

GALA: Toward Geometry-and-Lighting-Aware Object Search for Compositing

Sijie Zhu^{1,†}, Zhe Lin², Scott Cohen², Jason Kuen², Zhifei Zhang², Chen Chen¹

¹ Center for Research in Computer Vision, University of Central Florida

² Adobe Research

sizhu@knights.ucf.edu, {zlin, scohen, kuen, zzhang}@adobe.com, chen.chen@crcv.ucf.edu

Abstract. Compositing-aware object search aims to find the most compatible objects for compositing given a background image and a query bounding box. Previous works focus on learning compatibility between the foreground object and background, but fail to learn other important factors from large-scale data, i.e. geometry and lighting. To move a step further, this paper proposes GALA (Geometry-and-Lighting-Aware), a generic foreground object search method with discriminative modeling on geometry and lighting compatibility for open-world image compositing. Remarkably, it achieves state-of-the-art results on the CAIS dataset and generalizes well on large-scale open-world datasets, i.e. Pixabay and Open Images. In addition, our method can effectively handle non-box scenarios, where users only provide background images without any input bounding box. Experiments are conducted on real-world images to showcase applications of the proposed method for compositing-aware search and automatic location/scale prediction for the foreground object.

Keywords: Foreground Object Retrieval, Image Compositing

1 Introduction

Compositing-aware object search/retrieval [26] aims to find suitable source images for compositing [17,24]. Specifically, given a background image and a bounding box indicating the compositing location, the objective is to retrieve compatible foreground objects from a large reference database, so that the composite image appears realistic. Harmonization [7,29,22,2,15] is then applied to adjust the color and edge pixels, but it is extremely challenging to automatically adjust the semantics, lighting, or geometry of foreground objects. Although recent re-lighting method [19] can generate realistic lighting change for human portraits, it does not tackle general object categories and the required 3D reconstruction is not always available. Therefore, the quality of the final composite image highly depends on the performance of foreground retrieval system, and a good system should be aware of semantics, lighting, and geometry of foreground objects.

† This work was done during the first author’s internship at Adobe Research.

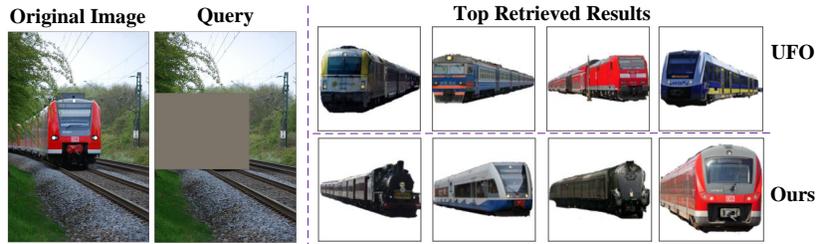


Fig. 1: Comparison between state-of-the-art method (UFO [27] in the first row) and the proposed geometry-and-lighting-aware search (second row). The retrieved objects in the first row do not respect the geometry of the background scene, while results of the proposed method have better geometry compatibility.

Early work [26] on foreground retrieval follows a constrained setting, which requires the user to specify the object category for retrieval. Specifying the category sets a limit on the search space, thus preventing the system from recommending diverse sets of objects. Later, UFO [27] proposes an unconstrained search method, i.e. objects from all categories are considered as candidates for retrieval. The unconstrained setting is closer to real-world scenarios, which require a large and diverse foreground object gallery to satisfy different users. However, UFO [27] focuses on finding semantically compatible foreground object and does not explicitly model lighting and geometry, which are critical factors for making object compositing realistic, as shown in Figs. 1 and Sec. 1 in supplementary material.

More recent work [13,23] either explores additional annotation [13] for constrained setting, or only focuses on indoor furniture [23] with fine-grained sub-categories. Therefore, how to encourage awareness on geometry and lighting is still unclear for general unconstrained foreground retrieval. One way to consider these factors is to explicitly estimate the 3D geometry and lighting of the scene, but it is extremely challenging to obtain realistic training data with 3D labels to do so. Instead, we aim to build a discriminative model which is sensitive to lighting and geometry mismatches using real-world image datasets, leading to a more generalizable and scalable solution.

In this paper, we propose a novel Geometry-And-Lighting-Aware (GALA) foreground object search system, which aims to retrieve objects that are compatible in terms of semantics, geometry, and lighting. Specifically, we design a model consisting of a foreground object encoder and background encoder, such that only the matching pairs of foreground object and background are closer to each other in the embedding space. To encode geometry and lighting sensitivity, we generate negative samples by augmenting the same objects with very different geometry and lighting conditions through homography transformation with left-right flip and non-linear illumination modification on foreground objects during training. Contrastive learning is then applied to push the transformed foreground object far away from the original one in the embedding space. However, the se-

Table 1: Comparison between our method and previous works on settings.

