
Understanding Self-Supervised Features for Learning Unsupervised Instance Segmentation

Paul Engstler* Luke Melas-Kyriazi* Christian Rupprecht Iro Laina
University of Oxford
{paule, lukemk, chrisr, iro}@robots.ox.ac.uk

Abstract

Self-supervised learning (SSL) can be used to solve complex visual tasks without human labels. Self-supervised representations encode useful semantic information about images, and as a result, they have already been used for tasks such as unsupervised semantic segmentation. In this paper, we investigate self-supervised representations for instance segmentation without any manual annotations. We find that the features of different SSL methods vary in their level of instance-awareness. In particular, DINO features, which are known to be excellent semantic descriptors, lack behind MAE features in their sensitivity for separating instances.

1 Introduction

Self-supervised visual representation learning has received significant attention over the past years. The goal is to learn strong feature representations from unlabeled images, which are transferable to downstream tasks. Most methods are currently designed for object-centric datasets like ImageNet [42], yet most downstream applications are scene-centric, containing multiple objects of interest.

We know that such representations can “store” several *different* concepts even when they appear simultaneously in an image [37]. It is thus reasonable to assume that they can be used to discover *multiple* objects in an image, without supervision. However, is it also possible to locate each concept spatially in the image? To answer this question, recent work has moved beyond concept discovery at the image level to pixel-level localization and segmentation of unsupervised concepts [2, 40, 45, 46, 54]. Different from fine-tuning models for detection and segmentation tasks, this line of work exploits self-supervised features directly without using any annotations, for example through clustering [24, 40, 50] for semantic segmentation.

Instance segmentation, on the other hand, has received less attention. This task is particularly challenging because it requires recognizing and segmenting each individual object in an image, while semantic segmentation treats multiple objects of the same category as one. In the unsupervised setting, instance segmentation relies on acquiring a notion of objectness, spanning diverse scales and appearances, that is more difficult to achieve than single-object discovery, which amounts to the most salient region in the image. As a result, prior work on this task mines a small set of pseudo-masks per image and uses that to bootstrap instance segmentation [51, 53, 54] or object detection [46] models.

In this work, we provide valuable insights with respect to the *choice of self-supervised features* that can be used to obtain these mask proposals. We find that the training objective in self-supervised models influences their ability to distinguish between objects, specifically different instances of the same semantic category.

*Equal contribution.

2 Related Work

Self-Supervised Representation Learning (SSL). The goal of SSL is to learn view-invariant representations from unlabeled data, which are then transferred to downstream tasks via task-specific finetuning. Following instance discrimination [21, 61], the majority of works focus on contrastive [6, 16, 18, 27, 30, 33, 36, 47], negative-free [9, 14, 17, 23, 68], and clustering-based [4, 12, 13, 35, 38] learning. Inspired by advances in NLP, masked image modeling [7, 25, 58, 65, 69] has emerged as an alternative approach. Meanwhile, the use of Vision Transformers (ViTs) [20, 49] has contributed significantly to the performance of self-supervised methods. Finally, there is a growing interest in learning dense [41, 43, 56, 64] or region-based representations [8, 31, 32, 59, 62, 63, 66, 67], since they may be better suited for dense prediction tasks.

Unsupervised Object Discovery. Unsupervised object discovery aims to localize objects in images by predicting bounding boxes. Earlier approaches were based on adversarial learning [1, 3, 11, 15], while, more recently, several works [2, 40, 45, 46, 51, 53, 57] have explored the use of self-supervised features. In particular, [14] first showed that the self-attention of DINO-ViT could be used for foreground segmentation. Follow-up works [40, 46, 51, 57] showed that it can also be used for salient object discovery, which in turn can be used to train unsupervised object detectors [10, 46, 53].

Unsupervised Semantic Segmentation. Recent work in unsupervised semantic segmentation can be split into two categories: (1) methods that utilize pre-trained self-supervised models for initialization [19, 50, 51, 60, 70] and (2) methods that directly exploit off-the-shelf self-supervised representations (e.g., DINO [14]) to obtain and cluster pseudo-masks and train a segmentation network [24, 40, 51]. These works demonstrate that self-supervised ViT features encode well-localized *semantics*, but do not investigate whether these features can discriminate *instances* of the same semantic class. We aim to answer this question in Sec. 3, finding a notable difference between models trained with discriminative (e.g., contrastive) and generative (e.g., autoencoding) objectives.

