

Unsupervised Salient Object Detection with Spectral Cluster Voting

Gyungin Shin¹ Samuel Albanie² Weidi Xie^{1,3}
¹ Visual Geometry Group, University of Oxford, UK
² Department of Engineering, University of Cambridge, UK
³ Shanghai Jiao Tong University, China
gyungin@robots.ox.ac.uk

Abstract

In this paper, we tackle the challenging task of unsupervised salient object detection (SOD) by leveraging spectral clustering on self-supervised features. We make the following contributions: (i) We revisit spectral clustering and demonstrate its potential to group the pixels of salient objects across various self-supervised features, e.g., MoCov2, SwAV, and DINO; (ii) Given mask proposals from multiple applications of spectral clustering on image features computed from different self-supervised models, we propose a simple but effective winner-takes-all voting mechanism for selecting the salient masks, leveraging object priors based on framing and distinctiveness; (iii) Using the selected object segmentation as pseudo groundtruth masks, we train a salient object detector, termed SELF-MASK, which outperforms prior approaches on three unsupervised SOD benchmarks. Code is publicly available at <https://github.com/NoelShin/selfmask>.

1. Introduction

Salient object detection (SOD)¹, which aims to group pixels that attract human visual attention, has been extensively studied in the field of computer vision due to its wide range of applications such as photo cropping [50, 37], re-targeting in images [18, 51] and video [45].

In the literature, early work tackled this problem by utilising low-level features (e.g., colour [14]) together with priors on salient regions in an object such as contrast priors [26], boundary priors [57] and centre priors [27]. Recent SOD models have approached this task from the perspective of representation learning, typically training deep neural networks (DNNs) on a large-scale dataset with manual annotations. However, the scalability of such supervised

¹In contrast to *object detection* (which aims to localise and recognise objects with bounding boxes), *salient object detection* aims to segment foreground objects by predicting pixel-wise masks for them.



Figure 1: In this work, we propose SELF-MASK, a framework for salient object detection that employs no human annotation. The figure depicts example segmentations produced by our model on DUT-OMRON [63], DUTS-TE [55] and ECSSD [47], with blue and orange regions denoting the intersection and the difference between the predicted mask and ground-truth label, respectively. Despite no supervision, SELF-MASK reliably segments a close approximation to the full spatial extent of salient regions. Best viewed in colour.

learning approach is limited because it is a costly process to collect ground-truth mask annotations.

To overcome the necessity of large-scale human annotation, many unsupervised methods for saliency detection/object segmentation have recently been proposed [6, 12, 53, 54, 38, 65, 68, 42]. Despite these efforts, the gap between unsupervised and fully supervised SOD methods remains significant.

Interestingly, however, it has been noted that recent self-supervised models such as DINO [10] exhibit significant potential for object segmentation despite the fact that their

training objective does not explicitly encourage pixel grouping. The focus of this work is to leverage this observation to propose a *simple yet effective mechanism for extracting object regions from self-supervised features* that can be employed for the task of unsupervised salient object detection.

To this end, we explore the use of spectral clustering [46], a classical graph-theoretic clustering algorithm, and find that it can generate useful segmentation candidates across a range of self-supervised features (i.e., DINO [10], MoCov2 [24], and SwAV [9]). Motivated by this finding, we propose a simple *winner-takes-all* voting strategy to select the salient object masks among a collection of clusters produced by repeated applications of spectral clustering to self-supervised features.² In particular, we base our voting strategy on two priors: The first is a *framing prior* that a salient object should not occupy the full image height or width; the second is a *distinctiveness prior* that assumes that salient regions are sufficiently distinctive that they will appear as clusters among an appropriately constructed collection of redundant re-clusterings of the data. We then show that the selected salient masks can be employed as *pseudo-labels* to train a saliency estimation network that achieves state-of-the-art results on a variety of benchmarks (see Fig. 1 for example visualisations).

In summary, we make the following contributions: (i) We revisit spectral clustering and highlight that, when applied to self-supervised features, it brings clear benefits over *k*-means clustering as a proposal mechanism on three popular salient object detection (SOD) datasets; (ii) We propose an effective voting strategy to select the most salient object mask in an image among multiple segmentations generated from different self-supervised features by leveraging saliency priors; (iii) Using the salient masks as pseudo ground truth masks (pseudo-masks), we train an object segmentation model, SELFMASK, that outperforms the previous unsupervised saliency detection approaches on three SOD benchmarks.

2. Related work

Our work relates to two themes in the literature: *self-supervised representation learning* and *unsupervised saliency detection*. Each of these is discussed next.

2.1. Self-supervised representation learning

There has been a great deal of interest in self-supervised approaches to learning visual representations that obviate the requirement for labels. These include techniques for solving proxy tasks, such as predicting patch locations [19], patch discrimination [21], grouping through co-occurrence [25], colourisation [29], jigsaw puzzles [43],

²In this work, we use the terms *cluster* and *mask* interchangeably. Specifically, by mask, we mean a (one-hot) mask which encodes the spatial extent of a cluster.

common fate principles [36], clustering [8, 4], and instance discrimination [58, 24]. Our work is inspired by recent work exploring the use of self-supervised transformers [61, 10], and by the analysis provided by [10, 41] who noted that the self-attention of Vision Transformers (ViT) [21] are capable of highlighting spatially coherent object regions in their input. While prior work [64] has demonstrated that self-supervised pretraining can be effective for semantic segmentation when coupled with end-to-end supervised metric learning on the target dataset, we instead seek a simple way to exploit self-supervised features for object segmentation without annotation.

