

Learning Co-segmentation by Segment Swapping for Retrieval and Discovery

Xi Shen¹, Alexei A. Efros², Armand Joulin³, and Mathieu Aubry⁴

^{1,4}LIGM (UMR 8049), École des Ponts ParisTech

²University of California, Berkeley

³Facebook AI Research

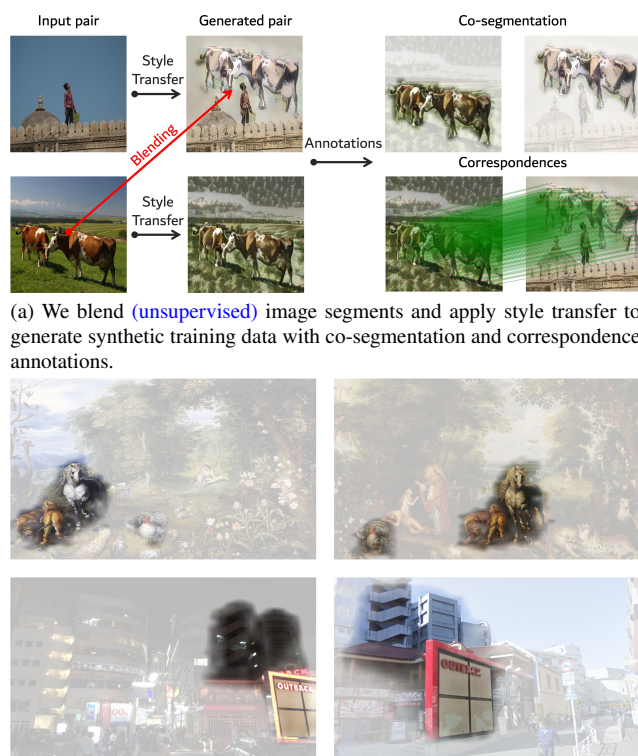
Abstract

The goal of this work is to efficiently identify visually similar patterns in images, e.g. identifying an artwork detail copied between an engraving and an oil painting, or recognizing parts a night-time photograph visible in its daytime counterpart. Lack of training data is a key challenge for this co-segmentation task. We present a simple yet surprisingly effective approach to overcome this difficulty: we generate synthetic training pairs by selecting segments in an image and copy-pasting them into another image. We then learn to predict the repeated region masks. We find that it is crucial to predict the correspondences as an auxiliary task and to use Poisson blending and style transfer on the training pairs to generalize on real data.

We analyse results with two deep architectures relevant to our joint image analysis task: a transformer-based architecture and Sparse Nc-Net, a recent network designed to predict coarse correspondences using 4D convolutions. We show our approach provides clear improvements for artwork details retrieval on the Brueghel dataset and achieves competitive performance on two place recognition benchmarks, Tokyo247 and Pitts30K. We also demonstrate the potential of our approach for unsupervised image collection analysis by introducing a spectral graph clustering approach to object discovery and demonstrating it on the object discovery dataset of [49] and the Brueghel dataset. Our code and data are available at <http://imagine.enpc.fr/~shenx/SegSwap/>.

1. Introduction

Identifying repeated patterns lies at the very heart of the computer vision problem, and is a key component of Intelligence itself. Yet, in practice, our best methods for performing such a fundamental task often leave a lot to be desired. While we now have good methods for discovering *exact* pattern matches (used extensively to find copyright infringe-



(b) Application on artworks and day/night photographs

Figure 1. We learn co-segmentation from synthetic images. Models trained on our data are effective for many applications including art detail retrieval, place recognition and object discovery.

ments), as well as approximate matches of *salient* objects (see object discovery and co-segmentation approaches in Section 2), detecting visually similar details within a larger visual context remains surprisingly difficult.

Spotting the repetition of visual detail has several applications. Identifying copied details in artworks allows art historians to discover influences, find provenance, and establish authorship [54]. Finding repeated details can

boost performances in visual localisation for place recognition [20]. Reliable pair-wise image co-segmentation and correspondence identification could also enable object discovery in image collections [11].

In this paper, we show it is possible to learn to detect repeated visual patterns – jointly predicting co-segmentation and correspondences – without any human-label. Instead, we generate synthetic correspondence pairs via automatic data augmentation. More precisely, we use a “segment swapping” approach, where we blend image segments in a random background using Poisson blending and apply style transfer to the resulting image to obtain challenging training image pairs (Figure 1a). On the generated image pairs, we have access to the ground-truth matchability masks as well as the correspondences which we use as supervisions for training a network (Figure 1a). Surprisingly, we find that models trained on such a dataset generalize well to real data and can be directly used for art detail retrieval and place recognition (Figure 1b). We consider using both annotated object segments and unsupervised segments for this data generation process. We obtain unsupervised segments by applying the Bilateral Solver [5] to the recent unsupervised object detection LOST [56], based on the self-supervised DINO [7] features, which we show significantly improves performance on unsupervised saliency detection benchmarks. Experimental results show that using unsupervised segments produces slightly lower but comparable results compared to using COCO [36] instance segments.

We experimented with two architectures which we adapt to predict co-segmentation and correspondences in image pairs: the recent Sparse Nc-Net [46] architecture, designed for predicting image coarse correspondences, and an architecture based on Transformers [66] which we refer to as cross-image transformer. We analyze the effectiveness of our data generation process, architectures and training strategy on two types of tasks. First, we perform retrieval tasks using the predicted pair-wise co-segmentation masks and correspondences. We show clear performance improvement for artwork details retrieval on the Brueghel [54] dataset and results comparable to state of the art for visual localization on two challenging place recognition benchmarks, Tokyo247 [61] and Pitts30K [62]. This last result is especially impressive, since these benchmarks are very competitive, and many dedicated methods leveraging geo-referenced images or real correspondence for supervision have been proposed. On the contrary, our approach is generic and relies solely on our synthetic “segment swapping” training. Second, we make use of the predicted masks and correspondences to build a candidate correspondence graph and introduce an approach to perform discovery in image collections with spectral clustering [32, 40]. We demonstrate results on par with state-of-the-art on the standard co-segmentation dataset of [49] and show qualitative




Figure 2. Data generation by “segment swapping”. Instead of directly copy-pasting (Figure 2c) a segment from a source image (Figure 2a) to a background (Figure 2b), we use Poisson blending [42] and add style transfer [23] to the result (Figure 2d).

results on the challenging Brueghel [54] dataset. Our code and data are available at <http://imagine.enpc.fr/~shenx/SegSwap/>.