Unsupervised Instance Segmentation. Unsupervised instance segmentation refers to the task of discovering and segmenting individual object instances in an image. State-of-the-art methods typically bootstrap class-agnostic instance segmentation networks using coarse masks extracted from self-supervised feature extractors. FreeSOLO [54] uses densely-learned features (DenseCL [56]), while Exemplar-FreeSOLO [34] additionally uses a pool of exemplar images to extract objects. MaskDistill [51] and CutLER [53] leverage DINO features; [51] follows a single-object discovery technique, while [53] obtains multiple masks through repeated applications of the NCut algorithm [44].

3 Analysis of *Instance-ness* in Self-Supervised Vision Transformers

We investigate whether features from self-supervised transformers exhibit a general notion of *instance-ness* that can be used to discover and discriminate between object instances rather than semantic categories. Among SSL methods, DINO features have been frequently used for dense unsupervised tasks. We note that different pre-training objectives may result in models with different properties and thus focus our investigation on models trained with a variety of objectives: contrastive (MoCo v3 [18]), self-distillation (DINO [14], MSN [5]), image reconstruction (MAE [25]), as well as patch-level (SelfPatch [67]) and dense (DenseCL [56]) pre-training. All models use ViT-B/16 backbones, except for SelfPatch (only available for ViT-S/16) and DenseCL (ResNet-101 [29]).

First, we examine all models with respect to their ability to encode instances in the learned representation, without further task-specific finetuning. Similar to prior work [2, 40], to decompose an image into a set of meaningful regions, we compute an affinity matrix using cosine similarity in feature space and apply k -way spectral clustering spatially to generate a set of masks. In particular, we extract features from the keys of the final self-attention block of ViTs. To account for the varying number of objects in images and to allow for the possibility of segmenting objects at different granularities (Fig. 1), we apply spectral clustering multiple times for different values of $k \in K$ and accumulate all masks. Finally, if the resulting masks have non-connected components, we split these off into separate masks.

It is already known from prior work that this process results in regions that likely correspond to semantic entities or objects. To understand the degree to which these regions overlap with real objects in images, we compute the mean average recall of these masks against ground truth mask annotations in common datasets (MS COCO val2017 [39], PASCAL VOC 2012 [22]) and report our findings




Figure 1: **Varying the value of k .** Depending on the number of instances and semantic classes in an image, different values of k capture scene elements at different levels of granularity. Higher values of k separate the different instances. This example demonstrates why we use multiple values of k to generate mask proposals and why it is necessary to filter these using a saliency map.

Table 1: **Analysis of SSL methods for instance segmentation.** We report the mean average recall (in %) of feature extractors for different numbers of instances of the same semantic class. For each image and feature extractor, we apply spectral clustering to the extracted features with $K = \{2, 3, 4, 5, 6\}$.

Features	Instances	COCO						PASCAL VOC					
		1	2	3	4	5	6+	1	2	3	4	5	6+
MAE (ViT-B/16) [25]		16.8	8.6	6.3	3.2	2.9	1.1	31.7	18.1	9.3	4.5	4.1	2.2
DINO (ViT-B/16) [14]		21.0	5.9	3.3	1.9	1.9	0.6	38.3	9.9	4.7	1.9	1.7	0.7
SelfPatch (ViT-S/16) [67]		17.4	7.7	5.5	2.8	2.6	1.0	32.8	13.9	7.1	4.3	3.4	1.5
MSN (ViT-B/16) [5]		19.8	5.7	3.4	1.9	1.8	0.6	35.6	10.1	4.9	1.7	1.3	0.7
MoCo-v3 (ViT-B/16) [18]		19.1	5.0	3.0	1.6	1.7	0.6	37.3	9.1	4.3	1.6	1.6	0.7
DenseCL (RN-101) [56]		13.3	3.2	1.8	0.9	0.9	0.3	24.3	5.2	2.3	1.0	0.8	0.3

in Table 1. We evaluate recall based on the number of instances of the same semantic category that co-occur in an image, *e.g.* in an image with two objects of the same class and one object of another class, we compute and report recall for the two instances and the remaining single instance separately.

Interestingly, we find that MAE is much better at discriminating between multiple object instances (2+) than other models. This finding may be explained by the fact that the pixel-wise reconstruction objective encourages the learning of representations of lower semanticity [25]. As a result, it shows a lower tendency to group objects of the same or similar semantic category together, compared to, *e.g.*, DINO. This holds true even against models trained with patch-level and pixel-level discrimination (*e.g.*, SelfPatch). We demonstrate this also qualitatively in the Appendix (Figure 3), comparing MAE and DINO masks obtained after spectral clustering while varying k . Especially for higher values of k , MAE features tend to *spatially* decompose an image, whereas DINO tends to divide objects further into semantic parts. The spatial bias in MAE likely explains its superiority at separating instances, though MAE masks remain of generally lower quality than DINO.

