

Unsupervised Discovery of the Long-Tail in Instance Segmentation Using Hierarchical Self-Supervision

Zhenzhen Weng, Mehmet Giray Ogut, Shai Limonchik, Serena Yeung
Stanford University

{zzweng, giray98, shailk, syyeung}@stanford.edu

Abstract

Instance segmentation is an active topic in computer vision that is usually solved by using supervised learning approaches over very large datasets composed of object level masks. Obtaining such a dataset for any new domain can be very expensive and time-consuming. In addition, models trained on certain annotated categories do not generalize well to unseen objects. The goal of this paper is to propose a method that can perform unsupervised discovery of long-tail categories in instance segmentation, through learning instance embeddings of masked regions. Leveraging rich relationship and hierarchical structure between objects in the images, we propose self-supervised losses for learning mask embeddings. Trained on COCO [34] dataset without additional annotations of the long-tail objects, our model is able to discover novel and more fine-grained objects than the common categories in COCO. We show that the model achieves competitive quantitative results on LVIS [17] as compared to the supervised and partially supervised methods.

1. Introduction

Instance segmentation is a crucial problem that has a wide range of applications in various real-world applications such as autonomous driving and medical imaging. Recent approaches [18, 19, 4, 35] have shown impressive results on large-scale datasets in parsing diverse real-world scenes into informative semantics or instance maps. Most of the existing works so far assume a fully supervised setting, where the instance level segmentation masks are available during training. However, relying on human annotated segmentation labels has a few obvious drawbacks. Since most of the existing segmentation datasets contain a small set of annotated categories (e.g. 20 in PASCAL VOC [10] and 80 in COCO [34]), models trained on these categories are not able to generalize well to novel and long-tail objects present in the real world. Despite the ease of cap-

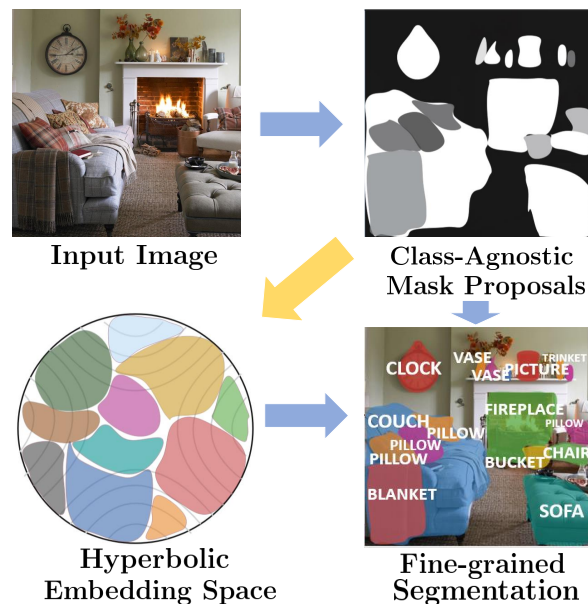


Figure 1. We propose an instance segmentation method that is able to discover the long-tail objects through self-supervised representation learning. Without having access to ground truth annotations of long-tail objects during training, our method is able to produce fine-grained segmentation result of novel objects.

turing a large number of images nowadays, extending the set of annotated categories is still very expensive and time-consuming. Recently, a dataset for Large Vocabulary Instance Segmentation (LVIS) [17] was released as an attempt to increase the coverage of the annotated categories in COCO from 80 to over 1200 categories including the long-tail categories that appear rarely in the dataset.

To overcome the bottleneck in obtaining large annotated datasets, previous works have attempted instance segmentation with weaker forms of supervision such as the object bounding boxes [7, 39], points on instances [2] or image-level labels [49]. Recent works [31, 25, 12, 48] aim to improve the generalization capability of instance segmentation models by employing a partially supervised setting

where only a subset of classes have instance-level mask annotations during training; the remaining classes have only bounding box annotations. However, these methods still require that the bounding boxes of all categories be known prior to training, which limits their ability in discovering novel objects that do not have bounding box annotations.

In this work we propose the first instance segmentation method that performs unsupervised discovery of the long-tail objects through representation learning using hierarchical self-supervision. To the best of our knowledge, this is the first method that eliminates the need for any type of instance level mask or bounding box annotation for the long tail of categories during training. Our idea is that since instance segmentation models (e.g. Mask R-CNN [19]) trained on a small set of categories can already produce good class-agnostic mask proposals, we can leverage these masks and use representation learning to separate these proposals into distinct categories. Therefore, much of our approach, after taking the mask proposals from a region proposal network, is focused on the representation learning. We present an effective approach that exploits the inherent hierarchical visual structure of the objects and enables the embedded features of all proposed masks to be easily differentiated through unsupervised clustering.

