times 11 \times 256$ for $|\mathcal{R}| = 4$, we get $\Phi_j(\mathcal{R})$ of size $44 \times 11 \times 256$. Intuitively, compared with Eq. (7), Eq. (8) makes the salient regions in \mathbf{I}_k become less salient across the group \mathcal{R} . We consider two ways of constructing \mathcal{R} .

One uses the transformation of \mathbf{I}_k , *e.g.*, rotation and up-down flip; the other employs other images. We will show that both can achieve effective co-saliency attack.

Optimization. With the above objective functions, we can calculate the adversarial exposure (*i.e.*, $\theta_e(\mathbf{a}, \mathbf{U})$) and noise (*i.e.*, θ_n) by optimizing

$$\begin{aligned} \arg \max_{\mathbf{a}, \mathbf{U}, \theta_n} J_{\text{co-cons/cons}}(\mathbf{a}, \mathbf{U}, \theta_n) + J_{\text{smooth}}(\mathbf{a}, \mathbf{U}), \\ \text{subject to } \|\theta_n\|_p \leq \epsilon. \end{aligned} \quad (9)$$

Compared with the objective function used by the adversarial noise attack, *i.e.*, Eq. (3), the new objective function has the following advantages: (1) It does not rely on CoSOD methods or the networks they use. This means that our method is a full black-box attack and could be very useful in practical applications. (2) Our method is unsupervised since it does not use the saliency map, *i.e.*, \mathbf{S}_k in Eq. (3), which also makes it more useful in the real world. (3) We use both exposure and noise for a more effective attack. We employ the sign gradient descent algorithm to optimize Eq. (9).

Algorithm. As shown in Fig. 2, our method takes a clean image as input \mathbf{I}_k . In addition to fetching references from a group of images, we design another approach in which \mathbf{I}_k is augmented by flipping, mirroring, left rotation, and right rotation to generate four reference images. We collect the original image and references together to build \mathcal{R} . Then, we initialize the coefficients \mathbf{a} of the exposure polynomial model, offset map \mathbf{U} , and pixel-wise noise θ_n , all with zero values. We follow Eq. (4) and Eq. (5) to perturb the clean image. Note that only the original image \mathbf{I}_k is perturbed. For initialization, the perturbation is an identity operation. We employ a sign gradient descent algorithm to optimize Eq. (9) and MI-FGSM is adopted in this paper. Finally, the optimized \mathbf{a} , \mathbf{U} , and θ_n are applied to the original clean image to generate a perturbed image as the output, which may fool CoSOD methods.

4. Experiments

4.1. Experimental setups

Datasets. We conduct experiments on Cosal2015 [48], iCoseg [2], CoSOD3k [13], and CoCA [52]. Cosal2015 and CoSOD3k are two large-scale datasets, containing 2,015 and 3,316 images of 50 and 160 groups respectively. iCoseg contains 643 images of 38 groups, in which objects are in the same scene within a group. CoCA is well-designed for evaluating CoSOD methods due to its multiple salient objects in each image. It contains 1,295 images of 80 groups.

Models. We employ three CoSOD models and one SOD model to evaluate the attacks. GICD [52] and GCAGC