	Zhao [26]	UFO [27]	Li [13]	Wu [23]	Ours
One model for all	✓	✓	✓	✗	✓
Unspecified Class	✗	✓	✗	✓	✓
Non-box Scenarios	✗	✗	✗	✗	✓
Large-scale Dataset	✗	✗	✗	✗	✓
Open-world Setting	✗	✗	✗	✗	✓

mantic compatibility of retrieved objects could degrade significantly if we train the foreground network jointly with the background network. Thus, we introduce an alternating training strategy to maintain semantic compatibility, while learning to respect geometry and lighting. GALA can also be extended to handle non-box scenarios, where the users only provide the background image without any bounding box or text input. Our model automatically retrieves foreground objects and predicts the best location and scale for compositing.

Furthermore, we perform experiments and validate GALA for open-world datasets with an arbitrarily large number of categories, targeting real-world unconstrained applications. Previous works are mostly implemented on medium scale datasets (< 100,000 objects) with limited categories, e.g. MSCOCO [16] only has 80 object categories. In contrast, we conduct experiments on large-scale real-world datasets, i.e. Pixabay [1] and Open Images [9], which contain significantly more categories and images with diverse contents. We show comparisons of various settings between the proposed method and previous works in Table 1. Our contributions are summarized as follows:

- A novel compositing-aware search method that discriminatively models geometry and lighting with contrastive training.
- An alternating training strategy to address the challenge of losing semantic compatibility when learning the foreground network in an unconstrained setting.
- Extensive experiments to demonstrate significant improvements over previous works on the CAIS dataset, as well as a new large-scale open-world dataset, i.e. Pixabay.
- An extension of the method for non-box conditioned scenarios which are never considered in previous works.

2 Related Work

Search for Compositing. The idea of searching for foregrounds to insert into a new background is first proposed by Lalonde et al. [10], by explicitly estimating the 3D geometry of foreground objects, and the lighting map of background scenes. Zhao et al. [26] propose a learning-based compositing-aware search method, where the users provide a specific class as text input. The positive samples are augmented based on shape and semantics. UFO [27] assumes that the user does not specify any category, which is the most relevant setting to our

work, but it only focuses on the semantic compatibility and ignores lighting and geometry. It trains a discriminator to distinguish real/fake object for a certain background and selects candidate positive samples for each background image. Li et al. [13] manually annotates multiple attributes as additional information to determine “Interpretable Foreground Object” with given categories. Wu et al. [23] propose a teacher-student framework for fine-grained indoor categories, which adopts multiple pre-trained features for foreground objects as teacher and trains the background network with KL-divergence. None of them learn discriminative features on geometry and lighting for general unconstrained foreground search.

Object placement. A thread of works [11,25,14] focus on object placement along with adversarial training, which predicts locations and scales to insert an object. Lee et al. [11] train a spatial transformer network to predict the location for placement and a shape mask to guide generation. Zhang et al. [25] train a location-scale prediction model using inpainted pure background images. Li et al. [14] first find candidate locations in indoor scenes, then predict the best human pose for a specific location. These methods either focus on street scenes with limited object categories, e.g. person, car, or indoor scenes with only human pose joints. Our location-scale prediction deals with a more challenging setting with general objects from diverse scenes by directly adopting our retrieval model.

3 Method

In this section, we introduce GALA, Geometry-and-Lighting-Aware foreground object search for compositing. We first formulate the problem and describe training data generation in Sec. 3.1. Then we present the details of contrastive learning with self-transformations and the alternating training strategy in Sec. 3.2 and 3.3. Finally, we show how GALA handles non-box scenarios in Sec. 3.4.

3.1 Problem Statement

Given a background image I_b with a bounding box (l, r, w, h) , our objective is to retrieve a set of most compatible foreground object images $\{I_f\}$, so that realistic composite images can be generated with simple harmonization techniques. Our framework aims to learn an embedding space so that compatible images are close to each other in this space and the ranking can be obtained by simply computing the cosine similarity or euclidean distance. Since foreground and background images have very different distributions and appearances, we use two different encoder networks N_b, N_f to generate the embedding features for I_b, I_f , respectively. During training, the only available annotation is semantic segmentation mask for each object. We use the mask to crop out the object as I_f . Then I_b is generated by applying a rectangle mask on the image covering the whole object. Since there is no manual annotation on positive and negative samples, we consider I_b and I_f generated from the same image as a positive sample to each

other, while other image pairs are considered as negative samples. With a hard-margin triplet loss [3], the optimization can be formulated as $\arg \min_{N_b, N_f} \mathcal{L}_t$, where

$$\mathcal{L}_t = [S(N_b(I_b), N_f(I_f^-)) - S(N_b(I_b), N_f(I_f^+)) + m]_+. \quad (1)$$

Here S means cosine similarity, and $[\cdot]_+$ denotes the hinge function. m is the margin for triplet loss and I_f^+, I_f^- indicate the positive and negative foreground objects for a given I_b .