2.2. Unsupervised saliency detection

Supervised object segmentation requires pixel-wise annotations which are time-consuming to acquire. Seeking to avoid this cost, many attempts have been made to solve the task in an unsupervised fashion. Prior to the dominance of deep neural networks, a broad range of handcrafted methods were proposed [33, 1, 28, 22, 62, 69, 15] based on one or more priors relating to foreground regions within an image such as the contrast prior [26], centre prior [27], and boundary prior [57]. However, these handcrafted approaches suffer from poor performance relative to recent DNN-based models, described next.

Generative models. A common approach for DNN-based unsupervised object segmentation is to utilise generative adversarial networks (GANs) [23]. Specifically, given an image, a generator is adversarially trained to produce an object mask which will be used to composite a realistic image by copying the corresponding object region in the image into a background which is either synthesised [6] or taken from a different image [2]. An alternative family of approaches aims to discover a direction in the latent space of a pre-trained GAN that can be used to segment foreground and background regions [53, 54, 38]. Then, a saliency detector is trained on a synthetic data set composed of pairs of images together with their foreground masks generated via the discovered latent space structure. In contrast, we seek to exploit representations learned via self-supervision by discriminative, rather than generative, models.

Noisy supervision with pseudo-labels. More closely related to our work, the use of weak pseudo saliency masks for training a DNN has been proposed. SBF [65], the first attempt to train saliency detector without human annotation, proposed to train a model with superpixel-level pseudo-masks generated by fusing weak saliency maps from multiple unsupervised methods (i.e., [67, 66, 48]). Similarly, USD [68] aims to learn from diverse noisy pseudo-labels obtained via distinct unsupervised handcrafted methods in such a way that a saliency detector trained with the pseudo-masks can predict a saliency map free from label noise.

DeepUSPS [42] proposed to refine the pseudo-masks for images produced by handcrafted saliency methods, by training segmentation networks via a self-supervised iterative refinement process. The refined pseudo-masks are then combined from different handcrafted methods to train a final segmentation network. In contrast to the methods above, we use neither handcrafted saliency methods nor an iterative refinement strategy, resulting in a simpler learning framework.

Object segmentation properties of self-supervised vision transformers. Another line of related work has sought to investigate the observation that self-supervised ViTs [10] exhibit object segmentation potential. LOST [49] propose to pick a seed patch from such a ViT that is likely to contain part of a foreground object, and then expand the seed patch to different patches sharing a high similarity with the seed patch. Concurrent work, TokenCut [56], proposes to use Normalised Cuts [46] to segment the salient object among the final layer self-attention key features of ViT. SELF-MASK differs from the TokenCut approach to salient object detection in two key ways: (1) While we similarly employ spectral clustering as part of our pipeline, we demonstrate the significant additional value of integrating cues from diverse re-clusterings via voting to bootstrap a pseudo-labelling process; (2) Rather than focusing solely on transformers, we leverage saliency cues from self-supervised convolutional neural networks (CNNs) as well as ViT architectures, and show the benefits of doing so. We compare our approach with theirs in section 4.

3. Method

In this section, we begin by formalising the problem scenario (Sec. 3.1) and briefly summarise spectral clustering (Sec. 3.2). Then, we introduce our approach to address unsupervised salient object detection by selecting pseudo ground truth masks via spectral clustering, and train a saliency prediction network, called SELF-MASK (Sec. 3.3).

3.1. Problem formulation

Here, we consider the task of unsupervised salient object detection (SOD), with the goal of training a segmenter (Φ_{seg}) that seeks to partition the image into two disjoint groups, namely foreground and background. That is,

$$\Phi_{\text{seg}}(\mathbf{I}; \Theta) = \mathbf{M}_{\text{seg}} \in \{0, 1\}^{H \times W} \quad (1)$$

where $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ refers to an input image, and Θ represents learnable parameters. \mathbf{M}_{seg} denotes a binary segmentation mask, with 1s denoting a foreground region and 0s denoting background.

Traditionally this has been treated as a clustering problem where the key challenge lies in designing effective fea-

tures for accurately describing salient regions. In this work, we instead look for a simple yet effective solution by leveraging self-supervised visual representations.

3.2. Segmentation with spectral clustering

Conceptually, segmentation is obtained via spectral clustering [46] with pixel-wise image features by projecting the features onto a representation space such that image partitions can be decided by directly comparing similarities between their corresponding features.

Concretely, we first extract dense features from the image, *i.e.*, $\mathbf{F} \in \mathbb{R}^{h \times w \times D}$ with a pre-trained convolutional or transformer-based encoder. Then, each feature vector $\mathbf{f}_i \in \mathbb{R}^D$ in the dense feature map \mathbf{F} can be seen as a vertex in an undirected graph with vertex set $\mathbf{V} = \{\mathbf{f}_1, \dots, \mathbf{f}_N\} \in \mathbb{R}^{N \times D}$, where $N = h \cdot w$. Each edge between two vertices \mathbf{f}_i and \mathbf{f}_j is associated with a non-negative weight $w_{ij} \geq 0$ defined by feature similarity. In particular, the weighted adjacency matrix \mathbf{W} of the graph is computed as

$$\mathbf{W} = (w_{ij}) = \frac{\mathbf{V}\mathbf{V}^T}{\|\mathbf{V}\| \|\mathbf{V}\|} \in \mathbb{R}^{N \times N} \quad (2)$$