	GICD				GCAGC				CBCD				PoolNet				
	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	
Cosal2015	Original	0.2352	0.8595	0.7800	0.0838	0.1881	0.8960	0.8275	0.0814	0.9504	0.6046	0.1530	0.2287	0.2625	0.8449	0.7626	0.1001
	Jadena _{single}	0.6864	0.7044	0.3831	0.1753	0.4868	0.7682	0.6162	0.1529	0.9663	0.5490	0.1199	0.2354	0.5538	0.6985	0.5152	0.1431
	Jadena _{group}	0.6849	0.7069	0.3838	0.1777	0.5221	0.7453	0.5842	0.1619	0.9826	0.5212	0.0715	0.2392	0.6060	0.6671	0.4650	0.1545
	Jadena _{augment}	0.6893	0.7016	0.3747	0.1749	0.5146	0.7432	0.5960	0.1592	0.9816	0.5137	0.0754	0.2394	0.6010	0.6648	0.4616	0.1513
	Jadena _{augment} w/o noise	0.3226	0.8297	0.7074	0.1037	0.2010	0.8860	0.8182	0.0870	0.9921	0.4853	0.0434	0.2422	0.3677	0.8032	0.6824	0.1145
Jadena _{augment} w/o exposure	0.5261	0.7536	0.5483	0.1437	0.3821	0.8185	0.7074	0.1243	0.9533	0.6028	0.1473	0.2290	0.4303	0.7680	0.6347	0.1240	
iCoseg	Original	0.2970	0.8189	0.7761	0.0851	0.1991	0.8721	0.8036	0.0788	0.7760	0.8005	0.4470	0.1744	0.2286	0.8845	0.7955	0.0709
	Jadena _{single}	0.6952	0.7063	0.4231	0.1666	0.3142	0.8369	0.7257	0.1111	0.8383	0.7422	0.3621	0.1849	0.5381	0.7677	0.5576	0.1346
	Jadena _{group}	0.6874	0.7070	0.4242	0.1663	0.3453	0.8167	0.7046	0.1166	0.9036	0.7073	0.2425	0.1979	0.5848	0.7474	0.4986	0.1450
	Jadena _{augment}	0.7092	0.7017	0.3966	0.1708	0.3733	0.8130	0.6989	0.1178	0.9051	0.7089	0.2324	0.1961	0.5785	0.7351	0.5002	0.1435
	Jadena _{augment} w/o noise	0.3997	0.7905	0.7171	0.1038	0.1975	0.8636	0.8015	0.0821	0.9425	0.6434	0.1361	0.2068	0.3515	0.8456	0.6965	0.0957
Jadena _{augment} w/o exposure	0.5490	0.7525	0.5593	0.1406	0.2457	0.8626	0.7736	0.0919	0.7916	0.7998	0.4307	0.1750	0.3670	0.8405	0.7051	0.1020	
CoSOD3k	Original	0.3555	0.7963	0.6981	0.0927	0.2995	0.8273	0.7672	0.0949	0.9717	0.5184	0.1400	0.2024	0.3667	0.7717	0.6947	0.1164
	Jadena _{single}	0.7630	0.6592	0.3282	0.1613	0.6092	0.6878	0.5481	0.1519	0.9795	0.4692	0.1080	0.2070	0.6622	0.6065	0.4425	0.1513
	Jadena _{group}	0.7672	0.6600	0.3158	0.1638	0.6441	0.6630	0.5125	0.1589	0.9873	0.4411	0.0632	0.2095	0.7156	0.5681	0.3823	0.1579
	Jadena _{augment}	0.7741	0.6544	0.3075	0.1630	0.6360	0.6601	0.5153	0.1586	0.9879	0.4400	0.0741	0.2096	0.7189	0.5614	0.3827	0.1590
	Jadena _{augment} w/o noise	0.4484	0.7657	0.6307	0.1075	0.3139	0.8154	0.7467	0.0993	0.9934	0.4119	0.0341	0.2095	0.4756	0.7192	0.6124	0.1265
Jadena _{augment} w/o exposure	0.6345	0.6959	0.4627	0.1404	0.5069	0.7361	0.6235	0.1316	0.9744	0.5158	0.1367	0.2030	0.5516	0.6779	0.5530	0.1376	
CoCA	Original	0.6680	0.5993	0.4939	0.1187	0.6062	0.5855	0.5970	0.1153	0.9606	0.3543	0.1738	0.1377	0.7259	0.4923	0.4707	0.1681
	Jadena _{single}	0.8046	0.5332	0.3650	0.1265	0.7660	0.4733	0.4893	0.1315	0.9737	0.3228	0.1455	0.1401	0.7954	0.4283	0.3829	0.1320
	Jadena _{group}	0.8340	0.5327	0.3335	0.1226	0.7915	0.4542	0.4584	0.1341	0.9876	0.2998	0.0974	0.1380	0.8355	0.3853	0.3221	0.1271
	Jadena _{augment}	0.8363	0.5284	0.3365	0.1224	0.7876	0.4588	0.4502	0.1347	0.9846	0.2973	0.0995	0.1392	0.8270	0.3909	0.3399	0.1284
	Jadena _{augment} w/o noise	0.6996	0.5852	0.4664	0.1120	0.6100	0.5737	0.5963	0.1179	0.9876	0.2833	0.0643	0.1352	0.7560	0.4571	0.4394	0.1541
Jadena _{augment} w/o exposure	0.7490	0.5535	0.4310	0.1250	0.7112	0.5244	0.5536	0.1222	0.9691	0.3526	0.1672	0.1385	0.7598	0.4588	0.4428	0.1405	

Table 1. Attack performance. “Original” means that we predict the co-saliency map of clean images using each CoSOD method. Jadena_{single, augment, group} indicates which variant of Jadena is used, and “w/o noise” or “w/o exposure” means that the attack is performed without additive noise or exposure tuning. We highlight the top results of each CoSOD method and each dataset in red.