Our method is motivated by the recent works [28, 36] that use hyperbolic embeddings to boost the performance of downstream computer vision tasks. These works draw inspiration from the key observation that objects in the real world exhibit hierarchical structure. To perform self-supervised learning in hyperbolic embedding space, we introduce three triplet losses for learning better mask features and capturing hierarchical relations between the masks. We show that our model outperforms the state-of-the-art partially supervised models [25, 31] even though they require box annotations of the long-tail objects during training and we do not. We also qualitatively show that our model is able to discover and segment additional novel object categories. Compared to the the fully supervised setting, our model is competitive in terms of its capability of detecting and segmenting rare and small objects.

We summarize our main contributions as the following:

- We propose an instance segmentation method that is able to discover the long-tail objects using hierarchical self-supervision. To the best of our knowledge, this is the first instance segmentation method that eliminates the need for any type of ground truth annotation of the long-tail categories during training.
- We leverage hyperbolic embeddings to capture the hierarchical structure in the segmented objects and demonstrate the effectiveness of our learned embeddings as compared to their Euclidean counterpart.
- We show that our method outperforms the state-of-the-art partially supervised instance segmentation methods

although using less supervision. We also provide the first set of baseline numbers on LVIS for such self-supervised methods.

2. Related Work

Fully supervised instance segmentation Fully supervised instance segmentation is well studied. State-of-the-art methods [19, 4] use top-down approaches that first generate object bounding boxes using a region proposal network and then perform instance segmentation on each detected bounding box. On the other hand, the bottom-up approaches [8, 13] segment each instance directly without referring to the detection results. These approaches rely on object-level mask annotations during training, which may restrict their applicability to real world problems.

Partially or weakly supervised instance segmentation Prior works [25, 31, 48, 12] assume a partially supervised setting where all the training images have ground truth bounding boxes, but only a subset of the images have segmentation annotations. While showing competitive segmentation results as compared to the fully supervised oracles, these works also showed great generalization ability in class-agnostic learning.

[25] takes a detection-based approach where a learned weight transfer function predicts how each class should be segmented based on parameters learned for detecting bounding boxes. To generalize to novel categories, it uses class-agnostic training which treats all categories as one foreground category. [31] takes a bottom-up approach where it begins with pixel-wise representations of the objects. Assuming that novel objects share similar shapes with the labeled common objects, [31] leverages aggregated shape priors as the intermediate representation to improve the generalization of the model to novel objects. [48] starts with a saliency map for each box detection and generates a shape representation which serves as shape prior to improve segmentation quality. [12] parses shape and appearance commonalities through instance boundary prediction and modeling pairwise affinities among pixels of feature maps to optimize the separability between instance and the background.

Other weakly supervised instance segmentation methods use even less annotation data, such as box-level annotations [27, 23], image-level annotations [49], and image groups [24]. To our knowledge, there are no prior instance segmentation methods that do not consume any kind of annotation data during training.

Unsupervised object discovery Unsupervised object discovery has long been attempted in computer vision but remains a challenging task. Earlier unsupervised methods ex-

plored partial correspondence and clustering of local features [16], clustering by composition [11]. All of these approaches have been successfully demonstrated in a restricted setting with a few distinctive object classes, but their localization results turn out to be far behind weakly supervised results on challenging benchmarks. [6] proposed an unsupervised approach using part-based matching and achieved state-of-the-art localization results on a challenging dataset Pascal VOC [10]. However, [6] requires that the dataset contain multiple dominant object classes and a few noisy images without any target objects, which is not a valid assumption in datasets with complex scenes and semantics such as COCO [34].

Self-supervised learning Self-supervised learning constructs pretext tasks by exploiting training signals directly from the input data without external supervision. Self-supervised learning is proven to be effective in a range of pretext computer vision tasks such as patch position prediction [9, 37], colorization [47, 32, 33], inpainting [40] and representation learning [42]. Previous works also attempted using self-supervision in instance segmentation. [3] utilized self-supervised signals from the temporal consistency of the frames in video data. [41] gains supervision signals through the interaction between an active agent and the segmented pixels. [26] proposed a self-supervised approach for part segmentation, where they use the intuition that parts provide a good intermediate representation of objects and devise several loss functions that aids in predicting part segments that are geometrically concentrated. However, self-supervised representation learning for instance segmentation for in-the-wild images is rarely explored. In this work, we present a novel method for instance segmentation through mask-level representation learning with self-supervised triplet losses.

Representation learning Object representations learned from large datasets can be used as intermediate representations on downstream tasks such as image classification. One popular class of methods for unsupervised representation learning uses generative models. Generative approaches [22, 30, 15] learn to generate or model pixels in the input space. However, pixel-level generation or reconstruction is computationally expensive and may not be necessary for learning intermediate object representations. In addition, pixel-based objectives can lead to such methods being overly focused on pixel-based details, rather than more abstract latent factors. Unlike generative methods that focus on pixel-level generation, contrastive methods [1, 21, 5] learn visual representations by contrasting positive and negative examples in the latent space. In our work, we learn the representations of the segmented foreground objects using a pretrained feature extractor and a hyperbolic embed-

ding space. We sample anchor objects and the corresponding positive and negative examples and use triplet loss to learn the object level representations. In addition, we argue that the representations of the foreground region within a bounding box should be closer to the full region relative to the background region, and thus by using triplet loss in the same embedding space, we are able to augment the mask features.