3.2 Contrastive Learning with Self-transformation

The standard loss in Eq. 1 only considers the contrastive information between the original foreground object and other objects, which mainly focuses on semantic information. We argue that a Geometry-and-Lighting-aware system should be able to tell the difference if the foreground object has exactly the same semantics with different lighting and geometry conditions. We thus perform transformation on the original positive foreground object so that it can be considered as negative.

As shown in Fig. 2, the transformed foreground objects have very different lighting or geometry conditions, which do not match with the background image anymore. We consider the transformed object image I_f^t as negative and formulate another triplet loss (Fig. 3):

$$\mathcal{L}_c = [S(N_b(I_b), N_f(I_f^t)) - S(N_b(I_b), N_f(I_f)) + m]_+. \quad (2)$$

The final loss is given by $\mathcal{L} = \mathcal{L}_t + \mathcal{L}_c$ from Eqs. 1 and 2. In Fig. 2, we show the pipeline of self transformations. For geometry transformation, we first apply random homography transformations, then left-right flip is applied with a 50% probability. For lighting transformation, we first find a random background image and apply Gaussian blur with a large radius, e.g. 100. The blurred map is then resized with interpolation to the size of foreground object, and masked with the segmentation mask. Finally, we enhance the variance of the lighting

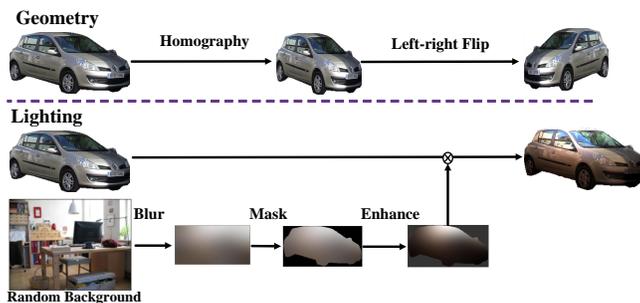


Fig. 2: Pipeline of generating lighting and geometry transformations used in contrastive learning.

map with an exponential function so that the largest value of the map is 5. In this way, we get a non-linear illumination map which highlights a random region in the object image, and it is multiplied by the original object image to generate the final transformed image. Although the lighting may not be as realistic as 3D relighting [18], it is enough to be considered as a negative sample with very different illuminations.

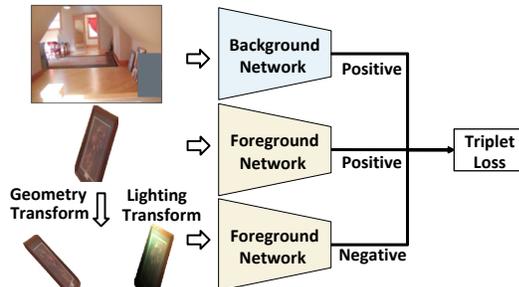


Fig. 3: Contrastive learning with self-transformations.

Table 2: Performance of different training strategies on CAIS.

	mAP
Fix Foreground	29.10
Direct Training	17.93
Aug	23.45
Aug + Alternating	31.20

3.3 Alternating Training

We notice that previous works [27,23] generally use pre-trained weights for foreground network N_f and freeze its parameters during training (denoted as “Fix Foreground”), which means N_f cannot learn from data. Since data augmentation like left-right flip is adopted for pre-training on ImageNet [4], the feature would be invariant to the left-right flip, which could change the geometry and direction of lighting. Therefore, the pre-trained weights are not discriminative to lighting and geometry changes, which will be further demonstrated in Sec. 4.2.

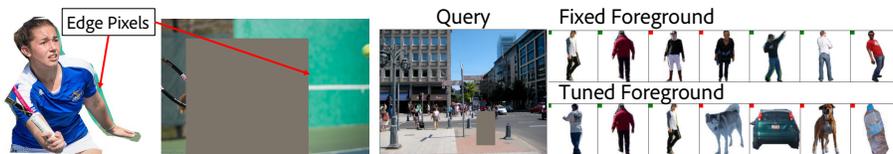


Fig. 4: Imperfect segmentation mask and edge pixel example. Fig. 5: An example of poor compatibility on semantics when training foreground network.

To learn geometry-and-lighting-aware representation, N_f must be learnable. However, performance drops significantly if N_f and N_b are directly trained with Eq. 1 (denoted as “Direct Training” in Table 2), partly due to the imperfect segmentation mask, as some annotations (MSCOCO) are based on polygon. When cropping with an imperfect mask (Fig. 4), some edge pixels of the object are actually background pixels, which are very likely to be the same as other background edge pixels. Direct training may optimize the model to match the edge pixels without compatibility on semantics, as the edge pixel is a very strong cue

for positive pairs. This issue could be tackled by additional mask augmentations (denoted as “Aug”) to prevent the model from using such cue. We randomly erode the foreground mask and extend the background mask so that the cue becomes random. Examples are included in the [supplementary material](#). “Aug” significantly improves the performance over direct training.