	GICD				GCAGC				CBCD				PoolNet			
	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow
Original	0.2352	0.8595	0.7800	0.0838	0.1881	0.8960	0.8275	0.0814	0.9504	0.6046	0.1530	0.2287	0.2625	0.8449	0.7626	0.1001
Noise ₈	0.2238	0.8602	0.7863	0.0841	0.2074	0.8893	0.8179	0.0868	0.9464	0.6057	0.1535	0.2283	0.2655	0.8432	0.7628	0.1011
Noise ₁₆	0.2367	0.8525	0.7734	0.0886	0.2323	0.8779	0.8009	0.0950	0.9504	0.6054	0.1476	0.2289	0.2844	0.8358	0.7521	0.1048
Noise ₃₂	0.3012	0.8216	0.7310	0.1048	0.2988	0.8395	0.7519	0.1200	0.9598	0.5953	0.1376	0.2313	0.3429	0.8144	0.7078	0.1145
FGSM _{BCE}	0.4422 *	0.7644 *	0.6366 *	0.1351 *	0.3355	0.8309	0.7315	0.1295	0.9543	0.5961	0.1503	0.2302	0.3454	0.8089	0.7043	0.1175
FGSM _{L1}	0.4561 *	0.7655 *	0.6098 *	0.1410 *	0.3300	0.8342	0.7384	0.1294	0.9573	0.5982	0.1458	0.2302	0.3524	0.8077	0.6998	0.1168
I-FGSM _{BCE}	0.7285 *	0.6510 *	0.3970 *	0.3337 *	0.2769	0.8455	0.7644	0.1146	0.9533	0.6042	0.1491	0.2286	0.2754	0.8339	0.7503	0.1057
I-FGSM _{L1}	0.6397 *	0.6561 *	0.5134 *	0.2700 *	0.2615	0.8643	0.7801	0.1063	0.9454	0.6059	0.1533	0.2280	0.2715	0.8364	0.7533	0.1052
MI-FGSM _{BCE}	0.6700 *	0.6697 *	0.4480 *	0.2868 *	0.3975	0.7909	0.6842	0.1532	0.9514	0.5995	0.1468	0.2293	0.3375	0.8009	0.6976	0.1215
MI-FGSM _{L1}	0.5921 *	0.6541 *	0.5543 *	0.2331 *	0.3598	0.8193	0.7130	0.1355	0.9548	0.6024	0.1508	0.2294	0.3266	0.8113	0.7142	0.1180
TI-MI-FGSM _{BCE}	0.5380 *	0.7206 *	0.5371 *	0.2112 *	0.3146	0.8197	0.7322	0.1316	0.9623	0.5648	0.1504	0.2318	0.3628	0.7894	0.6784	0.1236
TI-MI-FGSM _{L1}	0.4903 *	0.7247 *	0.6074 *	0.1789 *	0.3097	0.8378	0.7514	0.1234	0.9583	0.5677	0.1500	0.2315	0.3489	0.8050	0.6900	0.1190
ColorFool	0.4199	0.7962	0.6871	0.1213	0.2169	0.8851	0.8123	0.0903	0.9663	0.5830	0.1048	0.2340	0.2829	0.8366	0.7471	0.1097
Jadena _{single}	0.6864	0.7044	0.3831	0.1753	0.4868	0.7682	0.6162	0.1529	0.9663	0.5490	0.1199	0.2354	0.5538	0.6985	0.5152	0.1431
Jadena _{group}	0.6849	0.7069	0.3838	0.1777	0.5221	0.7453	0.5842	0.1619	0.9826	0.5212	0.0715	0.2392	0.6060	0.6671	0.4650	0.1545
Jadena _{augment}	0.6893	0.7016	0.3747	0.1749	0.5146	0.7432	0.5960	0.1592	0.9816	0.5137	0.0754	0.2394	0.6010	0.6648	0.4616	0.1513
Jadena _{augment} w/o exposure	0.5261	0.7536	0.5483	0.1437	0.3821	0.8185	0.7074	0.1243	0.9533	0.6028	0.1473	0.2290	0.4303	0.7680	0.6347	0.1240

Table 2. Comparison with existing attacks. Noise _{ϵ} apply random additive noise sampled from a uniform distribution $U(-\epsilon, \epsilon)$ in each channel of the images. The attacks of FGSM and *FGSMs are performed against co-saliency labels and have full access to the network structure and parameters of the GICD model. *_{BCE, L1} means that binary cross-entropy (BCE) or L_1 loss is adopted as the objective function. We mark white-box attacks with stars and highlight the top three results in red, green, and blue respectively.

[50] are two state-of-the-art CoSOD methods. We employ GCAGC with an HRNet backbone [40]. PoolNet [31] is a CNN-based SOD method, but we treat it as a CoSOD method due to its competitive performance on CoSOD datasets. CBCD [14] is a traditional method for both CoSOD and SOD, but we only adopt the co-saliency results.

Metrics. In our experiments, we employ four metrics to evaluate how adversarial attacks affect detection methods, *i.e.*, average precision (AP) [47], F-measure score F_β with $\beta^2 = 0.3$ [1], mean absolute error (MAE) [47], and success rate (S). The first three metrics are widely used for CoSOD evaluations; we can evaluate adversarial attacks by observing how they change from the original performance. The last metric, success rate, aims to evaluate the ability of an attack to make images undiscoverable. We treat an attack as successful when the IoU between the result of a perturbed image and the corresponding ground-truth map is less than

0.5. IoU is widely used in object detection and visual tracking tasks and defined as $\text{IoU} = \frac{\text{Area of Overlap}}{\text{Area of Union}}$. To calculate success rate, we simply divide the total number of results by the number of successful results. Furthermore, to measure the quality of the perturbed images and the perceptibility of the applied perturbations, we employ a no-reference image quality assessor, *i.e.*, BRISQUE [33]. A smaller BRISQUE score suggests better quality of an image.

4.2. Attack performance and ablation study

We apply our method on several datasets and use four CoSOD methods to predict result maps for the generated adversarial examples. We evaluate three variants of the proposed Jadena, with varied reference strategies to build a group around the target image, *i.e.*, no references, references from the same group, and augmented references of original images, denoted by Jadena_{single, group, augment}, respectively. The augmentation of original images includes