Hyperbolic embeddings Most existing metric learning methods in computer vision use Euclidean or spherical distances. Natural images often exhibits hierarchical structures [28, 36] which makes hyperbolic embeddings a better choice than their Euclidean counterpart. In this work, we learn a hyperbolic embedding space for the mask level features. Following previous works [14, 28, 38], we embed the features into a Poincaré ball model with negative curvature, which is a special instance of hyperbolic manifolds and is suitable for this particular task because of its simplicity and differentiable distance function.

3. Proposed Method

Our method is composed of three main parts: (i) mask proposal generation, (ii) representation learning with hyperbolic embeddings, and (iii) discovery of novel objects using unsupervised hyperbolic clustering. An overview of our method is shown in Figure 2. We first use a class-agnostic mask proposal generation network (pre-trained on common categories) to generate segmentation mask proposals of possible objects in the image. The goal of pre-training is for the mask proposal network learn to produce mask proposals that represent generic objectness from visual features. Although the mask proposal generation network is pre-trained with ground truth segmentation annotations of common categories, it has not been trained to detect or segment novel, long-tail categories.

Once we have the mask proposals from the segmentation head, we sample mask triplets according to the procedure in Section 3.2. Then, we extract the mask features and embed them into a Poincaré ball using contrastive triplet losses. We use sampling mechanisms to sample the masks such that the hierarchical structure in the proposed masks is fully exploited and encoded in the hyperbolic space. In the end, once we have trained a hyperbolic embedding space, we can perform unsupervised hyperbolic clustering to identify the distinct object categories of the embedded masks.

As compared to previous partially or weakly supervised methods that are only trained to detect and segment categories that exist in the ground truth annotations, our method considers a larger set of categories by keeping all the class-agnostic mask proposals. Therefore, while previous methods such as [31, 25, 12, 48] can only segment a pre-defined set of categories, our method is able to discover novel object