However, the performance is still much lower than using fixed foreground network. Another issue is that the foreground object images have very different appearances and distributions from regular images. And the “Direct Training” model N_f does not respect semantics very well as compared with using ImageNet pre-trained weights. As shown in Fig. 5, “tuned foreground” retrieves irrelevant categories for a sidewalk query, e.g. bottle, car. Our solution is to alternatively train N_b and N_f as shown in Fig. 6. Since there is only one network trained in one stage, the embedding features of the trained network will not have much drift. In this way, our method maintains semantic information in foreground feature, while allowing N_f to learn from data for other factors, i.e. lighting and geometry.

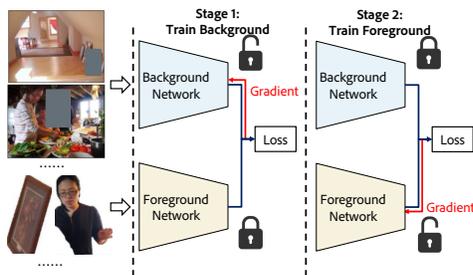


Fig. 6: The proposed alternating training strategy.

3.4 Extension to Non-box Scenarios

When the user only provides a background image without any bounding box, we adopt the following random seed and greedy search algorithm to retrieve matching foreground objects. We first sample multiple random locations. Each location is assigned with multiple bounding boxes with different aspect ratios and scales. We retrieve foreground objects based on each of the bounding boxes and re-rank all the results based on cosine similarity. We select the best one as

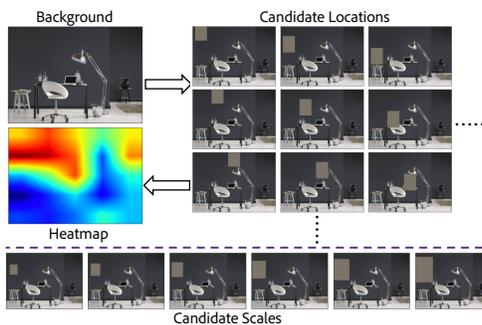


Fig. 7: Example of location/scale prediction.

the default object, then assign an initial bounding box with the same aspect ratio as the object. The area of the initial box is empirically assigned as $1/25$ of the query image. Then we apply $k \times k$ grid of locations (Fig. 7) with initial bounding box to cover the query image in a sliding window manner. By computing the similarity score between the object and background with all the

possible boxes, the box with the highest score is considered as the best location. A location heatmap is generated by interpolating the $k \times k$ score matrix to the size of query image. One can always increase k or apply further refinement with smaller stride to improve the location. The scale is then selected by applying a range of scale ratios on the initial box at the best location, as shown in Fig. 7.

4 Experiment

We first conduct experiment on CAIS [26], which is specifically annotated for compositing-aware search. Then we generate two large-scale datasets based on Pixabay [1] and Open Images [9] to demonstrate our model’s generalization ability on open-world setting.

CAIS: The dataset is generated by selecting 8 popular categories from three datasets, i.e. MS-COCO [16], PASCAL VOC 2012 [5], ADE20K [28]. The training set has 86,800 background images, and the original foreground object in each image is considered as ground-truth. The evaluation set has 80 manually selected background images and query bounding box to insert objects. About 16 ~ 140 compatible foreground objects are annotated for each background image.

Pixabay: “pixabay.com” is a stock image website containing tons of high-quality photos which are perfect for compositing. The images are highly diverse, free to use, and substantial in number. We first collect about 928,018 images, then apply object detection [21] and segmentation [12] to generate foreground objects and background images based on masks. In total, we get 5,771,912 foreground objects and 928,018 background images. But some objects are not well segmented or have extremely small size which is not likely to be used for compositing. Background images with overly large box is also not suitable as there is nothing left to tell what should be here. Therefore we only keep images with high confidence score (e.g. > 0.6) and proper bounding box size (e.g. box area in 5 ~ 50% of the area of whole image). Finally, we get 833,964 foreground and background pairs with 914 non-zero categories. They are then randomly split into training/evaluation set with portion of 90%/10%.

Open Images: Open Images originally contains about 9 million images with 9,605 trainable classes, thus is perfect for open-world evaluation. Up to 2.8 million objects from 350 categories are annotated with segmentation masks. We originally select 944,024 images with 2,686,666 objects, then apply the same filtering procedure as Pixabay based on box size only, as the mask is annotated. In total, we keep 1,374,344 background and foreground object pairs in our experiments. They are then randomly split into training/evaluation set with portion of 90%/10%. Open Images is only used in ablation study.