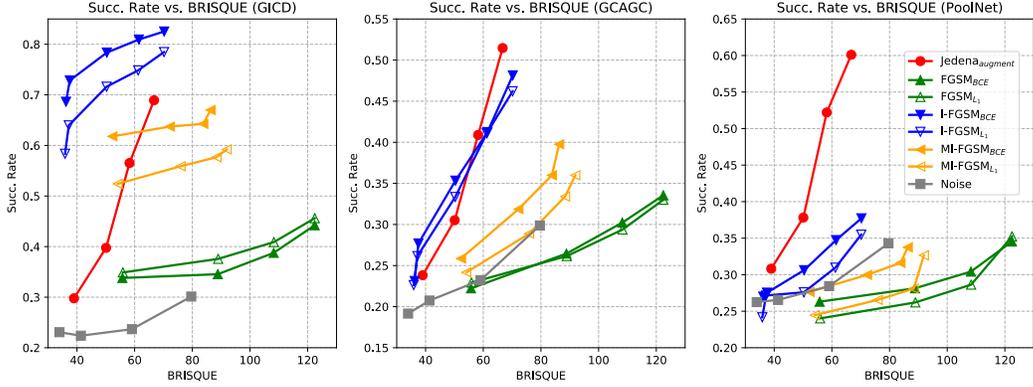


Figure 3. Success rate vs. BRISQUE. This figure shows the change in success rate of attacks with distortion degree. A larger BRISQUE suggests a worse quality of images, *i.e.*, more distortion. Note that, the baselines in the first plot are under white-box settings. The performance of ours is plotted in red.



Figure 4. Visualization of attack results. We show two visualized cases in this figure, with the original/perturbed image, the ground-truth of the co-saliency map and the results of GICD, GCAGC, and PoolNet in the columns, and highlight our method in green.

flipping, mirroring, left rotation, and right rotation, which are designed for full black-box attack, even when images of the same group are not provided. We optimize Eq. (7) for the first variant and Eq. (8) for the other two variants. Moreover, we apply Jedena_{augment} without additive noise perturbation or exposure tuning, and report them as Jedena_{augment} w/o noise and Jedena_{augment} w/o exposure.

For the feature extraction required for the objective function computation, we employ a general classification model, ResNet50 [22], and adopt the outputs of stage-{1, 2, 3}. For the optimization of our proposed objective function, we

use MI-FGSM and adopt $\mu = 1.0$ for momentum decay. For other hyperparameters of the proposed method, we set the iteration number $N = 20$, $\alpha_n = 1$ w.r.t. pixel values in $[0, 255]$ and maximum noise-perturbation $\epsilon = 16$, following standard additive-perturbation-based attack settings. For the optimization step size of exposure coefficients and the offset map, we use $\alpha_a = 0.1$ and $\alpha_U = 0.01$. To smooth the exposure tuning, we use $\lambda_b = 0.5$ for the “single” variant and $\lambda_b = 0.01$ for the “augment” or “group” variant. $\lambda_s = 0.01$ and polynomial degree $D = 10$, considering both success rate and imperceptibility of perturbations. We



Figure 5. Mixed perturbed images. #0, #2, #3, #8, #10, #13 are generated by the proposed method (#0 and #13 only contain exposure variation) and the others are generated by different baselines (FGSM: #1, #12, #15; I-FGSM: #5, #7, #17; MI-FGSM: #4, #11, #14; TI-MI-FGSM: #6, #9, #16).

report the performance across CoSOD methods in Table 1.

We observe that $Jedena_{augment}$ and $Jedena_{group}$ achieve the best performance on most metrics. Even without sampling images from the same group and only augmenting the original images, $Jedena_{augment}$ is competitive to $Jedena_{group}$ for GCAGC and PoolNet, and outperforms $Jedena_{group}$ for GICD. This demonstrates that providing the references is not necessary, as they may be obtained by augmentation. $Jedena_{augment}$ performs better than the “w/o noise” and “w/o exposure” versions, except for CBCD, which suggests that our proposed joint perturbation plays an important role in fooling CoSOD methods. While it is difficult to handle challenging CoSOD datasets, we find that the traditional method based on color histogram, *i.e.*, CBCD, is much more robust to noise than exposure tuning, since it shows a lower performance for “w/o noise” than “w/o exposure”. In contrast, the other deep methods are more sensitive to noise.

4.3. Comparison with existing attacks

Baselines. We consider a series of additive-perturbation-noise attacks as baselines, *i.e.*, FGSM [18], I-FGSM [26], MI-FGSM [10], TI-MI-FGSM [11], and ColorFool [36]. For the hyperparameters, we set the max perturbation $\epsilon = 16$ with pixel values in $[0, 255]$, as well as iteration number $N = 20$, and step size $\alpha_n = 1$ for iteration attacks. Momentum decay $\mu = 1.0$ for momentum methods, and the Gaussian kernel length is 15 for TI-MI-FGSM. Since no specific attacks for CoSOD methods have been proposed, we follow the white-box attack pattern and employ L_1 norm and binary cross-entropy (BCE) as objective functions.

Attack comparison. We launch baselines and our method on the Cosal2015 dataset. The hyperparameters of our method follow Sec. 4.2. The comparison results are shown in Table 2. Our proposed $Jedena$ outperforms all existing additive-perturbation-based attacks, except for white-box attacks, which indicates the better transferability of our method. Since we only borrow several shallow layers from the deep classification model, our proposed method has a lower computational cost than iterative attack methods

with a result-based objective function, which unavoidably back-propagate gradients through the whole deep model. To further clarify the advantages of the proposed attack method over baselines, we can compare our method, *i.e.*, $Jedena_{augment}$ w/o exposure that disables the exposure variation, with baseline methods under the same p -norm ball with $\epsilon_n = 16$. Our method still outperforms the baselines on success rates at most of the time.