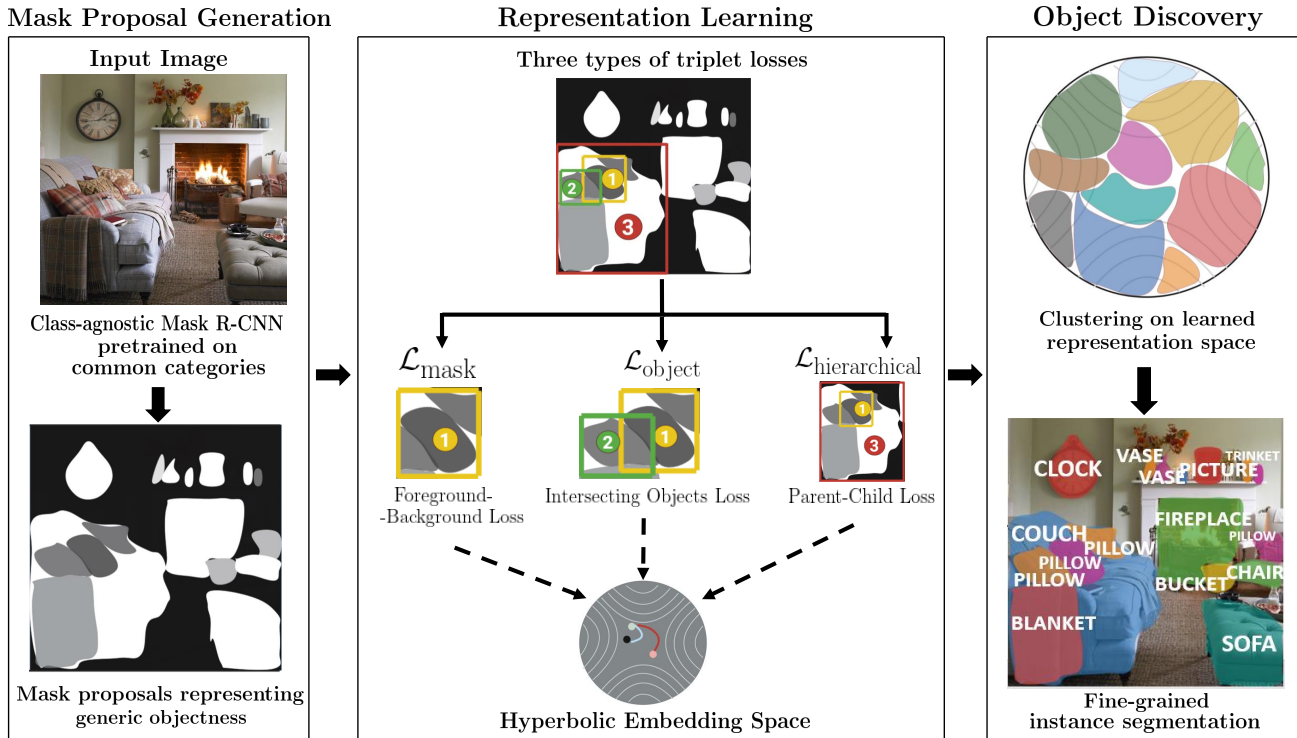


Figure 2. Overview of our method. We first use a class-agnostic mask proposal network (pre-trained on common categories) to generate masks for all possible objects in the image (Section 3.1). We sample the masks using sampling rules that exploits the relationship and hierarchical structure within the proposals. Then, we embed the sampled mask features into a Poincaré ball using self-supervised triplet losses (Section 3.2). Once we have the trained Poincaré embedding space, we perform unsupervised hyperbolic clustering to identify the distinct object categories of the embedded masks (Section 3.3). The output of our method consists of fine-grained instance segmentation masks.

categories, and we do this through effective representation learning in hyperbolic space using self-supervised signals. We will explain each part in detail in Section 3.1, 3.2 and 3.3.

3.1. Mask Proposal Generation

The first step of our method is to generate class-agnostic mask proposals of all possible objects in the input images. Concretely, our mask proposal generation network takes a batch of N training images $\mathcal{I} = \{I_1, I_2, \dots, I_N\}$. Given an image $I \in \mathcal{I}$, we first use a mask proposal network to generate k proposed bounding boxes $\mathcal{B} = \{B^1, B^2, \dots, B^k\}$ and the object segmentation masks within those bounding boxes $\mathcal{M} = \{M^1, M^2, \dots, M^k\}$. The mask proposal network is pre-trained on a large enough dataset with common categories (e.g. the 80 common categories in COCO). We refer to the segmented object within the bounding box B^j as foreground $x^{j,fg}$, and the rest background $x^{j,bg}$. We then learn the relationships between the detected objects through sampling and learning a hyperbolic embedding space using the sampled mask features. This process is illustrated in

Figure 2.

3.2. Representation Learning with Hyperbolic Embeddings

Once we have the foreground $x^{j,fg}$ and background $x^{j,bg}$ for the proposals generated from the previous step, we sample triplets and use a feature extractor f (e.g. ResNet [20]) to get their representations in a hyperbolic space, $f(x^{j,*}) = z^{j,*}$ for object j .

First, inspired by the intuition in [45] that the dominant features of a bounding box detection should come from the visual features of the segmentation within this bounding box whereas the background of the bounding box (i.e. bounding box area that is not part of the segmentation) should contain less meaningful features, we consider a triplet loss that enables us to learn better mask features. Specifically, for each proposed segmentation mask, we encourage the mask features and the bounding box features to be close to each other. In addition, the features of the background content within this bounding box should not be as close to the mask features as the bounding box features. This is particularly important in the context of the

long-tail nature of our task, since the foreground and background may often correspond to distinct objects in the visual hierarchy of the scene. The resulting triplet loss $\mathcal{L}_{\text{mask}}$ is

$$\mathcal{L}_{\text{mask}} = \sum_{j \in [k]} \max(0, \alpha - d(z^{j,\text{full}}, z^{j,\text{bg}}) + d(z^{j,\text{full}}, z^{j,\text{fg}})) \quad (1)$$

where α is the margin of the triplet loss, and d is the distance function in the respective embedding space. Formally, the distance function defined on a Poincaré ball with curvature -1 is defined as

$$d(x, y) = \cosh^{-1} \left(1 + 2 \frac{\|x - y\|^2}{(1 - \|x\|^2)(1 - \|y\|^2)} \right) \quad (2)$$

We use Euclidean distance for the respective ablative experiments in the Euclidean space.

Concurrently, we sample triplets for learning the object representations. For each segmented bounding box, we randomly choose an overlapping proposal that exceeds a predefined IoU threshold (0.4) to be the positive sample, and a non-overlapping proposal to be the negative sample. The negative sample can be from either the same image, or a different image in the training batch. The idea is that if multiple proposals suggested by the region proposal network have high overlapping area, then they are likely to be the same object. Note that although Non-maximum Suppression (NMS) does get rid of high overlapped boxes, it still leaves some remaining overlapping boxes. We notice that such high IoU proposals tend to exist when there are few dominant objects in the image. In such cases, the region proposal network would tend to make proposals of different parts of the same object repeatedly. This triplet loss comes from the intuition that since those proposals are essentially of the same object, their feature embeddings should also be close to each other in the embedding space. Hence, this loss term plays a similar role as the common transformation-based loss in contrastive learning. Formally, the triplet loss $\mathcal{L}_{\text{object}}$ is defined by

$$\mathcal{L}_{\text{object}} = \sum_{j \in [k]} \max(0, \alpha - d(z^{j,\text{fg}}, \hat{z}^{j,\text{fg}}) + d(z^{j,\text{fg}}, \bar{z}^{j,\text{fg}})) \quad (3)$$

where $\hat{z}^{j,\text{fg}}$ is the feature of the positive sample, and $\bar{z}^{j,\text{fg}}$ is the feature of the negative sample.