At the same time, our results in Table 1 indicate that DINO and other contrastive methods such as MSN and MoCo-v3 are superior at locating *single* objects; this finding matches prior work that uses DINO features for single-object image segmentation.

4 Self-Training for Instance Segmentation

To evaluate how these findings transfer to the task of unsupervised instance segmentation, we implement a simple method to generate mask proposals and use these to train a segmentation network, using the feature extractors considered above. An overview of our approach is shown in Figure 2.

We adopt a two-part approach to generate mask proposals, both parts based on SSL feature extractors. The first is the generation of pseudo-masks spanning the entire image area. For each of the feature extractors we follow the procedure outlined previously and apply multi- k -way clustering with $K = \{2, 3, 4, 5\}$ (resulting in 14 overlapping masks per image). Since we cluster features spatially over the entire image area, the resulting clusters (masks) will inevitably also enclose background regions. As we are interested in *instance* segmentation (*e.g.*, people, cars) rather than *stuff* segmentation (*e.g.*, grass, sky), we need to eliminate masks that are less likely to correspond to objects.

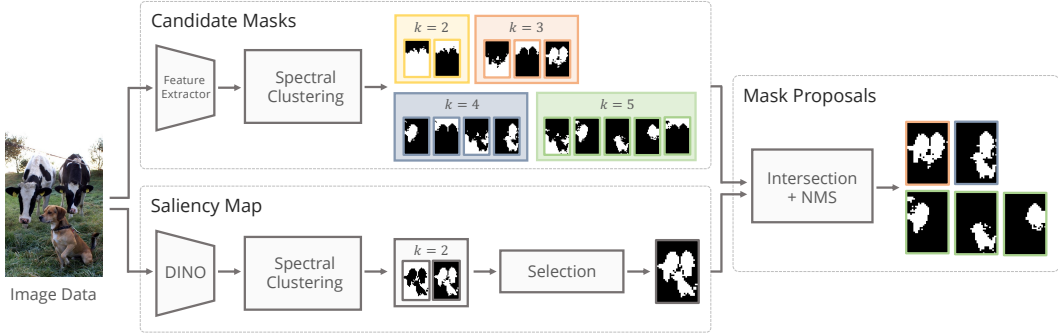


Figure 2: **Overview of our approach to produce mask proposals.** (1) Given an image and a self-supervised feature extractor, we generate a set of candidate masks by applying spectral clustering multiple times with different values for k , (2) we obtain a saliency map via spectral clustering with $k = 2$ on DINO features, and (3) we select the candidate masks that strongly intersect with the saliency map as our final mask proposals (followed by non-maximum suppression for deduplication).

Table 2: **Unsupervised instance segmentation performance (COCO val2017).** We evaluate segmentation models trained with mask proposals from different feature extractors.

Feature Extractor	AP ₅₀	AP ₇₅	AP	AR ₁	AR ₁₀	AR ₁₀₀
MAE [25]	12.1	3.7	5.2	3.8	10.9	18.3
DINO [14]	11.8	3.6	5.0	3.7	11.4	18.0
SelfPatch [67]	6.3	2.2	2.8	3.4	4.9	4.9
MSN [5]	9.0	2.7	3.9	3.9	6.2	6.2
MoCo v3 [18]	7.2	2.5	3.3	3.4	6.4	9.6
DenseCL [56]	7.7	2.5	3.3	3.3	7.3	10.2

Therefore, the second part is saliency-based masking to filter possibly noisy candidates, such as those corresponding to background regions. Although not *all* objects in a scene are necessarily salient, salient image regions are more likely to contain object instances than non-salient ones. Given the success of DINO features for this task [14, 40, 46, 57] and our findings in Table 1, we apply spectral clustering with $k = 2$ on DINO features. We select the mask that shares fewer pixels with the image boundary as the foreground and use it to filter the mask candidates from the first stage such that only masks for salient objects remain.

Finally, we use the resulting mask proposals (after non-maximum suppression (NMS) for deduplication) to train an instance segmentation network. We train a SOLOv2 [55] architecture using only pseudo-masks, following FreeSOLO [54]. For more training details, please refer to Appendix A.2.1.