Image quality comparison. We evaluate the image quality of generated adversarial examples and investigate how distortion degree affects the attack performance. For baselines attacks, we follow the default settings in this section. For our method, we adopt the “augment” variant, *i.e.*, $Jedena_{augment}$, which is practical for application in real-world settings. We provide visualizations of how different attack methods perturb an image in Fig. 4 and Fig. 5. Note that, in a real attack scenario, the original images would not be provided. To justify this phenomenon, we randomly mix images perturbed by different attacks in Fig. 5. For each attack method, we also plot success rate vs. BRISQUE in Fig. 3 by changing the maximum noise perturbation ϵ_n and λ_b . We do not adopt TI-MI-FGSM since it always applies unnatural coarse-grained noise on images and drastically reduces image quality. We have three findings: (1) Compared with white-box attacks for GICD, our method has competitive performance with MI-FGSM, however, hold worse performance than I-FGSM. (2) When transferred to GCAGC and PoolNet, our method outperforms most additive-perturbation-based methods. For SOD method PoolNet in particular, our method shows significant advantages over I-FGSM and MI-FGSM. This is reasonable for white-box attacks being adversarial against a CoSOD method instead of an SOD method and thus showing lower transferability. (3) Under the setup in Fig. 5, our adversarial images are imperceptible as the baselines.

Defense. We adopt a defense method, *i.e.*, HGD [30], to further demonstrate the effectiveness of the proposed method and report the results in Table 3. We find that the proposed method still outperforms the baselines except for the white-box settings. We report more defense results with other defense methods in the supplementary materials.

Attack comparison against SOD methods. As shown in Table 4, we also apply the proposed method against SOD methods and report the performance on the DUTS-TE [41] dataset. We employ white-box attacks against PoolNet [31] as the baseline and evaluate all attacks with PoolNet, EGNet [53], BASNet [35], and U²Net [34]. Our proposed methods achieve competitive performance and higher transferability across SOD methods than the white-box attacks. It indicates that our proposed loss function and joint perturbation can be adopted to attack SOD methods.

	GICD w/ HGD				GCAGC w/ HGD				CBCD w/ HGD				PoolNet w/ HGD			
	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow
Original	0.2352	0.8595	0.7800	0.0838	0.1881	0.8960	0.8275	0.0814	0.9504	0.6046	0.1530	0.2287	0.2625	0.8449	0.7626	0.1001
Noise ₈	0.2645	0.8510	0.7565	0.0905	0.1950	0.8917	0.8234	0.0840	0.9444	0.6068	0.1496	0.2279	0.2566	0.8495	0.7690	0.0992
Noise ₁₆	0.2710	0.8464	0.7571	0.0933	0.2283	0.8826	0.8063	0.0913	0.9499	0.6044	0.1512	0.2280	0.2645	0.8458	0.7644	0.1016
Noise ₃₂	0.3117	0.8250	0.7238	0.1051	0.2769	0.8585	0.7697	0.1059	0.9548	0.6017	0.1480	0.2292	0.3052	0.8284	0.7335	0.1097
FGSM _{BCE}	0.4194 *	0.7753 *	0.6439 *	0.1311 *	0.3141	0.8415	0.7476	0.1194	0.9543	0.5953	0.1492	0.2295	0.3136	0.8218	0.7244	0.1133
FGSM _{L1}	0.4447 *	0.7707 *	0.6205 *	0.1381 *	0.3002	0.8446	0.7570	0.1179	0.9509	0.5978	0.1496	0.2290	0.3176	0.8205	0.7216	0.1131
I-FGSM _{BCE}	0.7012 *	0.6639 *	0.4141 *	0.3114 *	0.2452	0.8583	0.7818	0.1055	0.9459	0.6049	0.1511	0.2282	0.2640	0.8435	0.7639	0.1022
I-FGSM _{L1}	0.6094 *	0.6692 *	0.5340 *	0.2519 *	0.2357	0.8735	0.7936	0.0988	0.9484	0.6056	0.1544	0.2277	0.2591	0.8451	0.7640	0.1019
MI-FGSM _{BCE}	0.6491 *	0.6773 *	0.4616 *	0.2749 *	0.3687	0.8000	0.7058	0.1424	0.9519	0.5998	0.1479	0.2289	0.3072	0.8154	0.7184	0.1156
MI-FGSM _{L1}	0.5826 *	0.6649 *	0.5562 *	0.2249 *	0.3310	0.8258	0.7363	0.1260	0.9514	0.6012	0.1490	0.2287	0.3037	0.8253	0.7326	0.1130
TI-MI-FGSM _{BCE}	0.5444 *	0.7148 *	0.5285 *	0.2066 *	0.3062	0.8162	0.7350	0.1354	0.9593	0.5670	0.1537	0.2310	0.3787	0.7834	0.6666	0.1259
TI-MI-FGSM _{L1}	0.4998 *	0.7215 *	0.5987 *	0.1776 *	0.2973	0.8360	0.7594	0.1228	0.9583	0.5680	0.1542	0.2308	0.3628	0.8022	0.6808	0.1207
Jadena _{single}	0.6988	0.7051	0.3721	0.1778	0.4675	0.7630	0.6331	0.1546	0.9648	0.5500	0.1194	0.2347	0.4739	0.7335	0.5884	0.1325
Jadena _{group}	0.6968	0.7106	0.3738	0.1797	0.4774	0.7499	0.6205	0.1625	0.9846	0.5203	0.0664	0.2388	0.5290	0.7040	0.5273	0.1415
Jadena _{augment}	0.6983	0.7034	0.3617	0.1779	0.4893	0.7452	0.6160	0.1611	0.9831	0.5167	0.0741	0.2390	0.5266	0.6969	0.5262	0.1399

Table 3. Defense results with HGD. The settings are the same as Table 2 except the CoSOD methods are equipped with HGD for defense. We mark white-box attacks with stars and highlight top 3 results with **red**, **green**, and **blue** respectively. The attack methods maintains similar performance rankings compared to the attack methods w/o defense which is reported in Table 2.