where the *degree* of a vertex $\mathbf{f}_i \in \mathbf{V}$ is defined as $d_i = \sum_{j=1}^N w_{ij}$, and the *degree matrix* \mathbf{D} is defined with the degrees d_1, \dots, d_N on the diagonal. Given the adjacency matrix \mathbf{W} and the degree matrix \mathbf{D} , the (un-normalised) graph Laplacian \mathbf{L} is defined as:

$$\mathbf{L} = \mathbf{D} - \mathbf{W} \quad (3)$$

Given \mathbf{L} , we can solve the generalised eigenproblem:

$$\mathbf{L}\mathbf{u} = \lambda\mathbf{D}\mathbf{u} \quad (4)$$

where $\mathbf{u} \in \mathbb{R}^N$ and λ represent an eigenvector and its eigenvalue. We take the k eigenvectors with the lowest eigenvalues and form a matrix $\mathbf{U} \in \mathbb{R}^{N \times k}$ which has the eigenvectors as its columns.

Finally, a set of clusters \mathcal{C} is obtained by running k -means algorithm on the row vectors of a matrix \mathbf{U} (see supplementary for more detail), producing regions with all pixels from the corresponding cluster. Note that, at this stage, the resulting clusters are composed of both object and background masks—the object mask itself will be selected by our selection strategy described next.

3.3. Supervision with pseudo-mask from spectral clusters

Here, we first introduce our voting strategy for selecting a salient mask from a set of spectral clusters from different features and multiple k values (*i.e.*, cluster numbers), which utilises a *framing prior* and *distinctiveness prior*. Then, we describe our model, SELF-MASK, which is trained by using the selected salient masks as pseudo-masks for supervision.

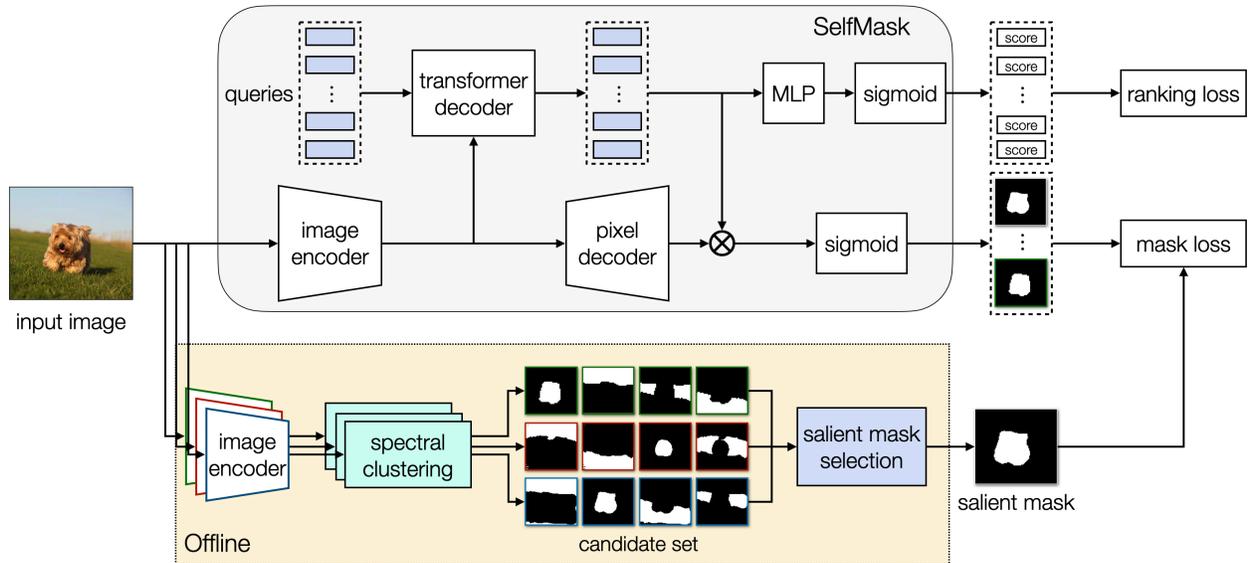


Figure 2: Overview of our approach. Given different self-supervised encoders, we first generate a set of pseudo-mask candidates per image using spectral clustering before the training step. In the figure we show 12 masks from clusterings from three different encoder features with $k=4$. We select the most salient mask among them via the proposed voting strategy and use it as a pseudo-mask for the image. Then we train our model to predict n_q queries (*i.e.*, predictions), all of which are encouraged to be similar to the salient mask. To make the model aware of the objectness of each prediction, we use the ranking loss which encourages the objectness score of a prediction closer to the salient mask to be higher. At inference time, we select the prediction with the highest objectness score. Please see the text for details.

3.3.1 Spectral cluster winner-takes-all voting

To choose the salient object among the mixture of foreground and background masks generated by spectral clustering, we propose a voting strategy based on two observations: (1) The spatial extent of an object rarely occupies the entire height and width of an image. (2) Salient object regions are likely to appear in multiple clusters across different self-supervised features as well as with different cluster numbers k . In other words, among k clusters from an application of clustering, we assume that at least one cluster encodes an object region within the image and that this holds for different features (*e.g.*, DINO, MoCov2, or SwAV). We call these priors the *framing prior* and *distinctiveness prior*, respectively. Note that the framing prior bears a resemblance to the centre prior [27], which states that salient objects are likely to be located near the center of an image, and the boundary prior [57], which presumes that the foreground object rarely touches the boundary of an image. However, the framing prior differs from these priors in that it is not related to a location of an object but rather the scale of an object within an image.