2. Related work

Learning correspondences between different images. SIFT-Flow [37] was an early method that aligns visually distinct scenes by incorporating visual features, such as SIFT, into optical flow-style approaches. More recently, many deep learning based approaches have been developed to predict correspondences from correlations of input features [39, 45, 53, 63–65]. Of particular interest, architectures based on attention mechanisms and Transformers [66] have been introduced to predict image correspondences. Super-Glue [50] is an attention-based graph neural network for key-point matching. Closer to this work, COTR [26] is a sequence-to-sequence transformer architecture that takes an image and 2D coordinates of a query points as inputs to predict correspondences. Finally, LoFTR [58] adopts a coarse-to-fine approach to matching with a transformer encoder. As opposed to our work, these transformer-based methods are trained on a large dataset with ground-truth poses and depth while we only train on a synthetic dataset. Additionally, our model is only composed of an encoder. More importantly, it outputs an accurate mask of the common regions along with the correspondences.

Learning correspondences without annotated data.

There is a large body of work that use synthetic images [15] or images with synthetic deformations [39, 45, 52, 63] to learn correspondences without real annotated training data.

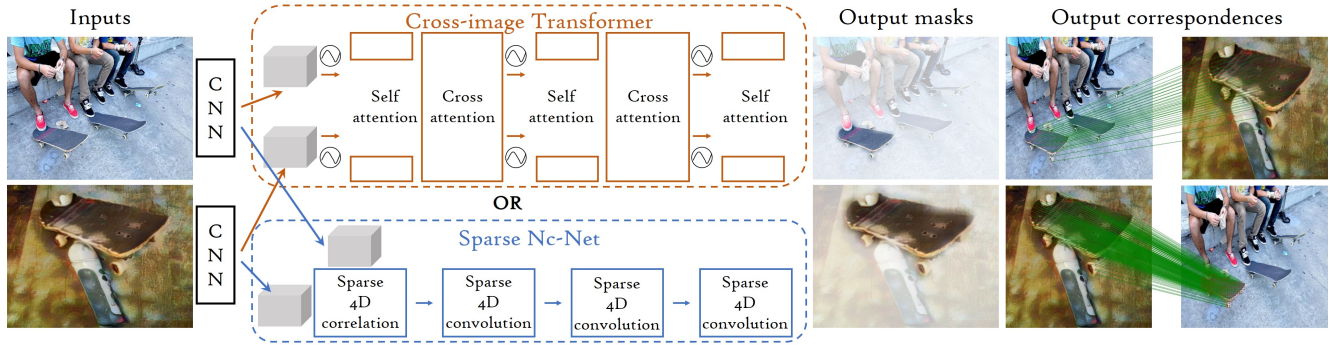


Figure 3. We train our cross-image transformer or Sparse Nc-Net [46] on the generated pairs. Both networks jointly predict masks and correspondences.

However, these approaches do not try to identify the matchable regions, which is essential to discover visual details. Some other approaches train directly on real images using proxy signals for correspondences, such as photometric or cycle consistency [24, 53, 64, 65, 72, 76]. Again, they focus on the quality of the correspondences and are not designed to predict matchable regions in vastly different images. On the contrary, the core of our approach is to discover these similar regions. This makes our approach particularly suited for retrieval tasks. We are also inspired by related data augmentation techniques, specifically, the CopyPaste augmentation used by [18] for instance segmentation and the stylised-ImageNet augmentation used in [17] to increase shape bias in neural networks.

Object discovery and co-segmentation. There is a wide variety of approaches aiming at discovering objects and their location from unlabelled images. Many methods [12, 59, 68, 69] use bounding box proposals and formulate the object discovery as an optimization problem. This relies on the quality of proposals which are typically not adapted for non-photorealistic data, such as artworks. Other methods [11, 22, 33, 35, 48, 49, 60, 67, 75] focus on predicting masks of salient objects. Some [33, 35, 75] require foreground masks for training, while others [11, 22, 27, 28, 33, 67] are designed to segment common repeated objects in an image collection. These approaches make assumptions about the frequency of appearance of an object, while, in many practical scenarios, repeated objects are rare and discovering them is about seeking a needle in a haystack [54]. Our approach is related to [49, 60], as we both leverage dense correspondences to discover objects. As opposed to our work, [60] focuses on a single pair of images while we also show results over an entire collection of images. [49] makes the assumption that the common object is also the most salient in the image. This works well with images from internet queries but does not apply to artworks where the common object can be a detail in a richer scene.

3. Co-segmentation by segment swapping

In Section 3.1, we introduce our “segment swapping” data generation process (Figure 2). We then present in Section 3.2 the two architectures we use (Figure 3). We discuss our loss and training strategy in Section 3.3.

3.1. Training data generation by segment swapping

Training pairs generation We generate training pairs using images from the MS COCO dataset [36]. We first sample a source image, from which we extract one or two segments (as explained below). We then build the target image by applying geometric transformations to the segments and blending them into a random background image using Poisson blending [42]. The geometric transformations include rotation, translation, scaling, and thin-plate spline (TPS). A style augmentation is then performed on both the source and target images using an AdaIN [23] model trained on the Brueghel dataset [54]. An example of training pair can be seen in Figure 2 and we provide more examples of training samples in the supplementary material [2].

Segments definition The simplest way to define segments for our data generation process is to use annotated object segments. For our experiments, we used the instance annotations from COCO [36]. We compare this approach to a completely unsupervised segment extraction, which we defined using the following strategy: (1) given an image, we employ LOST [56] to obtain an object segmentation, which is irregular and coarse on the boundary; (2) we refine the object segmentation using the Bilateral Solver [5].

We quantitatively validate our unsupervised segments on unsupervised saliency detection benchmarks in Section 5.2. More details about the unsupervised procedure and example of unsupervised segments and the associated image pairs are available in the supplementary material [2].

3.2. Architectures

Our networks take as input a source image \mathbf{I}^s and a target image \mathbf{I}^t , from which features maps \mathbf{F}^s and \mathbf{F}^t of spatial dimension $W \times H$ are extracted by a feature extraction backbone network. These feature maps are then processed either by our cross-image transformer or our modified sparse Nc-Net [46] architecture to predict both the masks of the repeated objects in the source and target images, $\mathbf{M}^s \in [0, 1]^{W \times H}$ and $\mathbf{M}^t \in [0, 1]^{W \times H}$ respectively, and the correspondences both from source to target $\mathbf{C}^{s \rightarrow t}$ and target to source $\mathbf{C}^{t \rightarrow s}$. $\mathbf{C}^{s \rightarrow t}$ and $\mathbf{C}^{t \rightarrow s}$ are represented as matrices of size $W \times H \times 2$. To simplify notation, we sometime use the masks as continuous 2D functions, which in practice is done by performing bilinear interpolation.