	PoolNet				EGNet				BASNet				U ² Net			
	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow	S \uparrow	AP \downarrow	F_β \downarrow	MAE \uparrow
Original	0.1368	0.9251	0.8546	0.0463	0.1748	0.9246	0.8340	0.0551	0.0809	0.9261	0.8992	0.0369	0.0822	0.9395	0.8928	0.0378
Noise ₈	0.1375	0.9263	0.8547	0.0467	0.1831	0.9243	0.8318	0.0555	0.0843	0.9237	0.8983	0.0384	0.0857	0.9386	0.8924	0.0394
Noise ₁₆	0.1472	0.9225	0.8469	0.0493	0.2032	0.9218	0.8219	0.0581	0.0940	0.9163	0.8884	0.0418	0.0974	0.9315	0.8861	0.0424
Noise ₃₂	0.2163	0.8982	0.8052	0.0616	0.2716	0.9007	0.7793	0.0696	0.1417	0.8912	0.8639	0.0513	0.1700	0.9046	0.8534	0.0560
FGSM _{BCE}	0.5563 *	0.6948 *	0.5837 *	0.1354 *	0.5308	0.7856	0.6054	0.1154	0.3117	0.8059	0.7743	0.0841	0.2999	0.8320	0.7637	0.0881
FGSM _{L1}	0.2191 *	0.8556 *	0.7740 *	0.0733 *	0.2716	0.8773	0.7701	0.0746	0.1679	0.8751	0.8459	0.0565	0.1638	0.8887	0.8402	0.0594
I-FGSM _{BCE}	0.1576 *	0.9158 *	0.8392 *	0.0518 *	0.2232	0.9104	0.8023	0.0631	0.1196	0.8986	0.8654	0.0502	0.1237	0.9146	0.8644	0.0505
I-FGSM _{L1}	0.1548 *	0.9202 *	0.8406 *	0.0499 *	0.2101	0.9153	0.8088	0.0609	0.1140	0.9073	0.8726	0.0471	0.1099	0.9228	0.8720	0.0477
MI-FGSM _{BCE}	0.9364 *	0.3619 *	0.3078 *	0.4503 *	0.5764	0.7799	0.5688	0.1169	0.3172	0.8125	0.7815	0.0822	0.3338	0.8249	0.7593	0.0899
MI-FGSM _{L1}	0.7146 *	0.6320 *	0.4873 *	0.2570 *	0.4340	0.8077	0.6556	0.1070	0.2668	0.8059	0.7718	0.0871	0.2875	0.8266	0.7582	0.0923
TI-MI-FGSM _{BCE}	0.8466 *	0.4869 *	0.3979 *	0.3467 *	0.6683	0.7432	0.4952	0.1310	0.3877	0.7761	0.7530	0.0922	0.3704	0.8107	0.7480	0.0940
TI-MI-FGSM _{L1}	0.5197 *	0.7156 *	0.5794 *	0.1790 *	0.4202	0.8047	0.6550	0.1089	0.2605	0.8190	0.7747	0.0820	0.2370	0.8474	0.7777	0.0827
Jadena _{single}	0.5771	0.7287	0.5179	0.1136	0.5722	0.7706	0.5337	0.1123	0.3891	0.7666	0.7837	0.0866	0.4181	0.7948	0.7405	0.0932
Jadena _{augment}	0.6745	0.6645	0.4185	0.1298	0.7091	0.7062	0.3997	0.1325	0.4983	0.7144	0.7412	0.1030	0.5121	0.7565	0.7006	0.1061

Table 4. Comparison with existing attacks against SOD. Noise _{ϵ} s apply random additive noise sampled from uniform distribution $U(-\epsilon, \epsilon)$ on each channel. The attacks of FGSM and *-FGSMs are launched against single saliency labels in a white-box manner. *_{BCE, L1} means that binary cross-entropy (BCE) or L_1 loss is adopted as the objective function. We mark white-box attacks with stars and highlight top 3 results with **red**, **green**, and **blue** respectively.

5. Conclusion

In this work, we have investigated a novel problem: how to effectively protect sensitive image content from being discovered by the state-of-the-art CoSOD methods. We have addressed this problem from the perspective of adversarial attacks and identified a novel task, *i.e.*, adversarial co-saliency attack. Given an image selected from a group of images containing some common and salient objects, the aim of this task is to generate an adversarial version of the image that can mislead CoSOD methods to predict incorrect co-salient regions, thus evading co-saliency detection and making the images undiscoverable by the co-saliency detection algorithms. We have proposed the very first black-box *joint adversarial exposure and noise attack (Jadena)*, where we jointly and locally tune the exposure and additive perturbations of the image according to a newly designed *high-feature-level contrast-sensitive loss* function. Our method, without needing any information on the state-of-the-art CoSOD methods, leads to significant performance degradation on various co-saliency detection datasets and makes the co-salient objects undetectable. The

effectiveness of the proposed method is thoroughly validated and showcased through various quantitative and qualitative experiments. We believe that the proposed method can greatly benefit the design and execution of privacy-oriented content sharing in the public domain where malicious extractions can be effectively inhibited.