To use the framing prior and distinctiveness prior in practice, we first form a candidate set of masks by repeatedly applying spectral clustering to different features with multiple k values. Then, we treat masks whose spatial extent is as long as the width or height of the image as background

masks, and eliminate them from the candidate set. Finally, we employ *winner-takes-all* voting: we pick the mask with the highest average pair-wise similarity w.r.t. IoU among all remaining masks as the final mask for salient objects (bottom of Fig. 2). There are two edge cases in the background elimination process we handle explicitly: (i) when no masks are left in the candidate set and (ii) when only two masks are left, sharing the same IoU. For the former case, we simply keep all the masks in the candidate set as every mask highlights regions spanning the spatial extent of the image, breaking the assumption of the framing prior. For the latter, we break ties randomly to pick one of the two masks.

3.3.2 SELFMask

Here, we describe our model architecture, training, and inference procedure.

Architecture. We base our salient object detection network, called SELFMask, on a variant of the MaskFormer architecture [13] which was originally proposed for the semantic/instance segmentation task.

SELFMask has two different sets of outputs: mask predictions and objectness scores for each mask. In detail, our model comprises an image encoder, a pixel decoder, and a transformer decoder (upper part of Fig. 2). The image encoder takes an image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ as input and outputs

feature maps $f(\mathbf{I}) \in \mathbb{R}^{h \times w \times D}$. Then, the feature maps are fed into the pixel decoder to produce upsampled features $g(f(\mathbf{I})) \in \mathbb{R}^{H \times W \times D}$. The feature maps are also passed to the transformer decoder which outputs n_q per-mask embeddings by using the feature maps as keys and values and the learnable embeddings as queries. The final mask predictions \mathbf{M} are produced via matrix multiplication between the upsampled features and per-mask embeddings, followed by an element-wise sigmoid function $\sigma(\cdot)$:

$$\mathbf{M} = \{\mathbf{M}_i \mid \mathbf{M}_i = \sigma(g(f(\mathbf{I})) \mathbf{q}_i), i = 1, \dots, n_q\} \quad (5)$$

where $\mathbf{q}_i \in \mathbb{R}^D$ denotes the i th query (*i.e.*, per-mask embedding). For each mask \mathbf{M}_i , an objectness score $o_i \in [0, 1]$ is estimated by feeding the corresponding per-mask embedding \mathbf{q}_i to a simple MLP with two hidden layers followed by a sigmoid function. Note that other encoder and decoder architectures can also be used here.

Objective. For training, we employ two objective functions: a mask loss and a ranking loss, denoted by $\mathcal{L}_{\text{mask}}$ and $\mathcal{L}_{\text{rank}}$, respectively. Given n_q mask predictions for an image from the model, we encourage all predictions to be similar to the pseudo-mask. Specifically, following [13], we use the Dice coefficient [39], which considers the class-imbalance between foreground and background regions within an image, as the mask loss. It is worth noting that, unlike [13], we do not include the focal loss [32] in the mask loss since we find that it hinders convergence.

To decide which prediction best highlights the salient region in the image among the proposed candidates when $n_q > 1$, we rank the predicted masks based on their objectness score. Specifically, we first re-order the indices of the predicted masks by their mask loss in ascending order such that

$$\mathcal{L}_{\text{mask}}(\mathbf{M}_i, \mathbf{M}_{\text{pseudo}}) \leq \mathcal{L}_{\text{mask}}(\mathbf{M}_j, \mathbf{M}_{\text{pseudo}}) \quad (6)$$

for any $i < j$, where $\mathbf{M}_{\text{pseudo}}$ denotes the target pseudo-mask for the image. Then, we enforce the objectness score o_i of the mask \mathbf{M}_i to be higher than the scores o_j for any $j > i$. As a consequence the model is encouraged to produce a higher score for a predicted mask that more closely resembles the pseudo-mask than other predictions. We instantiate this ranking loss as a *hinge loss* [16]:

$$\mathcal{L}_{\text{rank}} = \sum_{i=1}^{n_q-1} \sum_{j>i}^{n_q} \max(0, o_j - o_i). \quad (7)$$

Overall, our final objective function is as follows:

$$\mathcal{L} = \mathcal{L}_{\text{mask}} + \lambda \mathcal{L}_{\text{rank}} \quad (8)$$

where λ is a weighting factor, which is set to 1.0 across our experiments. Following [7, 13], we compute the loss for outputs from each layer of the transformer decoder.

Inference. During inference, given n_q predicted masks for an image and their objectness score, we pick the mask with the highest score as the salient object detection and binarise it with a fixed threshold of 0.5.

4. Experiments

In this section, we first describe the datasets used in our experiments (Sec. 4.1) and provide implementation details (Sec. 4.2). We then conduct ablation study (Sec. 4.3) and report our results for salient object detection (Sec. 4.4).

4.1. Datasets

We use *DUTS-TR* [55], which contains 10,553 images, to train our model with the pseudo-masks generated by following Sec. 3.3. We emphasize that only images are used for generating pseudo-masks and training, without the corresponding labels. For our ablation study and comparison to previous work, we consider five popular saliency datasets including *DUT-OMRON* [63], which comprises 5,168 images of varied content with ground-truth pixel level masks; *DUTS-TE* [55], containing 5,019 images selected from the SUN dataset [59] and ImageNet test set [17]; *ECSSD* [47] which contains 1,000 images that were selected to represent complex scenes; *HKU-IS* [30] which consists of 4,447 scene images with foreground/background sharing the similar appearances; *SOD* [40] which contains 300 images with many images having multiple salient objects.