Implementation Details. We implement our method based on PyTorch [19] with multiprocessing to support large-scale training and evaluation. The models are trained on 8 Tesla V100 GPUs. Batch size is 40 per card for training. Following [27], we use VGG-19 [20] pre-trained on ImageNet [4] as backbone, as it achieves the best performance among different networks. All foreground

objects are padded with white pixels as square images. The original object in background image is covered with rectangle mask with average value. Both foreground and background images are resized to 224×224 and normalized with average value before feed into networks. We use Adam [8] optimizer with learning rate of 0.00008 based on linear scaling rule [6] of multi-card training. The margins of Eq. 1, 2 are set as 0.3, 0.1.

Evaluation Metrics For retrieval, we use the most widely used metric mAP (mean Average Precision) and R@k (Recall@k). We also report mAP-100 which is the mAP for top 100 retrievals as a fair comparison with constrained methods, because constrained retrieval methods do not rank object in all categories. When there is only one ground-truth (the original object) for each background query, we report R@k as the percentage of background queries whose ground-truth foreground reference appears in top k retrievals. Although other objects may also be compatible, the original one should always have a good rank among all the objects.

4.1 Comparison with Previous Methods

In this section, we compare our method with previous works on CAIS [26] and Pixabay dataset. [13,23] are infeasible for comparison because they both have very different settings, and their codes and datasets are not released. **More qualitative results are included in supplementary materials.**

CAIS: We compare our method with previous state-of-the-art methods and their variants in Table 3. “Shape” [26] ranks all the foreground images by comparing their aspect ratio with the aspect ratio of query bounding box. “RealismCNN” [29] trains a discriminator to distinguish real/fake composite images by copy-paste original and irrelevant objects into a background image. The score of the discriminator measures the realism of each composite, which is used to rank all the objects. “CFO-C Search” [26] and “CFO-D Search” [26] are two constrained search methods evaluated in unconstrained scenarios. “CFO-C Search” first trains a classifier to specify the category, then apply constrained retrieval only from this class. The results will be completely wrong if the classification fails. “CFO-D Search” first apply constrained search to retrieve 100 samples from each category. Then it adopts the discriminator of “RealismCNN” to re-rank all these retrievals by compositing with each background. Note that this is computationally expensive and not scalable if there are hundreds of classes. “UFO Search” [27] applies the discriminator to generate extra positive samples.

In Table 3, we show the mAP-100 per class and the overall performance. The proposed method significantly (+5.33%) outperforms previous state-of-the-art methods on overall performance. It also achieves better performance than UFO [27] on most of the categories. “CFO-C Search” and “CFO-D Search” may performs well on two categories, but they are both not scalable, because it is infeasible to train one model for each category or scan all objects with discriminator for each background on large-scale datasets with hundreds of categories. The results indicate that geometry and lighting awareness can help the retrieval

Table 3: Comparison with previous works on CAIS in terms of mAP-100. “†” denotes constrained methods with multiple models, which are not scalable in practice.

Method	Boat	Bottle	Car	Chair	Dog	Painting	Person	Plant	Overall
Shape [26]	7.47	1.16	10.40	12.25	12.22	3.89	6.37	8.82	7.82
RealismCNN [29]	12.33	7.19	7.55	1.81	7.58	6.45	1.47	12.74	7.14
†CFO-C Search [26]	57.48	14.24	18.85	21.61	38.01	27.72	47.33	20.20	30.68
†CFO-D Search [26]	55.48	8.93	24.10	18.16	57.82	21.59	27.66	23.13	29.61
UFO Search [27]	59.73	21.12	36.63	19.27	36.51	25.84	27.11	31.19	32.17
Ours	70.58	19.41	40.22	24.17	37.81	28.20	44.72	34.91	37.50

Table 4: Retrieval performance of selected classes in terms of Recall@10 (%) on Pixabay.

Category	Majority #samples \approx 5000				Medium #samples \approx 50				Minority #samples $<$ 5
	Person	Flower	Birds	Vehicle	Cell Phone	Mandarin Orange	Christmas Tree	Boiled Egg	Last 50 Classes
UFO [27]	3.84	8.00	6.74	8.41	5.71	36.36	2.94	6.06	8.00
Ours	19.36	28.55	21.11	26.83	20.00	63.64	20.59	30.30	24.00

in general unconstrained retrieval.

Pixabay: Since none of the previous works conduct experiments on large-scale open-world setting, we implement the state-of-the-art unconstrained search method (UFO [27]) on Pixabay and provide detailed comparison. Table 4 shows the Recall@10 on selected categories in Pixabay. We select 4 majority and 4 medium classes, with about 5000 and 50 samples per class in evaluation set (Note that training set is 9 times larger). Then we compute the average Recall@10 on the last 50 long-tail classes. The proposed method significantly outperforms UFO [27] on all the classes. We also show the overall R@k on Pixabay in Table 5. The proposed method significantly outperforms UFO, which means the proposed method generalizes well on large-scale open-world setting.

Table 5: Overall retrieval accuracy in terms of Recall@k (%) on Pixabay.