Last but not the least, we include a hierarchical loss term which is essential in preserving the hierarchical structure among the embedded features. Specifically, for each anchor mask proposal, we find another mask that falls into the area of the anchor mask but is significantly smaller in size. As compared to the previous sampling procedure for $\mathcal{L}_{\text{object}}$, the anchor mask and the positive mask proposal in the hierarchical loss would have a much smaller IoU. We refer to the anchor mask as the “parent” and this sampled mask “child” since the sampled mask is likely to be a sub-component of the anchor object. We use a Poincaré

ball with negative curvature as the embedding space, because a Poincaré ball is able to capture this type of hierarchical parent-child relations within the data with arbitrarily low distortion as compared the Euclidean space. Therefore, we utilize this useful property of Poincaré ball and explicitly enforce that the mask features of the “parent” object to be closer to the origin of the tree than the mask features of the “child” object. Formally, we include a loss term $\mathcal{L}_{\text{hierarchical}}$,

$$\mathcal{L}_{\text{hierarchical}} = \sum_{j \in [k]} \max(0, \alpha - d(z^{j,\text{child}}, o) + d(z^{j,\text{fg}}, o)) \quad (4)$$

where o represents the origin of the Poincaré ball, and $z^{j,\text{child}}$ is the feature of the child mask of proposal j .

The total loss for image I is a weighted sum of three loss terms.

$$\mathcal{L} = \beta \mathcal{L}_{\text{mask}} + \gamma \mathcal{L}_{\text{object}} + \mathcal{L}_{\text{hierarchical}} \quad (5)$$

3.3. Discovery of Novel Objects

We validate our approach using LVIS [17] which contains ground truth annotations for the long-tail categories in the COCO images. In order to evaluate our approach against existing fully supervised or partially supervised benchmarks as well as showing that our instance embeddings have learned diverse object representations, we perform hyperbolic K-means clustering (Figure 2) with K being a much larger number than the number of categories in the annotations. We use a large K in order to account for the large number of long-tail or novel categories. We pick a few ground truth masks from each category and use those embedded ground truths as anchors in assigning labels to our detected clusters.

In our evaluation, we look at the clustered mask features and assign a label to each cluster by taking the label of the closest ground truth category. Once a cluster that is on average closest to a ground truth cluster has been assigned the corresponding label, other clusters are no longer eligible to be assigned to that label. This more rigorously considers our clusters to be good only if they have both high precision as well as recall, since the same label concept should not be split across multiple clusters if our goal is category discovery. With this process, clusters not assigned to any LVIS labels are considered “novel”.

In the COCO to LVIS experiment in Section 4, we show that with a mask proposal network pretrained on COCO categories, our method is effective in discovering long-tail objects that are not among the common categories in COCO.

4. Experiments

Datasets We pre-train our class-agnostic mask proposal network on COCO [34] that contains 80 common categories. We believe the true strength of our method lies in

Model	Supervision	mAP	mAP ₅₀	mAP ₇₅	mAP _r	mAP _c	mAP _f	mAP _s	mAP _m	mAP _l
Mask R-CNN	Fully Supervised	0.201	0.327	0.212	0.072	0.199	0.284	0.106	0.214	0.325
ShapeMask [31]	COCO masks+LVIS boxes	0.084	0.137	0.089	0.056	0.084	0.102	0.062	0.088	0.103
Mask ^X R-CNN [25]	COCO masks+LVIS boxes	0.056	0.095	0.058	0.024	0.051	0.079	0.031	0.056	0.078
Ours (rand. init. backbone)	COCO masks	0.096	0.139	0.104	0.051	0.092	0.168	0.075	0.107	0.139
Ours	COCO masks	0.109	0.160	0.113	0.087	0.105	0.174	0.092	0.129	0.147

Table 1. Quantitative results of the COCO to LVIS generalization experiment. We report mAP for the small/medium/large objects and the rare/common/frequent objects as defined by LVIS. Our ResNet-50 for mask proposals was pretrained on ImageNet [43]; a version with random initialization “Ours (rand. init. backbone)” instead drops performance slightly from 0.109 to 0.096 overall mAP. However, the crippled model still outperforms the prior work baselines, despite the fact that ShapeMask and Mask^X R-CNN do use backbones pretrained on ImageNet in addition to their extra utilization of bounding boxes of the LVIS categories.

its ability to find new objects besides the 80 common object classes. Therefore, we validate our approach using LVIS [17] which contains ground truth annotations for the long-tail categories in the COCO images. LVIS contains ground truth instance-level segmentations for over 1200 categories including sparsely populated categories that appear fewer than 5 times in the training set. All the 80 common categories in COCO can be mapped to LVIS.