Acknowledgment: This work was supported by the National Key Research and Development Program of China (No. 2020YFC1522700), the National Natural Science Foundation of China (No. 62072334), and the General Project of Tianjin (No. 18JCYBJC15200). It was also supported in part by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2020-019), Singapore National Cybersecurity R&D Program No. NRF2018NCR-NCR005-0001, National Satellite of Excellence in Trustworthy Software System No. NRF2018NCR-NSOE003-0001, NRF Investigatorship No. NRFI06-2020-0022-0001, and A*STAR AI3 HTPO Seed Fund (C211118012). We gratefully acknowledge the support of NVIDIA AI Tech Center (NVAITC) and AWS Cloud Credits for Research Award.

References

- [1] R. Achanta, S. Hemami, F. Estrada, and S. Susstrunk. Frequency-tuned salient region detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1597–1604, 2009. 2, 5
- [2] Dhruv Batra, Adarsh Kowdle, Devi Parikh, Jiebo Luo, and Tsuhan Chen. iCoseg: Interactive co-segmentation with intelligent scribble guidance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3169–3176, 2010. 4
- [3] Xiaochun Cao, Zhiqiang Tao, Bao Zhang, Huazhu Fu, and Wei Feng. Self-adaptively weighted co-saliency detection via rank constraint. *IEEE Transactions on Image Processing*, 23(9):4175–4186, 2014. 2
- [4] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*, pages 39–57, 2017. 2
- [5] Xiaojun Chang, Yi Yang, Eric Xing, and Yaoliang Yu. Complex event detection using semantic saliency and nearly-isotonic svm. In *International Conference on Machine Learning*, pages 1348–1357, 2015. 2
- [6] Zhaohui Che, Ali Borji, Guangtao Zhai, Suiyi Ling, Guodong Guo, and Patrick Le Callet. Adversarial attacks against deep saliency models. *arXiv preprint arXiv:1904.01231*, 2019. 2
- [7] M.-M. Cheng, N.J. Mitra, X. Huang, P. Torr, and S.-M. Hu. Global contrast based salient region detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3):569–582, 2015. 2
- [8] Ming-Ming Cheng, Niloy J Mitra, Xiaolei Huang, and Shi-Min Hu. Salienshape: Group saliency in image collections. *The Visual Computer*, 30(4):443–453, 2014. 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*, 2018. 2
- [10] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9185–9193, 2018. 2, 7
- [11] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 7
- [12] D.-P. Fan, T. Li, Z. Lin, G. P. Ji, D. Zhang, M. M. Cheng, H. Fu, and J. Shen. Re-thinking co-salient object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [13] Deng-Ping Fan, Zheng Lin, Ge-Peng Ji, Dingwen Zhang, Huazhu Fu, and Ming-Ming Cheng. Taking a deeper look at the co-salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 4
- [14] Huazhu Fu, Xiaochun Cao, and Zhuowen Tu. Cluster-based co-saliency detection. *IEEE Transactions on Image Processing*, 22(10):3766–3778, 2013. 1, 2, 5
- [15] Xueyang Fu, Delu Zeng, Yue Huang, Xiao-Ping Zhang, and Xinghao Ding. A weighted variational model for simultaneous reflectance and illumination estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2782–2790, 2016. 3
- [16] Ruijun Gao, Qing Guo, Felix Juefei-Xu, Hongkai Yu, and Wei Feng. AdvHaze: Adversarial haze attack. *arXiv preprint arXiv:2104.13673*, 2021. 2
- [17] Chenjie Ge, Keren Fu, Fanghui Liu, Li Bai, and Jie Yang. Co-saliency detection via inter and intra saliency propagation. *Signal Processing: Image Communication*, 44:69–83, 2016. 2
- [18] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *International Conference on Machine Learning*, 2015. 2, 7
- [19] Qing Guo, Ziyi Cheng, Felix Juefei-Xu, Lei Ma, Xiaofei Xie, Yang Liu, and Jianjun Zhao. Learning to adversarially blur visual object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10839–10848, 2021. 2
- [20] Qing Guo, Felix Juefei-Xu, Xiaofei Xie, Lei Ma, Jian Wang, Bing Yu, Wei Feng, and Yang Liu. Watch out! motion is blurring the vision of your deep neural networks. In *Advances in Neural Information Processing Systems*, volume 33, pages 975–985, 2020. 2
- [21] Qing Guo, Xiaofei Xie, Felix Juefei-Xu, Lei Ma, Zhongguo Li, Wanli Xue, Wei Feng, and Yang Liu. Spark: Spatial-aware online incremental attack against visual tracking. In *Proceedings of the European Conference on Computer Vision*, 2020. 2
- [22] 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. 6
- [23] Koteswar Rao Jerripothula, Jianfei Cai, and Junsong Yuan. Quality-guided fusion-based co-saliency estimation for image co-segmentation and colocalization. *IEEE Transactions on Multimedia*, 20(9):2466–2477, 2018. 2
- [24] Xiaojun Jia, Xingxing Wei, Xiaochun Cao, and Xiaoguang Han. Adv-watermark: A novel watermark perturbation for adversarial examples. *arXiv preprint arXiv:2008.01919*, 2020. 2
- [25] 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 ACM International Conference on Multimedia*, pages 1375–1382, 2019. 2
- [26] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *International Conference on Learning Representations*, 2017. 2, 7
- [27] Bo Li, Zhengxing Sun, Lv Tang, Yunhan Sun, and Jinlong Shi. Detecting robust co-saliency with recurrent co-attention neural network. In *Proceedings of the IEEE Conference on International Joint Conferences on Artificial Intelligence*, pages 818–825, 2019. 2