Cross-image transformer We built an architecture based on the classic transformer encoder [66] which alternates multi-headed attention and fully connected feed-forward networks (FFN) blocks. The FFN blocks contain two layers with a ReLU non-linearity. Similar to [50], we use two types of attention layers: one is the standard self-attention (SA) layer, the other one is a cross-attention (CA) layer where the attention is only computed between features from different images. We include the same 2D positional encoding as DeTR [6] on top of the feature map before SA. Our transformer alternates these two types of attention layers as shown in Figure 3, with a total of five attention and FFN blocks. Each attention layer has 2 heads and the dimension of the features is 256. Our last layer is followed by a sigmoid and has three outputs, that we interpret as masks and correspondences for each image. We provide an ablation study of this architecture in the supplementary material [2].

Sparse Nc-Net Nc-Net [46] is designed to learn coarse correspondences under weak supervision. It takes as input the correlations between \mathbf{F}^s and \mathbf{F}^t , seen as a 4D volume of affinities $\mathcal{A}_{input} \in \mathbb{R}^{W \times H \times W \times H}$, and processes them with 4D convolutions. The final 4D convolution predicts affinities $\mathcal{A}_{pred} \in \mathbb{R}^{W \times H \times W \times H}$, on which softmax functions are applied in dimensions corresponding to source and target giving $\mathcal{A}_{pred}^s(i, j, k, l) = \frac{\exp(\mathcal{A}_{pred}(i, j, k, l))}{\sum_{k, l} \exp(\mathcal{A}_{pred}(i, j, k, l))}$ and $\mathcal{A}_{pred}^t(i, j, k, l) = \frac{\exp(\mathcal{A}_{pred}(i, j, k, l))}{\sum_{i, j} \exp(\mathcal{A}_{pred}(i, j, k, l))}$. We use the maxima of these affinities as source and target masks, i.e., $\mathbf{M}^s(i, j) = \max_{k, l} \mathcal{A}_{pred}^s(i, j, k, l)$ and $\mathbf{M}^t(k, l) = \max_{i, j} \mathcal{A}_{pred}^t(i, j, k, l)$. Correspondences are obtained with soft-argmax:

$$\begin{aligned} \mathbf{C}^s(i, j) &= \left(\sum_{k, l} \frac{k}{W} \mathcal{A}_{pred}^s(i, j, k, l), \sum_{k, l} \frac{l}{H} \mathcal{A}_{pred}^s(i, j, k, l) \right) \\ \mathbf{C}^t(k, l) &= \left(\sum_{i, j} \frac{i}{W} \mathcal{A}_{pred}^t(i, j, k, l), \sum_{i, j} \frac{j}{H} \mathcal{A}_{pred}^t(i, j, k, l) \right) \end{aligned} \quad (1)$$

Since 4D convolutions are computational heavy, we instead use sparse 4D convolutions with the same architecture as Sparse Nc-Net [46].

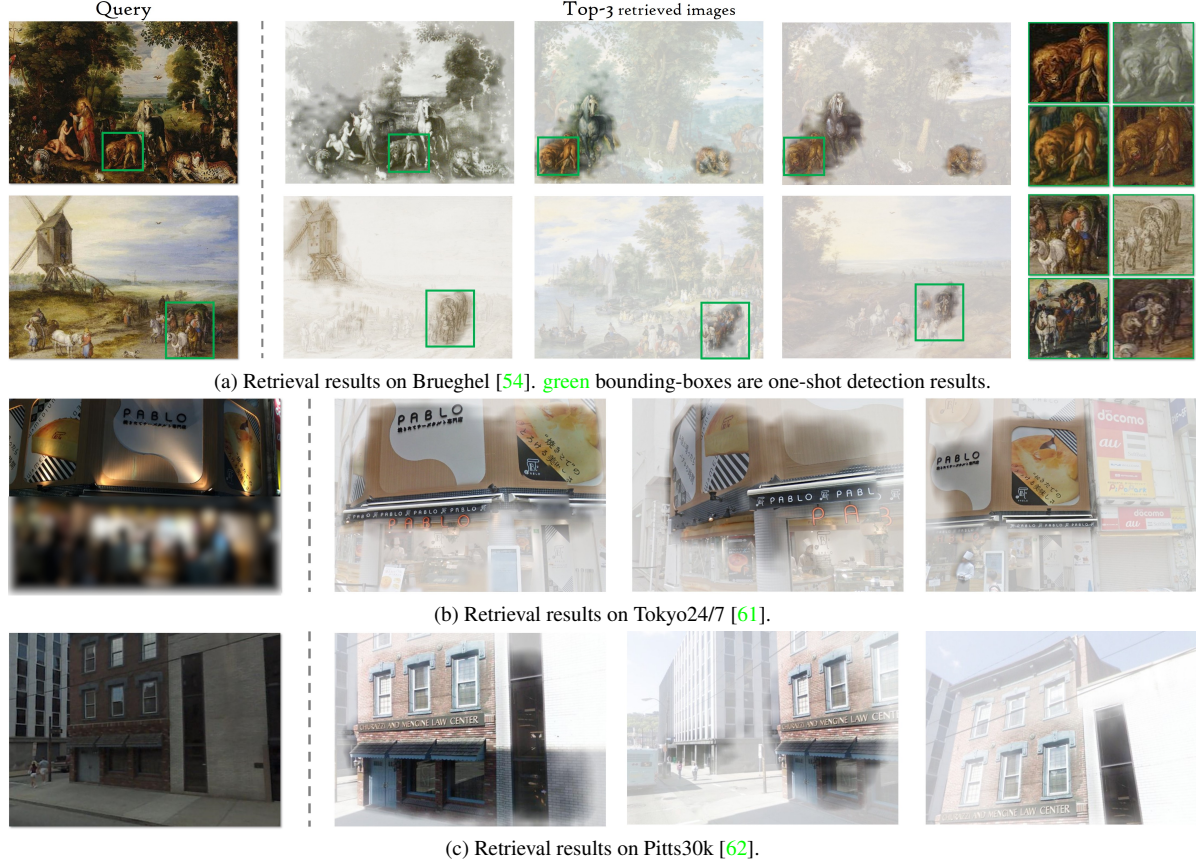
3.3. Loss and training

On our synthetic training data we have access to the ground truth masks \mathbf{M}_{gt}^s and \mathbf{M}_{gt}^t and ground truth correspondences $\mathbf{C}_{gt}^{s \rightarrow t}$ and $\mathbf{C}_{gt}^{t \rightarrow s}$ on the source and target images. Our loss is the sum of two symmetric terms for source and target, for simplicity we write only the source loss \mathcal{L}^s . It includes a cross-entropy (CE) loss on the predicted mask \mathcal{L}_m and the transported mask \mathcal{L}_{tm} , as well as a regression loss \mathcal{L}_{corr} on the correspondences:

$$\begin{aligned} \mathcal{L}^s &= \underbrace{CE(\mathbf{M}_{gt}^s, \mathbf{M}^s)}_{\mathcal{L}_m} + \underbrace{CE(\mathbf{M}_{gt}^s, \mathbf{M}^t(\mathbf{C}^{s \rightarrow t}))}_{\mathcal{L}_{tm}} \\ &+ \underbrace{\eta \frac{1}{\sum_{i, j} \mathbf{M}_{gt}^s(i, j)} \sum_{i, j} \mathbf{M}_{gt}^s(i, j) \|\mathbf{C}^{s \rightarrow t}(i, j) - \mathbf{C}_{gt}^{s \rightarrow t}(i, j)\|}_{\mathcal{L}_{corr}} \end{aligned} \quad (2)$$

where i and j correspond to the feature coordinates, η is a scalar hyper-parameter, and $CE(\mathbf{M}_{gt}, \mathbf{M}) = -\frac{1}{W \times H} \sum_{i, j} \mathbf{M}_{gt}(i, j) \log(\mathbf{M}(i, j)) + (1 - \mathbf{M}_{gt}(i, j)) \log(1 - \mathbf{M}(i, j))$. Note that this loss is computed both for positive pairs (source and target pairs generated by segment swapping) and negative pairs (sampled from two different pairs, without repeated objects) for which $\mathbf{M}_{gt}^s = \mathbf{M}_{gt}^t = 0$ and by convention $\mathcal{L}_{corr} = 0$.

Implementation details We implement our approach using the Pytorch library. We use as backbone features the *conv4* features of a ResNet-50 [21] trained on ImageNet [13] with MOCO-v2 [9]. We freeze the backbone during the training, as learning backbone features leads to overfitting on the synthetic training set. For all the experiments, we optimise the loss defined in Equation 2 with $\eta = 8$ and use the Adam optimiser [30] with momentum terms $\beta_1 = 0.5$ and $\beta_2 = 0.999$. At each iteration, we sample 5 positive and 15 negative pairs. For the transformer architecture, after training 200k iterations with learning rate $2e-4$, we train with hard negative pairs and learning rate $1e-5$ for 5k iterations. Hard negatives are obtained by sampling a pool of $N_{pool} = 500$ images from different synthetic pairs, computing predicted masks for all the pairs of images in the pool, and keeping those with mask prediction higher than a threshold $\tau = 0.04$ in a hard negative pair pool for $K_{hard} = 1000$ iterations of training. For Sparse Nc-Net [46] training 200k iterations with learning rate $2e-4$ without hard negative mining leads to the best performance. The entire trainings of the transformer and Sparse Nc-Net [46] take approximately 30 hours and 15 hours respectively on a single GPU Tesla-V100-16GB. An ablation



(a) Retrieval results on Brueghel [54]. green bounding-boxes are one-shot detection results.

(b) Retrieval results on Tokyo24/7 [61].

(c) Retrieval results on Pitts30k [62].

Figure 4. Visual results for retrieval on different datasets. For each query image (1st column), we show its 3 most similar images with the predicted masks as transparency. For Brueghel [54], we also show the detection results.

study of the architectures and more training details are provided in the supplementary material [2].

4. Application to Image Retrieval

In this section, we show how our model can be used for retrieval tasks. We first explain how we use it to compute an image similarity score in Section 4.1. We then present experimental results in Section 4.2, including art detail retrieval on the Brueghel dataset [54], place recognition on Pitts30k [62] and Tokyo 24/7 [61], validation of our unsupervised segments and ablation study. More visual results and ablation studies are in the supplementary material [2].

4.1. Score between a pair of images

We propose the following score \mathcal{S} to measure the similarity between a pair of images based on predicted correspondences and masks. \mathcal{S} is the sum of weighted local features similarities, where our predicted correspondences are used to associate features and the weight \mathbf{M}_{joint}^s is the product of the source and transported target mask

$$\mathbf{M}_{joint}^s(i, j) = \mathbf{M}^t(\mathbf{C}^{s \rightarrow t}(i, j)) \mathbf{M}^s(i, j):$$

$$\mathcal{S}(\mathbf{I}^s, \mathbf{I}^t) = \sum_{i, j} \underbrace{\mathbf{M}_{joint}^s(i, j)}_{\text{Mask}} \underbrace{\cos(\mathbf{F}^s(i, j), \mathbf{F}^t(\mathbf{C}^{s \rightarrow t}(i, j)))}_{\text{Feat. similarity}} \quad (3)$$

Ablations in supplementary material [2] show that the mask term is the key part of this score and that the feature similarity term provides an additional small boost.

4.2. Experiments

Qualitative results on our different datasets can be seen in Figure 4. The predicted masks, shown with transparency, are able to capture repeated regions even in challenging cases, such as large difference of scale, viewpoints, lighting conditions and depiction styles. More visual results are provided in the supplementary material [2].

Art detail retrieval We evaluate our approach on the Brueghel dataset [1, 54] in Table 1. Our score allows us to directly retrieve images from a selected query detail. To further compare with the detection performance in Art-

Feat. + Methods	mAP	
	Retrieval	Det.(IoU > 0.3)
Shen et al. [54] + cos [54]	75.5	75.3
Shen et al. [54] + discovery [54]	76.6	76.4
MocoV2 [9] + cos [54]	79.0	78.7
MocoV2 [9] + discovery [54]	80.8	79.6
Ours + Unsupervised segments		
Transformer	81.8	79.4
Sparse-Ncnet	82.8	73.4
Ours + COCO segments [36]		
Transformer	84.4	81.8
Sparse-Ncnet	83.3	73.7

Table 1. Art detail retrieval and detection on Brueghel [54]. For detection, we employ ArtMiner (Brueghel [54] + cos [54]) as a post-processing and reports results with IoU > 0.3 [54]

Miner [54], we crop a 320×320 patch around the predicted regions and use ArtMiner [54] as a post-processing to obtain the bounding box prediction. The correspondences are more accurate for the cross-image transformer which achieves much better results for detection. We also observe that, in this benchmark, the performances with unsupervised segments are close to the ones using COCO [36] instance annotations, which suggests that our approach does not depend on human annotations. Note that the best performance of ArtMiner is obtained with a discovery score which is expensive to compute and involves multi-scale feature matching and RANSAC. Our approach is thus simpler, faster and more effective.