Metrics Following previous works we use mean average precision (mAP) for quantitative evaluation. We split the LVIS categories into 2 subsets: the 80 common categories that exist in COCO annotation and the novel (long-tail) categories that are not in COCO. We report mAPs on the second subset (i.e. novel categories). To show that our method generalizes well to long-tail categories, we perform qualitative examination and cluster analysis on the discovered novel categories.

4.1. Implementation Details

Our mask proposal network is built upon the class-agnostic version of Mask R-CNN [19]. Specifically, all the bounding box proposals (after NMS) are passed to the class-agnostic mask branch directly without going through the classification branch. The mask proposal network is pretrained with the 80 common categories in COCO. We use ResNet-50 [20] as the backbone network. Note that the pre-training is only used for the model to learn to propose and segment all objects representing general objectness using visual features. We do not inject any information about the novel objects or the category information of the common objects during the pre-training stage. Since our mask proposal method is class-agnostic in the sense that we keep all the proposals regardless of the classes they might be in, we are not discarding any novel or long-tail categories as previous supervised instance segmentation methods do.

For each image, we keep 50 object proposals with the highest objectness score after using non-maximum suppression (NMS) with threshold 0.75. When calculating the

triplet losses, we sample 1 positive example and 3 negative examples for each anchor according to the procedure described in Section 3.2. Therefore, for each anchor mask, we calculate 3 triplet loss terms. Once we have the triplet samples, we use a randomly initialized ResNet-18 followed by 2 fully-connected layers of size 256 and 64 and finally an exponential map to map the features to the Poincaré ball. We use a 2-dimensional Poincaré ball with curvature -1 , and Adam [29] optimizer with learning rate 10^{-4} . The weight β and γ in the loss function (Equation 5) are both 0.2. We use hyperbolic K-Means with $k=1500$. We will justify our selection of these hyper-parameters in Section 4.3.

4.2. Comparison with Partially Supervised Baselines

We compare our model with ShapeMask [31] and Mask^X R-CNN [25] given the similar focus of segmenting novel objects. Both baselines consider a partially supervised setting, where the mask annotations for only a small subset of the categories and the box labels for all the categories are available during training. In the context of COCO to LVIS experiment, [31] and [25] assume that bounding boxes for all categories in LVIS are available, but masks for only the categories in COCO are available during training. However, we do not assume the availability of the bounding boxes of the LVIS categories. In addition to [31] and [25], we also compare our model with the fully supervised Mask R-CNN with the same backbone architecture. The quantitative results are reported in Table 1. Our method outperforms both [25] and [31] although we use less supervision. On the rare and small object categories, our model is comparable to the fully supervised Mask R-CNN model. The discrepancy of performance across different types of object categories is small compared to the fully supervised model, which can perform much worse on rare objects than on frequent categories.




Figure 3. Qualitative results showing model ablations. **From left to right:** Original image; segmentation masks obtained using only mask loss term; with mask loss and object loss; with all three loss terms. “L” means LVIS categories; w/o “L” means additional novel categories. We show that the hierarchical loss is essential for the model to discern the sub-components of an object (e.g. the slat and the bed).

Model	mAP	mAP ₅₀	mAP ₇₅
No. RP = 20	0.0318	0.0459	0.0332
No. RP = 50	0.1086	0.1597	0.1125
No. RP = 100	0.1083	0.1643	0.1117
$\alpha = 0.1$	0.0806	0.8680	0.0844
$\alpha = 0.2$	0.1086	0.1597	0.1125
$\alpha = 0.5$	0.0987	0.1455	0.0982
$\beta = 0.1$	0.0526	0.6280	0.0544
$\beta = 0.2$	0.1086	0.1597	0.1125
$\beta = 0.5$	0.0614	0.7030	0.0639
w/o $\mathcal{L}_{\text{mask}}$	0.0689	0.0842	0.0707
w/o $\mathcal{L}_{\text{object}}$	0.0374	0.0455	0.0396
w/o $\mathcal{L}_{\text{hierarchical}}$	0.0846	0.1082	0.0921
Euclidean	0.0382	0.0641	0.0433
Poincaré	0.1086	0.1597	0.1125

Table 2. Ablation Studies. We vary the number of region proposals (No. RP), margin of the triplet losses (α), weight of the mask loss term (β), and report the mean APs. In addition, we also report mAPs after taking out individual loss terms, as well as using an Euclidean embedding space instead of hyperbolic space.

4.3. Ablation Studies

Effect of the Number of Region Proposals Besides concerns of computational cost, the number of mask proposals to keep per image could influence the performance of our model. Keeping too few proposals would limit our ability to discover novel objects, and keeping too many proposals would introduce too much noise to the embeddings and worsen our cluster quality. By varying the number of region proposals (Table 2) we observe that increasing the number of proposals per image leads to better segmentation results. However, as the number of proposals gets too big, the improvement becomes marginal.