We train a segmenter for each of the feature extractors using their respective mask proposals. This allows us to assess their performance in unsupervised instance segmentation and examine whether our initial observations align with post-training results. We report average precision and recall on MS COCO val2017 in Table 2. We observe that the model trained with masks from MAE remains the best-performing one. Interestingly, the DINO-based model significantly narrows the performance gap after training, whereas other feature extractors perform worse. We hypothesize that this comes from the fact that DINO masks are generally cleaner and capture object boundaries better, which is important for learning. This could also explain why the vast majority of the current state of the art (see Table 3) in unsupervised instance segmentation performs well, despite leveraging DINO features.

5 Conclusion

In this work, we study the properties of SSL methods for the task of instance segmentation. We find that different learning signals result in features with varying characteristics. DINO learns highly semantic features and is thus widely used for semantic tasks. We find that image reconstruction, *e.g.* MAE, tasks are better suited to discriminate instances of the same class inside an image. This is an overlooked property that can potentially be used in many instance-specific downstream tasks.

Acknowledgements. P. E., C. R., and I. L. are supported by ERC-UNION- CoG-101001212. P. E. is also supported by Meta Research and a fellowship from the German Academic Exchange Service (DAAD). L. M. K. is supported by the Rhodes Trust. C. R. and I. L. are also supported by VisualAI EP/T028572/1.

References

- [1] Rameen Abdal, Peihao Zhu, Niloy Mitra, and Peter Wonka. Labels4free: Unsupervised segmentation using stylegan. In *Proc. ICCV*, 2021. 2
- [2] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2(3):4, 2021. 1, 2
- [3] Relja Arandjelović and Andrew Zisserman. Object discovery with a copy-pasting gan. *arXiv preprint arXiv:1905.11369*, 2019. 2
- [4] YM Asano, C Rupprecht, and A Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations*, 2019. 2
- [5] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *European Conference on Computer Vision*, pages 456–473. Springer, 2022. 2, 3, 4
- [6] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15509–15519, 2019. 2
- [7] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEit: BERT pre-training of image transformers. In *International Conference on Learning Representations*, 2022. 2
- [8] Amir Bar, Xin Wang, Vadim Kantorov, Colorado J Reed, Roei Herzig, Gal Chechik, Anna Rohrbach, Trevor Darrell, and Amir Globerson. Detreg: Unsupervised pretraining with region priors for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14605–14615, 2022. 2
- [9] Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-invariance-covariance regularization for self-supervised learning. In *International Conference on Learning Representations*, 2022. 2
- [10] Adam Bielski and Paolo Favaro. Move: Unsupervised movable object segmentation and detection. *arXiv preprint arXiv:2210.07920*, 2022. 2
- [11] Adam Jakub Bielski and Paolo Favaro. Emergence of object segmentation in perturbed generative models. *Advances in Neural Information Processing Systems (NIPS)*, 32, 2019. 2
- [12] M. Caron, P. Bojanowski, A. Joulin, and M. Douze. Deep clustering for unsupervised learning of visual features. In *Proc. ECCV*, 2018. 2
- [13] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [14] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 2, 3, 4, 11
- [15] Mickaël Chen, Thierry Artières, and Ludovic Denoyer. Unsupervised Object Segmentation by Redrawing. In *Proc. NeurIPS*, volume 32, 2019. 2
- [16] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2
- [17] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 2
- [18] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9640–9649, 2021. 2, 3, 4
- [19] Jang Hyun Cho, Utkarsh Mall, Kavita Bala, and Bharath Hariharan. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16794–16804, 2021. 2
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 2
- [21] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in neural information processing systems*, pages 766–774, 2014. 2
- [22] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 2
- [23] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray

- kavakcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284. Curran Associates, Inc., 2020. 2
- [24] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T. Freeman. Unsupervised semantic segmentation by distilling feature correspondences. In *Proc. ICLR*, 2022. 1, 2
- [25] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 2, 3, 4
- [26] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners, 2021. 11
- [27] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [28] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 12
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [30] Olivier Henaff. Data-efficient image recognition with contrastive predictive coding. In *International conference on machine learning*, pages 4182–4192. PMLR, 2020. 2
- [31] Olivier J Hénaff, Skanda Koppula, Jean-Baptiste Alayrac, Aaron Van den Oord, Oriol Vinyals, and João Carreira. Efficient visual pretraining with contrastive detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10086–10096, 2021. 2
- [32] Olivier J. Hénaff, Skanda Koppula, Evan Shelhamer, Daniel Zoran, Andrew Jaegle, Andrew Zisserman, João Carreira, and Relja Arandjelović. Object discovery and representation networks. In *Proc. ECCV*, pages 123–143, 2022. 2
- [33] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2019. 2
- [34] Taoseef Ishtiaq, Qing En, and Yuhong Guo. Exemplar-freesolo: Enhancing unsupervised instance segmentation with exemplars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15424–15433, June 2023. 2, 10
- [35] Xu Ji, Joao F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9865–9874, 2019. 2
- [36] Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. *Proc. NeurIPS*, 2020. 2
- [37] Iro Laina, Yuki M Asano, and Andrea Vedaldi. Measuring the interpretability of unsupervised representations via quantized reversed probing. In *Proc. ICLR*, 2022. 1, 8
- [38] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *International Conference on Learning Representations*, 2021. 2
- [39] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 8
- [40] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Deep spectral methods: A surprisingly strong baseline for unsupervised semantic segmentation and localization. In *Proc. CVPR*, pages 8364–8375, 2022. 1, 2, 4
- [41] Pedro O. Pinheiro, Amjad Almahairi, Ryan Benmalek, Florian Golemo, and Aaron C Courville. Unsupervised learning of dense visual representations. *Advances in Neural Information Processing Systems*, 33:4489–4500, 2020. 2
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015. 1, 10
- [43] Ramprasaath R Selvaraju, Karan Desai, Justin Johnson, and Nikhil Naik. Casting your model: Learning to localize improves self-supervised representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11058–11067, 2021. 2
- [44] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000. 2
- [45] Gyungin Shin, Samuel Albanie, and Weidi Xie. Unsupervised salient object detection with spectral cluster voting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3971–3980, 2022. 1, 2
- [46] 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 *Proc. BMVC*, November 2021. 1, 2, 4
- [47] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020. 2

- [48] Zhi Tian, Chunhua Shen, Xinlong Wang, and Hao Chen. Boxinst: High-performance instance segmentation with box annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5443–5452, 2021. 10
- [49] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. 2
- [50] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Unsupervised semantic segmentation by contrasting object mask proposals. In *International Conference on Computer Vision*, 2021. 1, 2
- [51] Wouter Van Gansbeke, Simon Vandenhende, and Luc Van Gool. Discovering object masks with transformers for unsupervised semantic segmentation. *arXiv preprint arXiv:2206.06363*, 2022. 1, 2, 10, 11
- [52] Weiyao Wang, Matt Feiszli, Heng Wang, and Du Tran. Unidentified video objects: A benchmark for dense, open-world segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10776–10785, 2021. 10
- [53] Xudong Wang, Rohit Girdhar, Stella X Yu, and Ishan Misra. Cut and learn for unsupervised object detection and instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3124–3134, 2023. 1, 2, 10, 11
- [54] Xinlong Wang, Zhiding Yu, Shalini De Mello, Jan Kautz, Anima Anandkumar, Chunhua Shen, and Jose M Alvarez. Freesolo: Learning to segment objects without annotations. In *Proc. CVPR*, pages 14176–14186, 2022. 1, 2, 4, 10, 11, 12
- [55] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. *Advances in Neural information processing systems*, 33:17721–17732, 2020. 4, 10, 11, 12
- [56] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 3, 4
- [57] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14543–14553, 2022. 2, 4
- [58] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14668–14678, 2022. 2
- [59] Fangyun Wei, Yue Gao, Zhirong Wu, Han Hu, and Stephen Lin. Aligning pretraining for detection via object-level contrastive learning. *Advances in Neural Information Processing Systems*, 34:22682–22694, 2021. 2
- [60] Xin Wen, Bingchen Zhao, Anlin Zheng, Xiangyu Zhang, and Xiaojuan Qi. Self-supervised visual representation learning with semantic grouping. *arXiv preprint arXiv:2205.15288*, 2022. 2
- [61] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. 2
- [62] Enze Xie, Jian Ding, Wenhai Wang, Xiaohang Zhan, Hang Xu, Peize Sun, Zhenguo Li, and Ping Luo. Detco: Unsupervised contrastive learning for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8392–8401, 2021. 2
- [63] Jiahao Xie, Xiaohang Zhan, Ziwei Liu, Yew Soon Ong, and Chen Change Loy. Unsupervised object-level representation learning from scene images. *Advances in Neural Information Processing Systems*, 34:28864–28876, 2021. 2
- [64] Zhenda Xie, Yutong Lin, Zheng Zhang, Yue Cao, Stephen Lin, and Han Hu. Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [65] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. 2
- [66] Ceyuan Yang, Zhirong Wu, Bolei Zhou, and Stephen Lin. Instance localization for self-supervised detection pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3987–3996, 2021. 2
- [67] Sukmin Yun, Hankook Lee, Jaehyung Kim, and Jinwoo Shin. Patch-level representation learning for self-supervised vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8354–8363, 2022. 2, 3, 4
- [68] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. 2
- [69] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Image BERT pre-training with online tokenizer. In *International Conference on Learning Representations*, 2022. 2
- [70] Adrian Ziegler and Yuki M Asano. Self-supervised learning of object parts for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14502–14511, 2022. 2