Place recognition In Table 2 we compare our approach to state of the art for place recognition on the Pitts30k [62] and Tokyo 24/7 [61] datasets. The descriptions of the datasets are in the supplementary material [2]. We follow the standard evaluation protocol [4, 16, 19, 51, 61, 62]. The query image is correctly localized if one of the top N retrieved database images is within $d = 25$ meters from the ground truth TUM coordinate of the query. The recall is then reported for $N = 1, 5, 10$. For Tokyo 24/7 we follow [16, 61] and perform spatial non-maximal suppression on ranked database images before evaluation. To enable fast evaluation, we follow PatchVlad [20] and evaluate our score on the top-100 images given by NetVLAD [3]. Although our approach is not specifically designed for place recognition, it achieves performances comparable to Patch-NetVLAD [20] without RANSAC. Note that the competing approaches either employ specific supervisions or more complicated process such as RANSAC, while our approach is trained only with our synthetic segment swapping data. Note that on this task where retrieving discriminative repeated regions is sufficient and correspondence accuracy is not critical, the Nc-Net architecture preforms better. Similar to the Brueghel

Method	Supervision	Tokyo 24/7 [61]			Pitts30k-test [62]		
		R@1	R@5	R@10	R@1	R@5	R@10
AP-GEM [20,44]	Image location	40.3	55.6	65.4	75.3	89.3	92.5
DenseVLAD [20,61]	Image location	59.4	67.3	72.1	77.7	88.3	91.6
NetVLAD [3,20]	Image location	73.3	82.9	86.0	86.0	93.2	95.1
CRN [16,29]	Image location	75.2	83.8	87.3	-	-	-
SARE [16,38]	Image location	79.7	86.7	90.5	-	-	-
IBL [16]	Image location	85.4	91.1	93.3	-	-	-
Re-ranking Top-100 from NetVLAD [3,20]							
Patch-NetVLAD [20]	Image location	81.9	85.7	87.9	88.6	94.5	95.8
Patch-NetVLAD [20] + RANSAC	Image location	86.0	88.6	90.5	88.7	94.5	95.9
SuperGlue [20,50]*	Pose+Depth	88.2	90.2	90.2	88.7	95.1	96.4
Ours + Unsupervised segments							
Transformer	Segment swapping	74.0	82.9	86.0	85.2	93.5	95.4
Nc-Net	Segment swapping	84.1	87.0	88.9	86.4	94.3	95.6
Ours + COCO segments [36]							
Transformer	Segment swapping	80.0	86.0	87.9	84.7	93.5	95.6
Nc-Net	Segment swapping	85.4	88.3	89.2	86.8	94.4	95.8

* uses learnt keypoint detector Superpoint [14]

Table 2. Image-based localization on Tokyo 24/7 [61] and Pitts30k [62]. We follow Patch-NetVLAD [20] and re-rank the top-100 images ranked by NetVLAD [3] features.

results, leveraging COCO [36] annotated segments leads to superior performance. Training with unsupervised segments still leads to competitive results using the NC-Net. However, it gives clearly worst results using the transformer architecture on Tokyo 24/7.

Validation of our unsupervised segments We evaluate our unsupervised segments for unsupervised saliency detection on three standard datasets ECSSD [55], DUTS [71] and DUT-OMRON [74]. The description of the datasets are provided in the supplementary material [2]. The results are shown in Table 4. Adding Bilateral Solver [5] largely improve the performance over LOST [56] and achieves a significant boost on all the datasets compared to state-of-the-art approaches. We hope that this simple method can serve as a strong baseline for unsupervised saliency detection. Qualitative results are in the supplementary material [2].

Ablation study An ablation study of our approach using the cross-image transformer architecture and COCO [36] annotated segments is shown in Table 3 on the Brueghel [1, 54] and Tokyo24/7 [61] datasets. We notice that: (i) Poisson blending [42] and style transfer [23] are both critical; (ii) the three terms of the loss are necessary for good performance. Removing any of these elements results in very strong performance loss. More analysis on the importance of learning correspondences, the similarity score and the architectures are provided in the supplementary material [2].

5. Application to Object Discovery and Co-segmentation

In this section, we introduce an approach to use our predicted masks and correspondences for object discovery. We first explain how we designed a correspondences graph on which it is possible to perform spectral analysis in Section 5.1. We then present experimental results on the dataset

Dataset		Losses			Cross-image Transformer	
Poisson blending [42]	Style transfer [23]	\mathcal{L}_m	\mathcal{L}_{tm}	\mathcal{L}_{corr}	Brueghel [1, 54] mAP	Tokyo 24/7 [61] R@1
✓	✓	✓	✓	✓	84.4	80.0
✓	✓	✓	✓	✓	75.1	60.0
✓	✗	✓	✓	✓	75.6	57.8
✓	✓	✗	✓	✓	80.9	67.8
✓	✓	✓	✗	✓	79.8	61.3
✓	✓	✓	✓	✗	8.5	13.3

Table 3. Ablation study. We report retrieval mAP on Brueghel [1, 54] and R@1 on Tokyo 24/7 [61] with our cross-image transformer using COCO segments [36].