Effect of the Margin Our representation learning process requires sampling of triplets. However, the number of pos-

sible triplets grows cubically as we increase the batch size. Furthermore, choosing hard negatives for tail objects is already difficult given their scarce number whereas it is much easier to find negative samples for common objects. In order to account for this heterogeneity in the possible hard negative samples, we instead vary the margin parameter (α) of the loss function since the margin parameter intrinsically represents the sampling effect. We observed from Table 2 increasing α indeed improves the performance of our model. However, after around the value of 0.2, increasing the margin seem to have minimal effect.

Effect of the Loss Terms Mask loss term in our loss function serves the important role of distinguishing between background and foreground objects. However, we wanted to be sure that the remaining losses are also important for clustering and that their contribution is no less important. Therefore, we gradually increased the weight of the mask loss term and noted that although a very low mask loss term results in worse background segmentation and hence distorted object embeddings, increasing the mask loss does not lead to a monotonic increase in the mean average precision. Overall, our loss function has three terms, and we show that each term is crucial for the performance of the model by taking out individual loss term at a time (Table 2). In Figure 3 we include qualitative results of the ablated loss function, showing that the three loss terms are essential in learning good representations.

Euclidean vs. Hyperbolic Space We opt for a hyperbolic embedding space because of its capability of embedding hierarchical data with arbitrarily low distortion. To validate our hypothesis of the suitability of hyperbolic space over its Euclidean counterpart, we compare our model’s performance with Euclidean versus Poincaré embedding space. In Table 2, we show that the Poincaré embedding space brings a massive improvement over the Euclidean embed-



Figure 4. Examples of the segmentation masks of the discovered categories. “L” means LVIS categories; without “L” means additional novel categories. Left sub-figure: Examples from “clock” and “book” categories in LVIS, as well as the “toggle switch” category that is not in LVIS, but discovered by our model. Right sub-figure: Examples from “umbrella” and “frisbee” categories in LVIS, as well as the “pivot” category that is not in LVIS, but discovered by our model.

ding space in terms of mean average precision.

Cluster Analysis After embedding mask features into the hyperbolic space, we use hyperbolic K-Means clustering to group the mask proposals. We use elbow method [46] to pick the number of clusters. Elbow method is a heuristic used for determining the optimal number of clusters using the explained variation as a function of the number of clusters. In addition, we use purity scores to measure the quality of the resulting clusters. In the context of cluster analysis, purity score is a measure of the extent to which clusters contain a single class [44] and is a commonly-used criterion of cluster quality. In Table 3, we show purity scores resulting from different number of clusters, including the optimal cluster number determined by elbow method. Table 3 suggests that the objects within the discovered clusters have strong semantic coherence, particularly when the clustering is done using the number of clusters chosen by elbow method.

No. of Clusters	LVIS	Purity _{avg}	Purity _r	Purity _c	Purity _f
k=1300	982	0.534	0.335	0.574	0.615
k*=1462 (Elbow)	1079	0.582	0.387	0.667	0.688
k=1500	1083	0.551	0.346	0.652	0.712
k=2000	1106	0.419	0.218	0.660	0.723

Table 3. Cluster purity analysis with different number of clusters. The “LVIS” column includes the number of discovered clusters determined to correspond to LVIS classes (and used to compute the purity scores).

4.4. Qualitative Examples of New Object Discovery

After obtaining clusters using hyperbolic K-means, we match these clusters with the nearest cluster obtained from ground-truth LVIS annotations. This helped us see how suc-

cessful we were in discovering the long-tail in the LVIS. We then mapped a few embeddings in various clusters to their original segmentation mask to get a sense of whether these objects were correctly grouped. We include qualitative examples of the discovered mask clusters in Figure 4. As shown, we are able to discover LVIS categories such as “clock”, “book” and “frisbee”. Beyond discovering LVIS categories that are not in COCO, we also found additional object classes that do not correspond to any LVIS categories (e.g. “toggle switch” and “pivot” in Figure 4). Since these additional object classes are not annotated in LVIS, they are therefore impossible to be found by any previous fully supervised or weakly/partially supervised methods.

5. Conclusion

In this paper we proposed an instance segmentation method that discovers long-tail objects through self-supervised representation learning. Trained on the common categories in COCO without consuming any annotations on the long-tail categories, we show that our model outperforms partially supervised methods that use box annotations of the long-tail objects. More importantly, through learning representations of more granular objects and the rich relationships between related regions, we demonstrated that our model is able to discover long-tail categories in LVIS and novel categories that are not in COCO nor LVIS.

Acknowledgements This material is based upon work supported by the National Science Foundation under Grant No. 2026498, as well as a seed grant from the Institute for Human-Centered Artificial Intelligence (HAI) at Stanford University.