A Appendix

A.1 Qualitative Analysis of Instance-ness in Self-Supervised Vision Transformers

In Section 3, we have measured the mean average recall of masks produced by clustering features from different self-supervised networks for a different number of instances of the same semantic class on MS COCO val2017 and PASCAL VOC 2012. We used this experiment to better understand the properties of the raw feature extractors, *i.e.* how they “see” objects, as well as a proxy for final instance segmentation performance. Our analysis indicates that features generated by MAE lead to an overall higher recall when multiple instances exist, *i.e.* features produced by MAE are better at discriminating multiple instances of the same semantic class, while DINO features are superior at detecting single instances.

We illustrate this observation further in Fig. 3 with qualitative examples. We show examples from COCO val2017 and compare the segments produced by MAE and DINO after spectral clustering while varying k .

Next, we record some qualitative observations from our experience with these methods, which we hope will be of use to researchers working with them in the future:

- Spectral clustering of MAE features has a tendency to separate instances of the same semantic class in the same image, whereas DINO features are more likely to produce a purely semantic grouping (*e.g.*, Fig. 3 (b)). Yet there is ambiguity as to how semantic classes are “seen” by self-supervised models. An example is shown in Fig. 3 (e), in which all objects correspond to the same semantic class *orange*. The image shows a whole orange, a halved orange, and three individual peeled segments, which represent different *states* of the orange. Since these states have distinct appearances, DINO is able to partially discriminate between these objects, and its behavior comes closer to MAE in resembling “instance awareness” in this example. This finding aligns with the observation of [37] and is likely true for most self-supervised models.
- For smaller values of k (especially if k is smaller than the number of objects in the image), the decomposition remains shallow. Even if k is sufficiently large to capture all objects, a number of clusters are often used to represent the nuances of the background, and as a result, different foreground objects group together in the remaining clusters.
- With an increasing number of clusters (over-clustering), MAE features tend to *spatially* decompose an image, whereas, with DINO, the level of semantic detail seems to increase, *e.g.* objects get subdivided into parts (Fig. 3 (a), (b), (d)).
- Applying spectral clustering to DINO features with $k = 2$ produces higher quality saliency candidates (Fig. 3 (b), (f)), which matches the results of the ablation in Tab. 7.
- Both feature extractors struggle with very small objects, *e.g.* the dog (Fig. 3 (c)) or the baseball bat and glove (Fig. 3 (f)). In these examples, small objects become part of larger, more prominent objects or part of the background.
- Complex scenes without prominent foreground objects, *e.g.* scenes “seen” from a distance (Fig. 3 (g)), yield segmentations that do not align with the valid object categories in COCO — for example, the image is partitioned into trees, buildings, and street.

A.2 Self-Training for Instance Segmentation, Addendum

A.2.1 Datasets and Training Details

MS COCO 2017. The Common Objects in COntext dataset comprises images of “complex everyday scenes containing common objects in their natural context” [39]. We use the 2017 release in our method, which consists of train2017, unlabeled2017, and val2017 sets. The training set contains 118,287 images with 860,001 polygon annotations for 80 semantic classes sourced from human annotators. The unlabeled image set comprises 123,403 images without any annotations. We use the train2017 and unlabeled2017 sets for generating mask proposals with our method to bootstrap an instance segmentation network (no ground truth annotations are used). We then evaluate it on val2017, which contains 5,000 images with 36,781 polygon annotations, in a class-agnostic (*i.e.*, classes ignored) setting.