Method	ECSSD [55]			DUTS [71]			DUT-OMRON [74]		
	max F_β	IoU	Acc.	max F_β	IoU	Acc.	max F_β	IoU	Acc.
HS [73]	0.673	0.508	0.847	0.504	0.369	0.826	0.561	0.433	0.843
wCtr [77]	0.684	0.517	0.862	0.522	0.392	0.835	0.541	0.416	0.838
WSC [34]	0.683	0.498	0.852	0.528	0.384	0.862	0.523	0.387	0.865
DeepUSPS [41]	0.584	0.440	0.795	0.425	0.305	0.773	0.414	0.305	0.779
BigBiGAN [70]	0.782	0.672	0.899	0.608	0.498	0.878	0.549	0.453	0.856
E-BigBiGAN [70]	0.797	0.684	0.906	0.624	0.511	0.882	0.563	0.464	0.860
LOST [56]	0.758	0.654	0.895	0.611	0.518	0.871	0.473	0.410	0.797
LOST [56] + Bilateral Solver [5] (Ours)	0.837	0.723	0.916	0.697	0.572	0.887	0.578	0.489	0.818

Table 4. Unsupervised saliency detection on ECSSD [55], DUTS [71] and DUT-OMRON [74].

of [49] for co-segmentation and Brueghel [1, 54] for discovery in Section 5.2.

5.1. Correspondences graph and clustering

In the spirit of [32], we see object discovery as a graph clustering problem, where the vertices \mathcal{V} of the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ are correspondences between images and the weights of the edges encodes consistency between the correspondences. Let us consider a set of N images (I_1, \dots, I_N). For every pair of images our network predicts correspondences that we add to the set of vertices \mathcal{V} if the associated mask value is higher than a threshold. Each vertex $v_i = (s_i, t_i, x_i^s, x_i^t, m_i)$ in the graph is thus associated to a predicted correspondence and defined by the indices s_i and t_i of the images it connects, the associated coordinates x_i^s and x_i^t and the predicted mask value m_i . We use the masks values and cycle consistency between the correspondences to define the weights of the edges between the different vertices. More precisely, we only connect correspondences which have exactly one image in common. For example, let's assume that we have two vertices v_i and v_j such that $s_i = s_j = s$ and $t_i \neq t_j$. We use our network to predict correspondence fields $\mathbf{C}^{t_i \rightarrow t_j}$ and $\mathbf{C}^{t_j \rightarrow t_i}$ and we define the weight $\epsilon_{i,j}$ of the edge between v_i and v_j as:

$$\epsilon_{i,j} = \frac{1}{2} m_i m_j \exp\left(-\frac{\|x_i^s - x_j^s\|}{\sigma}\right) \left(\exp\left(-\frac{\|x_i^t - \mathbf{C}^{t_j \rightarrow t_i}(x_j^t)\|}{\sigma}\right) + \exp\left(-\frac{\|x_j^t - \mathbf{C}^{t_i \rightarrow t_j}(x_i^t)\|}{\sigma}\right) \right) \quad (4)$$

where σ is a scalar hyper-parameter. The edges are defined similarly in the cases $s_i = t_j$, $t_i = s_j$ and $t_i = t_j$. More details about the way we define the graph and in particular strategies to limit the number of vertices are given in the supplementary material [2].

Method	Airplane		Car		Horse		Avg	
	\mathcal{P}	\mathcal{J}	\mathcal{P}	\mathcal{J}	\mathcal{P}	\mathcal{J}	\mathcal{P}	\mathcal{J}
DOCS [35]*	0.946	0.64	0.940	0.83	0.914	0.65	0.933	0.70
Sun et al. [57]	0.886	0.36	0.870	0.73	0.876	0.55	0.877	0.55
Rubinstein et al. [49]	0.880	0.56	0.854	0.64	0.828	0.52	0.827	0.43
Chen et al. [10]	0.902	0.40	0.876	0.65	0.893	0.58	0.890	0.54
Quan et al. [43]	0.910	0.56	0.885	0.67	0.893	0.58	0.896	0.60
Chang et al. [8]	0.726	0.27	0.759	0.36	0.797	0.36	0.761	0.33
Lee et al. [31]	0.528	0.36	0.647	0.42	0.701	0.39	0.625	0.39
Jerripothula et al. [25]	0.905	0.61	0.880	0.71	0.883	0.61	0.889	0.64
Hsu et al. [22]	0.936	0.66	0.914	0.79	0.876	0.59	0.909	0.68
Chen et al. [11]	0.941	0.65	0.940	0.82	0.922	0.63	0.935	0.70
Ours + Unsupervised segments								
transformer	0.925	0.65	0.914	0.79	0.909	0.60	0.916	0.68
Nc-Net	0.746	0.25	0.874	0.68	0.836	0.38	0.819	0.44
Ours + COCO segments [36]								
transformer	0.941	0.67	0.928	0.82	0.916	0.60	0.928	0.70
Nc-Net	0.655	0.23	0.857	0.61	0.873	0.43	0.795	0.42


* learned with strong supervision (i.e., manually annotated object masks)

Table 5. Co-segmentation on the dataset of [49]. We report pixel level precision \mathcal{P} and Jaccard index \mathcal{J}


Given the correspondence graph, we use the spectral decomposition of its adjacency matrix [32, 40] either to obtain clusters of correspondences for object discovery, or a foreground potential for co-segmentation. For discovery we first compute N_{eig} principal eigenvectors then performing K-means with $K_{cluster}$ clusters. For co-segmentation, we directly use the first eigenvector to define a foreground potential. Note that because we only consider in the graph correspondences with mask values higher than a threshold, the full graph is extremely sparse that the eigen-decomposition can be efficiently computed.

5.2. Experiments

Object co-segmentation on the dataset of [49] We build the correspondences graph using for each image only the correspondences in the five most similar images according to the retrieval score of Equation 3. We then use the principal eigen-vector of the correspondence graph to define a seed for GrabCut [47]. More precisely, for every image we associate to each position the sum of the eigen-vector values for the correspondences at this position. Note that GrabCut [47] is crucial to achieve good performance on this dataset, and is widely used by competing approaches such as [11, 22, 25, 43, 49]. More details about the GrabCut [47] can be found in the supplementary material [2]. We follow the standard evaluation protocol [11, 49] and report pixel-level precision \mathcal{P} and the Jaccard index \mathcal{J} on three subsets: Airplane, Car, Horse. The precision \mathcal{P} measures pixel accuracy. The Jaccard index \mathcal{J} is the IoU between the segmented object and ground truth object. Quantitative results are presented in Table 5 and qualitative results in Figure 5a. Our cross-image transformer obtains performance comparable to the state of the art unsupervised approaches. Again, the performances using annotated COCO [36] segments and unsupervised segments are close, which demonstrates that



(a) Co-segmentation results in the dataset of [49] for the Horse, Airplane and Car categories.



(b) Examples of discovered clustered on Brueghel [1, 54].