Method	R@1	R@5	R@10	R@1%
UFO [27]	2.04	6.66	10.24	61.76
Ours	7.75	20.13	28.20	85.61

4.2 Ablation Study

Geometry and Lighting Sensitivity: Since our goal is to gain awareness on geometry and lighting, a good model should be discriminative/sensitive to the change of geometry and lighting on foreground images. We conduct experiment on Pixabay [1] and Open Images [9] to verify this point. We randomly select 2,000 foregrounds along with their corresponding background images. Then for

Table 6: Evaluation of sensitivity (discriminative ability) to lighting and geometry transformation. The “sensitivity” denotes the squared L2 distance between the normalized embedding of original and transformed foreground images.

Ablation	Lighting			Geometry				
	Sensitivity(\uparrow)	R@5	R@10	R@15	Sensitivity(\uparrow)	R@5	R@10	R@15
Pixabay								
UFO [27]	0.27	53.70	67.90	75.50	0.39	58.25	69.90	77.95
Baseline	0.27	53.40	67.30	75.35	0.39	58.10	69.70	77.65
No Contrastive	0.51	55.60	70.70	79.70	0.72	61.30	74.30	82.75
Overall	0.57	60.55	74.70	82.85	1.12	98.55	99.45	99.70
Open Images								
Baseline	0.24	51.70	65.45	74.60	0.40	60.10	71.80	78.45
No Contrastive	0.53	54.80	71.55	80.10	0.98	71.70	82.60	89.10
Overall	0.56	59.35	73.90	81.80	1.58	99.50	99.75	99.90

each foreground, we use the geometry and lighting transformations in Sec. 3.2 to generate transformed image. In total, we generate 100 transformed images for each foreground image, 50/50 for geometry/lighting transformation. Then we rank the original one along with the 50 geometry or lighting transformed images to compute the Recall@k in Table 6. We also measure the discriminative ability as the sensitivity to these transformations, i.e. the square euclidean distance between normalized embedding features of the original and transformed foregrounds. With L2 normalization, the square euclidean distance is $d = 2 - 2s$, where s is the cosine similarity. Therefore higher sensitivity value means larger distance between the features of original and transformed objects, thus indicates stronger ability on distinguishing geometry or lighting transformations.

In Table 6, the overall proposed method has much higher sensitivity to



Fig. 8: Qualitative comparison on Pixabay and Open Images. Each of the four examples contains three rows of retrieval results, corresponding to: “Baseline” (1st row), “No Contrastive” (2nd row), “Overall” model (3rd row).

both geometry and lighting transformations for both datasets. UFO [27] and “Baseline” both use fixed foreground features, thus have a low sensitivity. “No Contrastive” removes the proposed contrastive transformation but keeps the alternating training, which has a lower sensitivity and retrieval performance than “Overall”, but better than “Baseline”. The results demonstrate the effectiveness of both alternating training and contrastive learning on improving discriminative ability. We also show qualitative results of our ablations in Fig. 8 with diverse scenes, viewpoints and lighting. Our overall method better respects the geometry and lighting than baseline methods.

Components Ablations: We conduct detailed ablation study on CAIS to show the effectiveness of each component. “Direct Training” mean directly training two networks with loss in Eq. 1. “Aug” denotes direct training along with mask augmentations in Sec. 3.3 to prevent the model from overfitting on edge pixels. “Fix+Aug” is also used as “Baseline” in other tables. It uses a fixed foreground network with mask augmentations. Our “Overall” method adopt both alternating training and contrastive learning with mask augmentations. “No Alternating” removes alternating training from “Overall”, and “No Contrastive” removes the contrastive learning loss in Eq. 2. As shown in Table 7, “Overall” performs the best among all the ablations, indicating the effectiveness of each component.

Table 7: Ablation study on components.

Ablations	mAP	mAP-100
Direct Training	17.93	24.02
Aug	23.45	31.78
Fix+Aug	28.65	30.33
No Alternating	29.99	32.13
No Contrastive	31.20	36.30
Overall	32.67	37.50

Table 8: Ablation study on different alternating training strategies.

Ablations	mAP	mAP-100
Zero Round	28.65	30.33
One Round	31.20	36.30
Two Rounds	30.35	36.22
Reverse Order	28.82	36.14

Table 9: Ablation study on different transformations.

Ablations	mAP	mAP-100
No Contrastive	31.20	36.30
Geometry	32.11	36.68
Geometry+Color	30.80	35.06
Geometry+Lighting	32.67	37.50

Training Strategy Ablations: We conduct ablation study on different alternating training strategies and select the best strategy for our overall method. All these ablations adopt mask augmentations without using contrastive learning. “Zero Round” means no alternating training, which is equivalent to the “Baseline”. “One Round” ($\text{argmin}_{N_b} \mathcal{L}$, then $\text{argmin}_{N_f} \mathcal{L}$) trains foreground network after the background network is trained. “Two Round” trains the background network for another round based on the “One Round” model. “Reverse Order” ($\text{argmin}_{N_f} \mathcal{L}$, then $\text{argmin}_{N_b} \mathcal{L}$) first trains the foreground network, then trains the background network. Table 8 shows that “One Round” strategy performs the best among all the strategies, indicating that training more rounds is not beneficial. We thus use one round for our overall method.