4.2. Implementation details

Networks. We use the ViT-S/8 architecture [20] for the encoder, a bilinear upsampler with a scale factor of 2 for the pixel decoder, and 6 transformer layers [52] for the transformer decoder. For the MLP applied to per-mask queries (with a dimensionality of 384) that outputs a scalar value for the objectness score, we use three fully-connected layers with a ReLU activation between them. We set the same number of units for the hidden nodes as the input (*i.e.*, 384) and output a single value followed by a sigmoid.

Training details. We train our models for 12 epochs and optimise all parameters including the backbone encoder using AdamW [35] with a learning rate of 6e-6 and the Poly learning rate policy [44, 60, 11]. For data augmentation, we use random scaling with a scale range of [0.1, 1.0], random cropping to a size of 224×224 and random horizontal flipping with a probability of 0.5. In addition, photometric transformations include random color jittering, color dropping, and Gaussian blurring are applied. We run each model with three different seeds and report the average.

Metrics. In our experiments, we report intersection-over-union (IoU), pixel accuracy (Acc) and maximal F_β score (max F_β) with β^2 set to 0.3 following [56]. Please refer to the supplementary for more details on these metrics.



Figure 3: Sample visualisations of the pseudo-masks and predictions from our model on the ECSSD, DUT-OMRON, and DUTS-TE benchmarks. From left to right, the input image, ground-truth mask, a pseudo-mask decided by the proposed voting-based salient mask selection, and the prediction of our model are shown. Blue and orange coloured regions denote the intersection and difference between a ground-truth and a predicted mask. Best viewed in colour.

Model	Arch.	Cluster.	DUT-OMRON $k=\{2, 3, 4\}$	DUTS-TE $k=\{2, 3, 4\}$	ECSSD $k=\{2, 3, 4\}$
Convolutional Nets					
MoCov2 [24]	ResNet50	k -means	.375	.415	.500
MoCov2 [24]	ResNet50	spectral	.387	.454	.627
SwAV [9]	ResNet50	k -means	.399	.444	.542
SwAV [9]	ResNet50	spectral	.401	.458	.590
Vision Transformer					
DINO [10]	ViT-S/16	k -means	.377	.392	.541
DINO [10]	ViT-S/16	spectral	.394	.417	.577
DINO [10]	ViT-S/8	k -means	.369	.377	.551
DINO [10]	ViT-S/8	spectral	.398	.411	.587

Table 1: **Spectral clustering dominates k -means for self-supervised features.** We report upper bound IoUs to compare the quality of masks produced by k -means and spectral clustering on *self*-supervised features with two different encoder architectures (i.e., convolution- and transformer-based encoder). We report the average of the results from $k=\{2, 3, 4\}$.

4.3. Ablation study

In this section, we first conduct experiments to compare the effectiveness of spectral clustering and k -means when applied to self-supervised image encoders. Next, we quantitatively verify the performance of our winner-takes-all voting strategy for foreground mask selection and compare it to different saliency selection methods. Lastly, we investigate the effect of different number of queries on SELFMask.

4.3.1 Spectral clustering vs k -means clustering

We compare spectral clustering against a k -means [34] clustering baseline on three salient object detection benchmarks. As the resulting segmentations from each algorithm are agnostic to foreground/background regions, we consider a best-case evaluation for both algorithms. In detail, given a

Model	Arch.	Cluster.	DUT-OMRON $k=\{2, 3, 4\}$	DUTS-TE $k=\{2, 3, 4\}$	ECSSD $k=\{2, 3, 4\}$
Convolutional Nets					
ResNet [24]	ResNet50	k -means	.337	.354	.444
ResNet [24]	ResNet50	spectral	.310	.327	.437
Vision Transformer					
ViT [20]	ViT-S/16	k -means	.394	.411	.542
ViT [20]	ViT-S/16	spectral	.380	.400	.551

Table 2: **Spectral clustering and k -means perform comparably for fully-supervised features.** We report upper bound IoUs to provide a comparison of mask quality comparison between k -means and spectral clustering applied to *fully*-supervised features. We report the average of the results from $k=\{2, 3, 4\}$.

groundtruth mask, we pick the cluster with the highest IoU w.r.t. the groundtruth. Such IoUs act as an upper-bound score among the clusters (i.e., average best overlap in [3]).

To account for the effect of the cluster number k on performance of the clustering algorithms, we consider different k values from 2 to 4 and average the results. For the full results with each k value, please refer to the supplementary.

As shown in Tab. 1, we observe that object masks from spectral clustering consistently outperform k -means masks by a large margin. Interestingly, however, when using fully-supervised image encoders, the performance gain of spectral clustering diminishes (Tab. 2). These findings boil down to a simple summary: *while using self-supervised visual representations for grouping, spectral clustering is considerably superior to k -means, regardless of the choice of encoder architecture and self-supervised learning algorithm.*

4.3.2 Voting for salient object masks

Here, we conduct experiments to assess our voting method for selecting a foreground mask among the mask candidates. Since these experiments include ablations across hy-

perparameter choices, we conduct them on the HKU-IS [30] and SOD [40], rather than the benchmarks used to compare to prior work.