Figure 5. Visual results for object discovery. We show co-segmentation results on the dataset of [49] in Figure 5a and some discovered clusters in Brueghel [1, 54] in Figure 5b.

the success of our approach does not come from implicitly leveraging annotated object segmentations. Sparse Nc-Net performances are clearly worse for this task. This can be understood by looking at qualitative results: the segmentation masks predicted by Nc-Net tend to be more localized in discriminative regions.

Discovery on Brueghel dataset [1, 54] To compute correspondences and build our correspondence graph, we first resize all images to 640×640 , as many repeated details in Brueghel [1, 54] are small. We also remove duplicate images and images with similar borders, easily detected with our algorithm, to focus on more interesting repeated details. Again, we only include in the graph the correspondences from the five most similar images according to the retrieval score to limit the size of the graph and we perform K-means for $K_{cluster} = 500$ clusters with $N_{eig} = 100$ principal eigen vectors. Figure 5b presents some interesting clusters that are not covered by ArtMiner [54]¹. More results and details are in the supplementary material [2].

¹ <http://imagine.enpc.fr/~shenx/ArtMiner/visualRes/brueghel/brueghel.html>

Limitation The current version of our approach is not scalable to large datasets, as our approach needs to compare pairs of images for retrieval and the dimensions of the graph is quadratic on number of images.

6. Conclusion

In this work, we presented a “segment swapping” approach to generate pairs of images with repeated patterns from which we show it is possible to train networks to predict co-segmentation. We evaluated two architectures, a cross-image transformer we introduced and a modified Sparse Nc-Net [46]. We also compared using annotated segments in COCO [36] and segments extracted in a completely unsupervised way, which shows that our approach is not reliant on COCO [36] object annotations. We demonstrated the interest and generality of the trained co-segmentation networks by showing competitive or better performance compared to specialized baseline on a wide range of datasets and tasks, including art detail retrieval, place recognition and object discovery.

Acknowledgement This work was supported in part by ANR project EnHerit, project Rapid Tabasco, and IDRIS under the allocation AD011011160R1 made by GENCI.

References

- [1] Brueghel family: Jan brueghel the elder." the brueghel family database. university of california, berkeley. <http://www.janbrueghel.net/>. Accessed: 2018-10-16. 5, 6, 7, 8
- [2] Supplementary material of segswap. <http://imagine.enpc.fr/~shenx/SegSwap/suppMat.pdf>. 3, 4, 5, 6, 7, 8
- [3] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *CVPR*, 2016. 6
- [4] Relja Arandjelović and Andrew Zisserman. Dislocation: Scalable descriptor distinctiveness for location recognition. In *ACCV*, 2014. 6
- [5] Jonathan T Barron and Ben Poole. The fast bilateral solver. In *ECCV*, 2016. 2, 3, 6, 7
- [6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 4
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 2
- [8] Haw-Shiuan Chang and Yu-Chiang Frank Wang. Optimizing the decomposition for multiple foreground cosegmentation. *CVIU*, 2015. 7
- [9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv*, 2020. 4, 6
- [10] Xinlei Chen, Abhinav Shrivastava, and Abhinav Gupta. Enriching visual knowledge bases via object discovery and segmentation. In *CVPR*, 2014. 7
- [11] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. Show, match and segment: Joint weakly supervised learning of semantic matching and object co-segmentation. *PAMI*, 2020. 2, 3, 7
- [12] Minsu Cho, Suha Kwak, Cordelia Schmid, and Jean Ponce. Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals. In *CVPR*, 2015. 3
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 4
- [14] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *CVPRW*, 2018. 6
- [15] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015. 2
- [16] Yixiao Ge, Haibo Wang, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-supervising fine-grained region similarities for large-scale image localization. In *ECCV*, 2020. 6
- [17] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2019. 3
- [18] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, 2021. 3
- [19] Petr Gronat, Guillaume Obozinski, Josef Sivic, and Tomas Pajdla. Learning and calibrating per-location classifiers for visual place recognition. In *CVPR*, 2013. 6
- [20] Stephen Hausler, Sourav Garg, Ming Xu, Michael Milford, and Tobias Fischer. Patch-netvlad: Multi-scale fusion of locally-global descriptors for place recognition. In *CVPR*, 2021. 2, 6
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [22] Kuang-Jui Hsu, Yen-Yu Lin, Yung-Yu Chuang, et al. Co-attention cnns for unsupervised object co-segmentation. In *IJCAI*, 2018. 3, 7
- [23] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. 2, 3, 6, 7
- [24] Joel Janai, Fatma Guney, Anurag Ranjan, Michael Black, and Andreas Geiger. Unsupervised learning of multi-frame optical flow with occlusions. In *ECCV*, 2018. 3
- [25] Koteswar Rao Jerripothula, Jianfei Cai, and Junsong Yuan. Image co-segmentation via saliency co-fusion. *TOMM*, 2016. 7
- [26] Wei Jiang, Eduard Trulls, Jan Hosang, Andrea Tagliasacchi, and Kwang Moo Yi. Cotr: Correspondence transformer for matching across images. *arXiv*, 2021. 2
- [27] Armand Joulin, Francis Bach, and Jean Ponce. Discriminative clustering for image co-segmentation. In *CVPR*, 2010. 3
- [28] Armand Joulin, Francis Bach, and Jean Ponce. Multi-class cosegmentation. In *CVPR*, 2012. 3
- [29] Hyo Jin Kim, Enrique Dunn, and Jan-Michael Frahm. Learned contextual feature reweighting for image geo-localization. In *CVPR*, 2017. 6
- [30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv*, 2014. 4
- [31] Chulwoo Lee, Won-Dong Jang, Jae-Young Sim, and Chang-Su Kim. Multiple random walkers and their application to image cosegmentation. In *CVPR*, 2015. 7
- [32] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *ICCV*, 2005. 2, 7
- [33] Bo Li, Zhengxing Sun, Qian Li, Yunjie Wu, and Anqi Hu. Group-wise deep object co-segmentation with co-attention recurrent neural network. In *ICCV*, 2019. 3
- [34] Nianyi Li, Bilin Sun, and Jingyi Yu. A weighted sparse coding framework for saliency detection. In *CVPR*, 2015. 7
- [35] Weihao Li, Omid Hosseini Jafari, and Carsten Rother. Deep object co-segmentation. In *ACCV*, 2018. 3, 7
- [36] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 3, 6, 7, 8