Transformation Ablations: To show the effectiveness of each transformation, we train our method with different transformations in contrastive learning. “No Contrastive” removes contrastive learning loss in Eq. 2. “Geometry” only adopts geometry transformation in Fig. 2, and “Geometry+Lighting” means each foreground is applied with both transformations. “Geometry+Color” uses linear color jittering instead of our non-linear lighting transformation on the top of “Geometry”. In Table 9, the “Geometry+Lighting” outperforms other ablations, indicating that both transformations are necessary in our framework, and linear color jittering does not well simulate lighting changes.

4.3 Location and Scale Prediction

Since CAIS has annotated bounding boxes on intact background images, we evaluate our location and scale prediction method based on these boxes for non-box scenarios. We compare the proposed method with random strategy and “Baseline” model. For each background image, we randomly select 5 of its annotated compatible foregrounds for evaluation and compute the average value to report. Qualitative results are included in supplementary material.

For location prediction, given a background along with its annotated foreground object, we first apply 10 grid search with sliding window method in Sec. 3.4. To evaluate location and scale separately, the scale here is fixed as the scale of annotated box. Then bilinear interpolation is adopted to generate the heatmap, as shown in Fig. 9. The heatmap value is then normalized with the maximum and minimal value in the grid matrix, so that all values are in $[0, 1]$. Higher value means the model assigns high compatibility on the corresponding location. Note that the annotated box location may not be the only good location, other boxes with zero IOU (Intersection Over Union) can also have compatible locations, e.g. all the non-occluded locations on the wall in Fig. 9. However, the ground-truth location should always have a good compatibility score as compared with scores of other locations. Therefore, we compute the similarity score on the annotated location and normalize it with the same normalization factors as the heatmap. A good location prediction model should give a high normalized similarity (NS) on the annotated locations. For “Random” strategy, the normalized similarity is randomly selected in $[0, 1]$ since the heatmap is random.

For scale prediction, given a background along with its annotated foreground

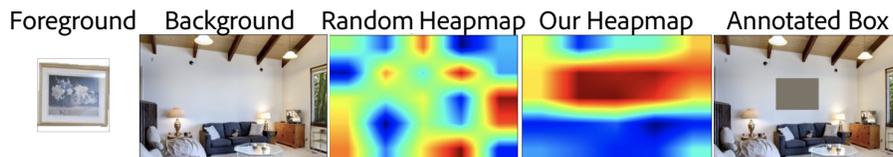


Fig. 9: Example heatmaps for location prediction evaluation.

Table 10: Evaluation on location and scale prediction. NS (normalized similarity) denotes the cosine similarity of the ground truth location normalized by the maximum and minimum similarity of the other sliding window locations.

Method	Location			Scale		
	$NS > 0.99$	$NS > 0.95$	$NS > 0.9$	$IOU > 0.9$	$IOU > 0.75$	$IOU > 0.5$
Random	5.50	9.25	12.25	4.25	14.50	38.25
Baseline	15.25	25.25	38.25	5.75	19.75	48.50
Ours	22.25	31.50	39.00	6.75	24.25	51.75

object, all methods use the annotated ground-truth locations as the center of the box. With a fixed initial size with $1/25$ area of the whole background image, we assign 9 different scales by $scale = 1.2^{k-4}$, $k = 0, 1, \dots, 8$. The box size is then multiplied by the scale to get 9 candidate boxes. ‘‘Random’’ strategy selects one scale randomly, while the other two methods use the similarities between the background with boxes and the foreground to rank all the scales. We then compute the IOU between the predicted box and the annotated box.

As shown in Table 10, we report the percentage of correct predictions with different thresholds on two evaluation metrics, i.e. NS and IOU. Our method performs the best on both location and scale prediction, indicating the superiority of the proposed method on non-box scenarios.

5 Limitations and Potential Social Impact

Current framework only uses 2D transform for contrastive learning, which is efficient and scalable, but not as accurate as 3D transform. This limitation could be addressed in the future by adopting 3D transform as better pre-processing w/o changing the proposed learning framework. Another limitation is that the search space is currently bounded by the gallery, and hence there may not be any perfectly compatible object images in the database even with the large-scale open-world database setting we adopt in our work. One solution is to augment the search space by allowing transformation of objects in the database.

Retrieval-based compositing could be indirectly used to generate fake images, but we can make sure all the gallery object images do not have ethical or privacy issues. Also, fake images usually have incompatible geometry or lighting conditions which can be detected using our model.