- [28] Guanbin Li and Yizhou Yu. Deep contrast learning for salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [29] Hongliang Li, Fanman Meng, Bing Luo, and Shuyuan Zhu. Repairing bad co-segmentation using its quality evaluation and segment propagation. *IEEE Transactions on Image Processing*, 23(8):3545–3559, 2014. 2
- [30] Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, and Jun Zhu. Defense against adversarial attacks using high-level representation guided denoiser. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 7
- [31] Jiang-Jiang Liu, Qibin Hou, Ming-Ming Cheng, Jiashi Feng, and Jianmin Jiang. A simple pooling-based design for real-time salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1, 5, 7
- [32] Yan Luo, Ming Jiang, Yongkang Wong, and Qi Zhao. Multi-camera saliency. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(10):2057–2070, 2015. 2
- [33] Anish Mittal, Anush K Moorthy, and Alan C Bovik. Blind/referenceless image spatial quality evaluator. In *Asilomar Conference on Signals, Systems and Computers*, pages 723–727, 2011. 5
- [34] Xuebin Qin, Zichen Zhang, Chenyang Huang, Masood Dehghan, Osmar Zaiane, and Martin Jagersand. U2-Net: Going deeper with nested u-structure for salient object detection. *Pattern Recognition*, 106:107404, 2020. 7
- [35] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. BASNet: Boundary-aware salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 7
- [36] Ali Shahin Shamsabadi, Ricardo Sanchez-Matilla, and Andrea Cavallaro. ColorFool: Semantic adversarial colorization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 7
- [37] Kevin Tang, Armand Joulin, Li-Jia Li, and Li Fei-Fei. Co-localization in real-world images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1464–1471, 2014. 2
- [38] Binyu Tian, Felix Juefei-Xu, Qing Guo, Xiaofei Xie, Xiaohong Li, and Yang Liu. AVA: Adversarial vignetting attack against visual recognition. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2021. 2
- [39] Richard Tran, David Patrick, Michael Geyer, and Amanda Fernandez. SAD: Saliency-based defenses against adversarial examples. *arXiv preprint arXiv:2003.04820*, 2020. 2
- [40] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 5
- [41] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 7
- [42] Linzhao Wang, Lijun Wang, Huchuan Lu, Pingping Zhang, and Xiang Ruan. Saliency detection with recurrent fully convolutional networks. In *European Conference on Computer Vision*, pages 825–841, 2016. 2
- [43] Run Wang, Felix Juefei-Xu, Qing Guo, Yihao Huang, Xiaofei Xie, Lei Ma, and Yang Liu. Amora: Black-box adversarial morphing attack. In *Proceedings of the ACM International Conference on Multimedia*, 2020. 2
- [44] W. Wang, Q. Lai, H. Fu, J. Shen, H. Ling, and R. Yang. Salient object detection in the deep learning era: An in-depth survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [45] Zheng-Jun Zha, Chong Wang, Dong Liu, Hongtao Xie, and Yongdong Zhang. Robust deep co-saliency detection with group semantic and pyramid attention. *IEEE Transactions on Neural Networks and Learning Systems*, 2020. 2
- [46] Liming Zhai, Felix Juefei-Xu, Qing Guo, Xiaofei Xie, Lei Ma, Wei Feng, Shengchao Qin, and Yang Liu. It's Raining Cats or Dogs? Adversarial Rain Attack on DNN Perception. *arXiv preprint arXiv:2009.09205*, 2020. 2
- [47] 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*, 9(4):1–31, 2018. 1, 5
- [48] 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. 2, 4
- [49] Dingwen Zhang, Deyu Meng, Chao Li, Lu Jiang, Qian Zhao, and Junwei Han. A self-paced multiple-instance learning framework for co-saliency detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 594–602, 2015. 2
- [50] Kaihua Zhang, Tengpeng 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 Conference on Computer Vision and Pattern Recognition*, pages 9050–9059, 2020. 1, 2, 5
- [51] Qing Zhang, Yongwei Nie, Lei Zhu, Chunxia Xiao, and Wei-Shi Zheng. Enhancing underexposed photos using perceptually bidirectional similarity. *IEEE Transactions on Multimedia*, pages 1–1, 2020. 3
- [52] Zhao Zhang, Wenda Jin, Jun Xu, and Ming-Ming Cheng. Gradient-induced co-saliency detection. In *European Conference on Computer Vision*, 2020. 1, 2, 4
- [53] Jia-Xing Zhao, Jiang-Jiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. EGNet: Edge guidance network for salient object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 7
- [54] Hongyuan Zhu, Fanman Meng, Jianfei Cai, and Shijian Lu. Beyond pixels: A comprehensive survey from bottom-up to semantic image segmentation and cosegmentation. *Journal of Visual Communication and Image Representation*, 34:12–27, 2016. 2