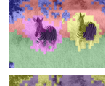





	Image	GT Annotations	Features	k				
				2	3	4	5	6
(a)			MAE					
			DINO					
(b)			MAE					
			DINO					
(c)			MAE					
			DINO					
(d)			MAE					
			DINO					
(e)			MAE					
			DINO					
(f)			MAE					
			DINO					
(g)			MAE					
			DINO					

Figure 3: **Spectral clustering of self-supervised features from MAE and DINO with $k \in \{2, \dots, 6\}$.** We show a sample of images from COCO val2017 with different numbers of ground-truth (GT) annotations. For each feature extractor, we show all masks for individual values of k . Each mask is shown with a different color, where *no color* represents the absence of annotations. Best viewed on screen.

Table 3: **Unsupervised class-agnostic instance segmentation.** We report the performance of state-of-the-art methods on this task on common benchmarks. For FreeSOLO, we report additional results marked with † based on our reproduction. FreeSOLO and Exemplar-FreeSOLO use DenseCL features, while MaskDistill and CutLER use DINO. Our setting is most comparable to FreeSOLO.

Method	PASCAL VOC			COCO20k			COCO (val2017)			UVO		
	AP ₅₀	AP ₇₅	AP	AP ₅₀	AP ₇₅	AP	AP ₅₀	AP ₇₅	AP	AP ₅₀	AP ₇₅	AP
FreeSOLO [54]	9.8 [†]	0.2 [†]	2.3 [†]	10.2 [†]	3.5 [†]	4.6 [†]	9.8	2.9	4.0	12.7	3.0	4.8
Exemplar-FS [34]	-	-	-	-	-	-	13.2	6.3	8.4	14.2	7.3	9.2
MaskDistill [51]	24.3	6.9	9.9	6.8	2.1	2.9	-	-	-	-	-	-
CutLER [53]	-	-	-	19.6	10.0	9.2	18.9	9.7	9.2	22.8	8.0	10.1
Ours (MAE)	18.8	0.6	4.6	12.3	3.7	5.3	12.0	3.7	5.2	12.9	2.6	4.7

Training Details. We generate the mask proposals using pre-trained ViT-B/16 vision transformers (except for SelfPatch, which is only available for ViT-S/16, and DenseCL, which is a ResNet-101 model) and a ViT-B/16 DINO model for saliency estimation. These networks are all pre-trained on ImageNet [42] without supervision. When extracting features, we do not apply any data augmentation to the images. We choose $K = \{2, 3, 4, 5\}$ as the set of values for k to generate candidate masks and further split non-connected components into separate masks. To obtain the final set of mask proposals, we first discard masks with an intersection score of less than 0.5 when compared to the saliency map. Then, we apply NMS with a threshold of 0.8 to remove any duplicates from the masks that remain.

Unless otherwise specified, we use SOLOv2 [54, 55] as the instance segmentation network, initialized with self-supervised weights (DenseCL, ResNet-101). We employ copy & paste augmentation as well as the BoxInst loss [48, 54]. We train the network with a frozen backbone for 60k steps with our mask proposals as targets. The stochastic gradient descent optimizer is used with a learning rate of 0.00025 and a batch size of 4. We follow the default parameters and loss setup of [54], and perform only a single round of training (which differs from [54] that employs two rounds with 30k steps each). The trained network returns a confidence score for each detection/mask, which allows us to set a threshold and retain the most confident predictions. Training takes 48 hours on an NVIDIA A40 GPU. We will publish code and pre-trained models to fully reproduce our results.

A.2.2 State-Of-The-Art on Unsupervised Instance Segmentation

In Tab. 3, we present the performance of recent works for the task of unsupervised class-agnostic instance segmentation and provide a comparison to our simple approach (using MAE features to generate mask proposals).

Apart from COCO val2017 (discussed above), we show results for:

1. PASCAL VOC, a standard benchmark for objection detection with annotations spanning 20 object categories that overlap with the COCO categories. It contains 5,011 images and 15,662 object annotations.
2. COCO20k, a subset of the COCO 2014 trainval split, with 19,817 images (143,951 object annotations) and excludes objects flagged as *crowd*.
3. The Unidentified Video Objects dataset, which provides a benchmark for “open-world class-agnostic object segmentation in videos” [52]. It includes dense segmentation annotations of video frames from human annotators without any class information to evaluate the degree to which models recognize object categories not seen during training. For evaluation, the validation set of UVO-Sparse 1.0 is used, which comprises a total of 65,588 temporally sparse frame annotations for 8,124 frames of 2,708 videos.