- [37] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *PAMI*, 2010. 2
- [38] Liu Liu, Hongdong Li, and Yuchao Dai. Stochastic attraction-repulsion embedding for large scale image localization. In *ICCV*, 2019. 6
- [39] Iaroslav Melekhov, Aleksei Tiulpin, Torsten Sattler, Marc Pollefeys, Esa Rahtu, and Juho Kannala. Dgc-net: Dense geometric correspondence network. In *WACV*, 2019. 2
- [40] Andrew Ng, Michael Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. *NeurIPS*, 2001. 2, 7
- [41] Duc Tam Nguyen, Maximilian Dax, Chaithanya Kumar Mummadi, Thi Phuong Nhung Ngo, Thi Hoai Phuong Nguyen, Zhongyu Lou, and Thomas Brox. Deepusps: Deep robust unsupervised saliency prediction with self-supervision. In *NeurIPS*, 2019. 7
- [42] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. In *SIGGRAPH*, 2003. 2, 3, 6, 7
- [43] Rong Quan, Junwei Han, Dingwen Zhang, and Feiping Nie. Object co-segmentation via graph optimized-flexible manifold ranking. In *CVPR*, 2016. 7
- [44] Jerome Revaud, Jon Almazán, Rafael S Rezende, and Cesar Roberto de Souza. Learning with average precision: Training image retrieval with a listwise loss. In *ICCV*, 2019. 6
- [45] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *CVPR*, 2017. 2
- [46] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. Efficient neighbourhood consensus networks via submanifold sparse convolutions. In *ECCV*, 2020. 2, 3, 4, 8
- [47] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. "grabcut" interactive foreground extraction using iterated graph cuts. *ACM transactions on graphics (TOG)*, 2004. 7
- [48] Carsten Rother, Tom Minka, Andrew Blake, and Vladimir Kolmogorov. Cosegmentation of image pairs by histogram matching-incorporating a global constraint into mrfs. In *CVPR*, 2006. 3
- [49] Michael Rubinstein, Armand Joulin, Johannes Kopf, and Ce Liu. Unsupervised joint object discovery and segmentation in internet images. In *CVPR*, 2013. 1, 2, 3, 7, 8
- [50] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *CVPR*, 2020. 2, 4, 6
- [51] Torsten Sattler, Tobias Weyand, Bastian Leibe, and Leif Kobbelt. Image retrieval for image-based localization revisited. In *BMVC*, 2012. 6
- [52] Paul Hongsuck Seo, Jongmin Lee, Deunsol Jung, Bohyung Han, and Minsu Cho. Attentive semantic alignment with offset-aware correlation kernels. In *ECCV*, 2018. 2
- [53] Xi Shen, François Darmon, Alexei A Efros, and Mathieu Aubry. Ransac-flow: generic two-stage image alignment. In *ECCV*, 2020. 2, 3
- [54] Xi Shen, Alexei A Efros, and Mathieu Aubry. Discovering visual patterns in art collections with spatially-consistent feature learning. In *CVPR*, 2019. 1, 2, 3, 5, 6, 7, 8
- [55] Jianping Shi, Qiong Yan, Li Xu, and Jiaya Jia. Hierarchical image saliency detection on extended cssd. *TPAMI*, 2015. 6, 7
- [56] Oriane Siméoni, Gilles Puy, Huy V. Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. 2021. 2, 3, 6, 7
- [57] Jian Sun and Jean Ponce. Learning dictionary of discriminative part detectors for image categorization and cosegmentation. *IJCV*, 2016. 7
- [58] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loft: Detector-free local feature matching with transformers. *arXiv*, 2021. 2
- [59] Kevin Tang, Armand Joulin, Li-Jia Li, and Li Fei-Fei. Co-localization in real-world images. In *CVPR*, 2014. 3
- [60] Tatsunori Tanai, Sudipta N Sinha, and Yoichi Sato. Joint recovery of dense correspondence and cosegmentation in two images. In *CVPR*, 2016. 3
- [61] Akihiko Torii, Relja Arandjelovic, Josef Sivic, Masatoshi Okutomi, and Tomas Pajdla. 24/7 place recognition by view synthesis. In *CVPR*, 2015. 2, 5, 6, 7
- [62] Akihiko Torii, Josef Sivic, Tomas Pajdla, and Masatoshi Okutomi. Visual place recognition with repetitive structures. In *CVPR*, 2013. 2, 5, 6
- [63] Prune Truong, Martin Danelljan, and Radu Timofte. Glunet: Global-local universal network for dense flow and correspondences. In *CVPR*, 2020. 2
- [64] Prune Truong, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning accurate dense correspondences and when to trust them. *arXiv*, 2021. 2, 3
- [65] Prune Truong, Martin Danelljan, Fisher Yu, and Luc Van Gool. Warp consistency for unsupervised learning of dense correspondences. *arXiv*, 2021. 2, 3
- [66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2, 4
- [67] Sara Vicente, Carsten Rother, and Vladimir Kolmogorov. Object cosegmentation. In *CVPR*, 2011. 3
- [68] Huy V Vo, Francis Bach, Minsu Cho, Kai Han, Yann LeCun, Patrick Pérez, and Jean Ponce. Unsupervised image matching and object discovery as optimization. In *CVPR*, 2019. 3
- [69] Huy V Vo, Patrick Pérez, and Jean Ponce. Toward unsupervised, multi-object discovery in large-scale image collections. In *ECCV*, 2020. 3
- [70] Andrey Voynov, Stanislav Morozov, and Artem Babenko. Object segmentation without labels with large-scale generative models. In *ICML*, 2021. 7
- [71] 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 *CVPR*, 2017. 6, 7
- [72] Yang Wang, Yi Yang, Zhenheng Yang, Liang Zhao, Peng Wang, and Wei Xu. Occlusion aware unsupervised learning of optical flow. In *CVPR*, 2018. 3

- [73] Qiong Yan, Li Xu, Jianping Shi, and Jiaya Jia. Hierarchical saliency detection. In *CVPR*, 2013. 7
- [74] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via graph-based manifold ranking. In *CVPR*, 2013. 6, 7
- [75] Ze-Huan Yuan, Tong Lu, Yirui Wu, et al. Deep-dense conditional random fields for object co-segmentation. In *IJCAI*, 2017. 3
- [76] Tinghui Zhou, Yong Jae Lee, Stella X Yu, and Alexei A Efros. Flowweb: Joint image set alignment by weaving consistent, pixel-wise correspondences. In *CVPR*, 2015. 3
- [77] Wangjiang Zhu, Shuang Liang, Yichen Wei, and Jian Sun. Saliency optimization from robust background detection. In *CVPR*, 2014. 7