In detail, we construct an initial corpus of mask candidates by clustering different self-supervised features with different number of clusters, as described in Sec. 3.2. For this, we build the candidate set with spectral clusters from different combinations of self-supervised features, e.g., MoCov2/DINO or SwAV/MoCov2. It is important to do so to allow our voting-based method to leverage the distinctiveness prior across different features. In addition, for each combination, we experiment with 3 different k value settings: $k=2$, $\{2, 3\}$ or $\{2, 3, 4\}$ to account for various object scales, e.g., a lower k tends to cover large regions, while a higher k segments smaller objects. We then evaluate the selected masks on the *HKU-IS* set. For reference we also compute an upper bound IoU, which is computed in a similar way as done in the previous section.

Effectiveness of clustering various self-supervised models with different number of clusters. As shown in Tab. 3, we make two observations: (i) IoU of both selected masks and upper bound masks, denoted by pseudo-mask and UB improves by increasing k across all feature combinations; (ii) using all three features (i.e., DINO, MoCov2 and SwAV) results in better pseudo-masks than using two of the three features (e.g., MoCov2 and SwAV). These support the distinctiveness prior, which assumes that at least one cluster represents a foreground region, and its application to our voting-based saliency selection method.

Effectiveness of the proposed voting approaches. We further validate the effectiveness of our voting scheme by comparing to different selection methods, i.e., random selection and a centre prior [27]-based strategy. Specifically, we form a candidate set using DINO/MoCov2/SwAV features with $k = \{2, 3, 4\}$, which amounts to 27 masks in total. For the random strategy, we simply pick one of the masks uniformly from the candidates. For the centre prior selection strategy, we choose a mask whose average Euclidean distance to the image centre from its constituent pixel locations is lowest. In addition, we consider each method with or without utilising the framing prior to assess the influence of filtering out background mask. We evaluate the IoU of each case on the *HKU-IS* and *SOD* benchmarks.

As can be seen in Tab. 4, deploying the framing prior boosts IoU in all considered selection methods, with the proposed voting selection method performing best. The framing prior plays a crucial role in the voting process: voting without this prior performs much worse than its counterparts in both random and center-based selections on *HKU-IS*, and performs similarly to the random strategy on *SOD*. This is caused in large part by mistakes when selecting background masks as salient objects.

Features			k	Pseudo-mask UB	
DINO [10]	MoCov2 [24]	SwAV [9]			
\times	\checkmark	\checkmark	2	.508	.562
			2, 3	.561	.626
			2, 3, 4	.580	.658
\checkmark	\times	\checkmark	2	.473	.553
			2, 3	.538	.644
			2, 3, 4	.559	.682
\checkmark	\checkmark	\times	2	.459	.546
			2, 3	.536	.648
			2, 3, 4	.566	.688
\checkmark	\checkmark	\checkmark	2	.511	.584
			2, 3	.567	.664
			2, 3, 4	.590	.698

Table 3: **Forming a candidate set with various self-supervised features and multiple k values improves IoU of both pseudo-masks and upper bound masks (UB).** We compare cases with different combinations of self-supervised features and cluster numbers of $k=2$, $\{2, 3\}$ or $\{2, 3, 4\}$ on the *HKU-IS* benchmark.

Selection	Framing prior	HKU-IS [30]	SOD [40]
random	\times	.206	.197
	\checkmark	.464	.277
center	\times	.362	.122
	\checkmark	.442	.392
voting (ours)	\times	.081	.200
	\checkmark	.590	.447

Table 4: **Winner-takes-all voting and the framing prior both significantly improve mask quality.** We compare our voting strategy to different selection strategies along with the effect of framing prior under the IoU metric. Selection is performed from a candidate set including DINO, MoCov2 and SwAV features with $k = \{2, 3, 4\}$.

4.3.3 The influence of the number of queries

As described in Sec. 3.3, we train SELFmask using the selected salient masks as pseudo-masks. Here, we investigate the effect of the number of queries n_q in the Transformer decoder. For this, we train our model with $n_q=\{5, 10, 20, 50, 100\}$ on *DUTS-TR* and evaluate performance on the *HKU-IS* benchmark in terms of $\max F_\beta$ for two settings, (i) using ground-truth masks to pick the best mask out of n_q mask predictions, denoted as the SelfMask upper bound (UB); (ii) taking the mask with the highest object score, denoted SelfMask.

As shown in Figure 4, both SelfMask and SelfMask UB are fairly robust to the number of queries, initially increas-

Model	DUT-OMRON [63]			DUTS-TE [55]			ECSSD [47]		
	Acc	IoU	$\max F_\beta$	Acc	IoU	$\max F_\beta$	Acc	IoU	$\max F_\beta$
HS [62]	.843	.433	.561	.826	.369	.504	.847	.508	.673
wCtr [69]	.838	.416	.541	.835	.392	.522	.862	.517	.684
WSC [31]	.865	.387	.523	.862	.384	.528	.852	.498	.683
DeepUSPS [42]	.779	.305	.414	.773	.305	.425	.795	.440	.584
BigBiGAN [54]	.856	.453	.549	.878	.498	.608	.899	.672	.782
E-BigBiGAN [54]	.860	.464	.563	.882	.511	.624	.906	.684	.797
Melas-Kyriazi et al. [38]	.883	.509	-	.893	.528	-	.915	.713	-
LOST [49]	.797	.410	.473	.871	.518	.611	.895	.654	.758
LOST [49] + Bilateral solver [5]	.818	.489	.578	.887	.572	.697	.916	.723	.837
TokenCut [56]	.880	.533	.600	.903	.576	.672	.918	.712	.803
TokenCut [56] + Bilateral solver [5]	.897	.618	.697	.914	.624	.755	.934	.772	.874
pseudo-masks (Ours)	.811	.403	-	.845	.466	-	.893	.646	-
SELFMASK (Ours)	.901	.582	.680	.923	.626	.750	.944	.781	.889
SELFMASK (Ours) + Bilateral solver [5]	.919	.655	.852	.933	.660	.882	.955	.818	.956

Table 5: **Comparison to state-of-the-art unsupervised saliency detection methods on three salient object detection benchmarks.** For all metrics, higher number indicates better results. The best score per column is highlighted in bold. We observe that SELFMASK yields improved performance over prior state-of-the-art approaches across all benchmarks.