As our approach follows the setup of FreeSOLO [54] (architecture, loss formulation, and hyper-parameters), it permits direct comparison and shows clear improvements across datasets. Exemplar-FreeSOLO [34], with its addition of a randomly drawn pool of exemplars used in a contrastive learning loss, shows stronger improvements.

MaskDistill [51] trains separate models for each dataset, as opposed to all previously mentioned methods that are solely trained on COCO train2017 and unlabeled2017. Therefore, comparability is limited.

CutLER [53] trains on ImageNet, exploiting its object-centric prior, as most images contain a single object in the center of the frame. Due to its strong instance discrimination abilities, CutLER is the current state-of-the-art method for this task.

A.3 Ablations

We perform various ablations to investigate the effectiveness of our design choices surrounding the self-training setup. For simplicity, we focus on annotations produced by the best-performing feature extractor, *i.e.* MAE features (filtered with saliency maps obtained from DINO features), instead of considering all feature extractors individually. All models are evaluated after training an instance segmentation network for 30k steps on COCO val2017.

Choice of k . The set K of values for k determines the number of objects that can be discovered but also the number of spurious masks that might appear. For images with a single object, large values of k might over-segment objects into parts, whereas for images with multiple objects, too small values might lead to masks that do not capture individual instances or fail to segment any objects at all. In Tab. 4, we show that for complex datasets such as COCO, where images contain various numbers of objects, no single choice of k emerges as the best option. Generating masks from multiple values of k , however, provides an opportunity to deal with this variance as evidenced by the significantly higher recall.

Generating the Saliency Map. The saliency map, which we use to retain masks that are likely to correspond to objects in an image while discarding those that are not, is a central piece in our self-training setup. In Tab. 5, we compare saliency maps generated from MAE and DINO features, showing that DINO features are far superior.

Non-Connected Component Splitting. Although feature extractors yield mask proposals that segment individual instances, they may still generate mask candidates with multiple objects in a single mask, especially for smaller values of k (Fig. 1). If such masks contain disconnected components, we assume they are separate objects and thus split them apart. This is particularly helpful as we apply NMS as part of our mask generation process: If a larger k does indeed capture these objects separately, the split components can be deduplicated, effectively eliminating the poor original mask candidate. In Tab. 6, we show that this splitting of non-connected components in masks is beneficial.

Segmentation Architecture. As the final part of our self-training design, we use an instance segmentation network, which refines masks and helps to increase the number of objects that are detected. We observe that the setup of SOLOv2 [55], adapted to unsupervised segmentation [54], is superior to an off-the-shelf Mask R-CNN for unsupervised instance segmentation, as seen in Tab. 7.

Table 4: **Ablation of different values for k .** Multi-way clustering ($K = \{2, 3, 4, 5\}$) performs best.

k	AP	AP _S	AP _M	AP _L	AR	AR _S	AR _M	AR _L
2	0.2	0.0	0.0	0.3	0.5	0.0	0.0	2.3
3	2.7	0.7	1.5	8.8	4.3	0.1	1.3	16.1
4	3.7	0.8	2.4	12.3	6.0	0.1	2.3	21.7
5	3.7	0.7	3.0	11.4	6.5	0.2	3.3	22.1
2-5	4.5	1.4	6.0	12.8	16.5	1.7	19.1	38.6

Table 5: **Ablation of the saliency map.** We also compare different self-supervised models (MAE, DINO) for generating the map. Without the saliency filter (None) performance drops significantly.

Saliency Mask	AP ₅₀	AP ₇₅	AP	AR ₁	AR ₁₀	AR ₁₀₀
None	2.3	0.2	0.6	0.9	3.8	9.4
Ours (MAE [26])	5.9	1.3	2.2	2.2	5.1	7.3
Ours (DINO [14])	10.7	2.9	4.5	3.5	9.0	16.5

Table 6: **Ablation of non-connected component splitting (NCCS)** during the generation of candidate masks.

NCCS	AP ₅₀	AP ₇₅	AP	AR ₁	AR ₁₀	AR ₁₀₀
X	7.2	2.5	3.3	3.4	6.4	9.6
✓	10.7	2.9	4.5	3.5	9.0	16.5

Table 7: **Ablation of the segmentation architecture.** Both networks are trained using MAE mask proposals.

Architecture	AP ₅₀	AP ₇₅	AP	AR ₁	AR ₁₀	AR ₁₀₀
Mask R-CNN [28]	6.8	1.8	2.7	2.2	7.3	11.9
SOLOv2 [54, 55]	10.7	2.9	4.5	3.5	9.0	16.5