6 Conclusion

We propose a novel unconstrained foreground search method for compositing (GALA), as the first to learn geometry-and-lighting awareness from large-scale data for real-world scenarios. GALA achieves state-of-the-art results on CAIS and Pixabay dataset. It also tackles non-box scenarios by automatic location-scale prediction, which is not explored by previous works.

References

1. <https://pixabay.com/>, <https://pixabay.com/> 3, 8, 10
2. Azadi, S., Pathak, D., Ebrahimi, S., Darrell, T.: Compositional gan: Learning image-conditional binary composition. *International Journal of Computer Vision* **128**(10), 2570–2585 (2020) 1
3. Chechik, G., Sharma, V., Shalit, U., Bengio, S.: Large scale online learning of image similarity through ranking. *Journal of Machine Learning Research* **11**(3) (2010) 5
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009) 6, 8
5. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International journal of computer vision* **88**(2), 303–338 (2010) 8
6. Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch sgd: Training imagenet in 1 hour. arXiv preprint arXiv:1706.02677 (2017) 9
7. Jiang, Y., Zhang, H., Zhang, J., Wang, Y., Lin, Z., Sunkavalli, K., Chen, S., Amirghodsi, S., Kong, S., Wang, Z.: Ssh: A self-supervised framework for image harmonization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4832–4841 (2021) 1
8. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) 9
9. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., et al.: The open images dataset v4. *International Journal of Computer Vision* **128**(7), 1956–1981 (2020) 3, 8, 10
10. Lalonde, J.F., Hoiem, D., Efros, A.A., Rother, C., Winn, J., Criminisi, A.: Photo clip art. *ACM transactions on graphics (TOG)* **26**(3), 3–es (2007) 3
11. Lee, D., Liu, S., Gu, J., Liu, M.Y., Yang, M.H., Kautz, J.: Context-aware synthesis and placement of object instances. arXiv preprint arXiv:1812.02350 (2018) 4
12. Lee, Y., Park, J.: Centermask: Real-time anchor-free instance segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13906–13915 (2020) 8
13. Li, B., Zhuang, P.Y., Gu, J., Li, M., Tan, P.: Interpretable foreground object search as knowledge distillation. In: European Conference on Computer Vision. pp. 189–204. Springer (2020) 2, 3, 4, 9
14. Li, X., Liu, S., Kim, K., Wang, X., Yang, M.H., Kautz, J.: Putting humans in a scene: Learning affordance in 3d indoor environments. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12368–12376 (2019) 4
15. Lin, C.H., Yumer, E., Wang, O., Shechtman, E., Lucey, S.: St-gan: Spatial transformer generative adversarial networks for image compositing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9455–9464 (2018) 1
16. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014) 3, 8
17. Niu, L., Cong, W., Liu, L., Hong, Y., Zhang, B., Liang, J., Zhang, L.: Making images real again: A comprehensive survey on deep image composition. arXiv preprint arXiv:2106.14490 (2021) 1

18. Pandey, R., Escolano, S.O., Legendre, C., Haene, C., Bouaziz, S., Rhemann, C., Debevec, P., Fanello, S.: Total relighting: learning to relight portraits for background replacement. *ACM Transactions on Graphics (TOG)* **40**(4), 1–21 (2021) [6](#)
19. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32**, 8026–8037 (2019) [8](#)
20. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014) [8](#)
21. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9627–9636 (2019) [8](#)
22. Tsai, Y.H., Shen, X., Lin, Z., Sunkavalli, K., Lu, X., Yang, M.H.: Deep image harmonization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3789–3797 (2017) [1](#)
23. Wu, Z., Lischinski, D., Shechtman, E.: Fine-grained foreground retrieval via teacher-student learning. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 3646–3654 (2021) [2](#), [3](#), [4](#), [6](#), [9](#)
24. Zhang, H., Zhang, J., Perazzi, F., Lin, Z., Patel, V.M.: Deep image compositing. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 365–374 (2021) [1](#)
25. Zhang, L., Wen, T., Min, J., Wang, J., Han, D., Shi, J.: Learning object placement by inpainting for compositional data augmentation. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII* 16. pp. 566–581. Springer (2020) [4](#)
26. Zhao, H., Shen, X., Lin, Z., Sunkavalli, K., Price, B., Jia, J.: Compositing-aware image search. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 502–516 (2018) [1](#), [2](#), [3](#), [8](#), [9](#), [10](#)
27. Zhao, Y., Price, B., Cohen, S., Gurari, D.: Unconstrained foreground object search. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2030–2039 (2019) [2](#), [3](#), [6](#), [8](#), [9](#), [10](#), [11](#), [12](#)
28. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 633–641 (2017) [8](#)
29. Zhu, J.Y., Krahenbuhl, P., Shechtman, E., Efros, A.A.: Learning a discriminative model for the perception of realism in composite images. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 3943–3951 (2015) [1](#), [9](#), [10](#)