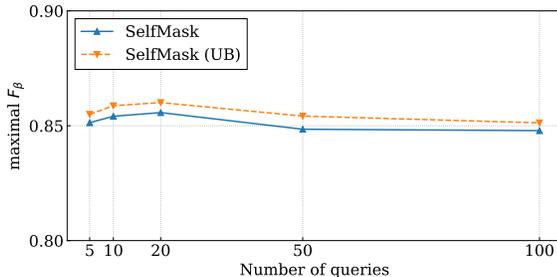


Figure 4: **Effect of number of queries on the performance of SelfMask on the HKU-IS dataset.** The model’s prediction and its upper bound, denoted by SelfMask and SelfMask (UB) each, are shown.

ing slightly with this hyperparameter (*i.e.*, predictions) before degrading after 20 queries. We conjecture that this is because a handful of queries are enough to localise the salient objects, while further predictions may make it challenging to appropriately rank the objectness of each prediction. For this reason, in the section that follows, we consider SELFMASK with 20 queries, and pick the query with the highest objectness as our prediction during inference.

4.4. Comparison to state-of-the-art unsupervised saliency detection methods

To compare with existing works on unsupervised SOD, we evaluate on three popular SOD benchmarks in terms of Acc., IoU, and $\max F_\beta$. Following [56], we also report results after post-processing predictions with the bilateral solver [5]. As shown in Tab. 5, while the pseudo-masks from spectral cluster voting already perform reasonably well compared to previous models, our self-trained model outperforms all existing approaches on all bench-

marks. This suggests both that the model can learn to generalise effectively from noisy masks, and that the objectness score trained with the ranking loss is effective for picking the best salient mask.

4.5. Broader Impact

This work contributes a new framework to deliver performant unsupervised salient object detection. As such, it offers the potential to underpin a range of societally beneficial applications that are bottlenecked by annotation costs. These include improved low-cost medical image segmentation, crop measurement from aerial imagery, and wildlife monitoring. However, low-cost segmentation is a powerful dual-use technology, and we caution against its deployment as a tool for unlawful surveillance and oppression.

5. Conclusion

In this work, we address the challenging problem of unsupervised salient object detection (SOD). For this, we first observe that self-supervised features exhibit significantly greater object segmentation potential with spectral clustering than with k -means. Inspired by this observation, we extract foreground regions among multiple masks generated with multiple types of features, and varying cluster numbers based on winner-takes-all voting. By using the selected masks as pseudo-masks, we train a saliency detection network and show promising results compared to previous unsupervised methods on various SOD benchmarks.

Acknowledgements. GS is supported by AI Factory, Inc. in Korea. WX is supported by Visual AI (EP/T028572/1). SA would like to thank Z. Novak and N. Novak for enabling his contribution. GS would like to thank Jaesung Huh for proof-reading.

References

- [1] Radhakrishna Achanta, Sheila Hemami, Francisco Estrada, and Sabine Susstrunk. Frequency-tuned salient region detection. In *CVPR*, 2009.
- [2] Relja Arandjelović and Andrew Zisserman. Object discovery with a copy-pasting gan. *arXiv:1905.11369*, 2019.
- [3] Pablo Arbeláez, Jordi Pont-Tuset, Jonathan Barron, Ferran Marques, and Jitendra Malik. Multiscale combinatorial grouping. In *CVPR*, 2014.
- [4] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *ICLR*, 2020.
- [5] Jonathan T. Barron and Ben Poole. The fast bilateral solver. In *ECCV*, 2016.
- [6] Adam Bielski and Paolo Favaro. Emergence of object segmentation in perturbed generative models. In *NeurIPS*, 2019.
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [8] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018.
- [9] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020.
- [10] 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.
- [11] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018.
- [12] Mickaël Chen, Thierry Artières, and Ludovic Denoyer. Unsupervised object segmentation by redrawing. In *NeurIPS*, 2019.
- [13] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *NeurIPS*, 2021.
- [14] Ming-Ming Cheng, Niloy J. Mitra, Xiaolei Huang, Philip H. S. Torr, and Shi-Min Hu. Global contrast based salient region detection. *TPAMI*, 2015.
- [15] Ming-Ming Cheng, Niloy J. Mitra, Xiaolei Huang, Philip H. S. Torr, and Shi-Min Hu. Global contrast based salient region detection. *TPAMI*, 2015.
- [16] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 1995.
- [17] 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.
- [18] Yuan Yuan Ding, Jing Xiao, and Jingyi Yu. Importance filtering for image retargeting. In *CVPR*, 2011.
- [19] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015.
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [21] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *TPAMI*, 2015.
- [22] Stas Goferman, Lihi Zelnik-Manor, and Ayellet Tal. Context-aware saliency detection. *TPAMI*, 2012.
- [23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.
- [24] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.
- [25] Phillip Isola, Daniel Zoran, Dilip Krishnan, and Edward H Adelson. Learning visual groups from co-occurrences in space and time. *arXiv:1511.06811*, 2015.
- [26] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *TPAMI*, 1998.
- [27] Tilke Judd, Krista Ehinger, Frédo Durand, and Antonio Torralba. Learning to predict where humans look. In *ICCV*, 2009.
- [28] Dominik A. Klein and Simone Frintrop. Center-surround divergence of feature statistics for salient object detection. In *ICCV*, 2011.
- [29] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *CVPR*, 2017.
- [30] Guanbin Li and Yizhou Yu. Visual saliency based on multi-scale deep features. In *CVPR*, 2015.
- [31] Nianyi Li, Bilin Sun, and Jingyi Yu. A weighted sparse coding framework for saliency detection. In *CVPR*, 2015.
- [32] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. In *ICCV*, 2017.
- [33] Tie Liu, Jian Sun, Nan-Ning Zheng, Xiaoou Tang, and Heung-Yeung Shum. Learning to detect a salient object. In *CVPR*, 2007.
- [34] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- [35] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.
- [36] Aravindh Mahendran, James Thewlis, and Andrea Vedaldi. Cross pixel optical-flow similarity for self-supervised learning. In *ACCV*, 2018.
- [37] Luca Marchesotti, Claudio Cifarelli, and Gabriela Csurka. A framework for visual saliency detection with applications to image thumbnailing. In *ICCV*, 2009.
- [38] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Finding an unsupervised image segmenter in each of your deep generative models. *arXiv:2105.08127*, 2021.

- [39] F. Milletari, N. Navab, and S. Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, 2016.
- [40] Vida Movahedi and James H. Elder. Design and perceptual validation of performance measures for salient object segmentation. In *CVPR Workshops*, 2010.
- [41] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *arXiv:2105.10497*, 2021.
- [42] Tam Nguyen, Maximilian Dax, Chaithanya Kumar Mummadi, Nhung Ngo, Thi Hoai Phuong Nguyen, Zhongyu Lou, and Thomas Brox. Deepusps: Deep robust unsupervised saliency prediction via self-supervision. In *NeurIPS*, 2019.
- [43] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016.
- [44] Yassine Ouali, Celine Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *CVPR*, 2020.
- [45] Michael Rubinstein, Ariel Shamir, and Shai Avidan. Improved seam carving for video retargeting. *ACM Transactions on Graphics (SIGGRAPH)*, 2008.
- [46] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *TPAMI*, 2000.
- [47] Jianping Shi, Qiong Yan, Li Xu, and Jiaya Jia. Hierarchical image saliency detection on extended cssd. *TPAMI*, 2015.
- [48] Jianping Shi, Qiong Yan, Li Xu, and Jiaya Jia. Hierarchical image saliency detection on extended cssd. *TPAMI*, 2016.
- [49] 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. In *BMVC*, 2021.
- [50] Bongwon Suh, Haibin Ling, Benjamin B. Bederson, and David W. Jacobs. Automatic thumbnail cropping and its effectiveness. In *ACM Symposium on User Interface Software and Technology*, 2003.
- [51] Jin Sun and Haibin Ling. Scale and object aware image retargeting for thumbnail browsing. In *ICCV*, 2011.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [53] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. In *ICML*, 2020.
- [54] Andrey Voynov, Stanislav Morozov, and Artem Babenko. Object segmentation without labels with large-scale generative models. In *ICML*, 2021.
- [55] 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.
- [56] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *CVPR*, 2022.
- [57] Yichen Wei, Fang Wen, Wangjiang Zhu, and Jian Sun. Geodesic saliency using background priors. In Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid, editors, *ECCV*, 2012.
- [58] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018.
- [59] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010.
- [60] Shuai Xie, Zunlei Feng, Y. Chen, Songtao Sun, Chao Ma, and Ming-Li Song. Deal: Difficulty-aware active learning for semantic segmentation. In *ACCV*, 2020.
- [61] Kaiming He Xinlei Chen, Saining Xie. An empirical study of training self-supervised vision transformers. *arXiv:2104.02057*, 2021.
- [62] Qiong Yan, Li Xu, Jianping Shi, and Jiaya Jia. Hierarchical saliency detection. In *CVPR*, 2013.
- [63] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via graph-based manifold ranking. In *CVPR*, 2013.
- [64] Xiaohang Zhan, Ziwei Liu, Ping Luo, Xiaoou Tang, and Chen Loy. Mix-and-match tuning for self-supervised semantic segmentation. In *AAAI*, 2018.
- [65] Dingwen Zhang, Junwei Han, and Yu Zhang. Supervision by fusion: Towards unsupervised learning of deep salient object detector. In *ICCV*, 2017.
- [66] Jianming Zhang and Stan Sclaroff. Exploiting surround-ness for saliency detection: A boolean map approach. *TPAMI*, 2016.
- [67] Jianming Zhang, Stan Sclaroff, Zhe Lin, Xiaohui Shen, Brian Price, and Radomír Mech. Minimum barrier salient object detection at 80 fps. In *ICCV*, 2015.
- [68] Jing Zhang, Tong Zhang, Yuchao Dai, Mehrtash Harandi, and Richard Hartley. Deep unsupervised saliency detection: A multiple noisy labeling perspective. In *CVPR*, 2018.
- [69] Wangjiang Zhu, Shuang Liang, Yichen Wei, and Jian Sun. Saliency optimization from robust background detection. In *CVPR*, 2014.