References

- [1] 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. 3
- [2] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *European conference on computer vision*, pages 549–565. Springer, 2016. 1
- [3] Alex Bewley, Vitor Guizilini, Fabio Ramos, and Ben Uppcroft. Online self-supervised multi-instance segmentation of dynamic objects. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1296–1303. IEEE, 2014. 3
- [4] Liang-Chieh Chen, Alexander Hermans, George Papandreou, Florian Schroff, Peng Wang, and Hartwig Adam. Masklab: Instance segmentation by refining object detection with semantic and direction features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4013–4022, 2018. 1, 2
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 3
- [6] Minsu Cho, Suha Kwak, Cordelia Schmid, and Jean Ponce. Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1201–1210, 2015. 3
- [7] Jifeng Dai, Kaiming He, and Jian Sun. Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1635–1643, 2015. 1
- [8] Bert De Brabandere, Davy Neven, and Luc Van Gool. Semantic instance segmentation with a discriminative loss function. *arXiv preprint arXiv:1708.02551*, 2017. 2
- [9] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. 3
- [10] 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. 1, 3
- [11] Alon Faktor and Michal Irani. “clustering by composition”—unsupervised discovery of image categories. In *European Conference on Computer Vision*, pages 474–487. Springer, 2012. 3
- [12] Qi Fan, Lei Ke, Wenjie Pei, Chi-Keung Tang, and Yu-Wing Tai. Commonality-parsing network across shape and appearance for partially supervised instance segmentation. In *European Conference on Computer Vision*, pages 379–396. Springer, 2020. 1, 2, 3
- [13] Alireza Fathi, Zbigniew Wojna, Vivek Rathod, Peng Wang, Hyun Oh Song, Sergio Guadarrama, and Kevin P Murphy. Semantic instance segmentation via deep metric learning. *arXiv preprint arXiv:1703.10277*, 2017. 2
- [14] Octavian-Eugen Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. *arXiv preprint arXiv:1805.09112*, 2018. 3
- [15] 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*, pages 2672–2680, 2014. 3
- [16] Kristen Grauman and Trevor Darrell. Unsupervised learning of categories from sets of partially matching image features. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 1, pages 19–25. IEEE, 2006. 3
- [17] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5356–5364, 2019. 1, 5, 6
- [18] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Simultaneous detection and segmentation. In *European Conference on Computer Vision*, pages 297–312. Springer, 2014. 1
- [19] 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. 1, 2, 6
- [20] 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. 4, 6
- [21] Olivier J Hénaff, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019. 3
- [22] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006. 3
- [23] Cheng-Chun Hsu, Kuang-Jui Hsu, Chung-Chi Tsai, Yen-Yu Lin, and Yung-Yu Chuang. Weakly supervised instance segmentation using the bounding box tightness prior. In *Advances in Neural Information Processing Systems*, pages 6582–6593, 2019. 2
- [24] Kuang-Jui Hsu, Yen-Yu Lin, and Yung-Yu Chuang. Deepco3: Deep instance co-segmentation by co-peak search and co-saliency detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8846–8855, 2019. 2
- [25] Ronghang Hu, Piotr Dollár, Kaiming He, Trevor Darrell, and Ross Girshick. Learning to segment every thing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4233–4241, 2018. 1, 2, 3, 6
- [26] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 869–878, 2019. 3

- [27] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2017. [2](#)
- [28] Valentin Khrukov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. Hyperbolic image embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6418–6428, 2020. [2](#), [3](#)
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [30] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [3](#)
- [31] Weicheng Kuo, Anelia Angelova, Jitendra Malik, and Tsung-Yi Lin. Shapemask: Learning to segment novel objects by refining shape priors. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9207–9216, 2019. [1](#), [2](#), [3](#), [6](#)
- [32] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *European conference on computer vision*, pages 577–593. Springer, 2016. [3](#)
- [33] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6874–6883, 2017. [3](#)
- [34] 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. [1](#), [3](#), [5](#)
- [35] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8759–8768, 2018. [1](#)
- [36] Emile Mathieu, Charline Le Lan, Chris J Maddison, Ryota Tomioka, and Yee Whye Teh. Continuous hierarchical representations with poincaré variational auto-encoders. *arXiv preprint arXiv:1901.06033*, 2019. [2](#), [3](#)
- [37] T Nathan Mundhenk, Daniel Ho, and Barry Y Chen. Improvements to context based self-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9339–9348, 2018. [3](#)
- [38] Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. *arXiv preprint arXiv:1705.08039*, 2017. [3](#)
- [39] G Papandreou, L-Ch Chen, K Murphy, and AL Yuille. Weakly-and semi-supervised learning of a dcnn for semantic image segmentation. *arXiv preprint arXiv:1502.02734*, 10, 2015. [1](#)
- [40] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016. [3](#)
- [41] Deepak Pathak, Yide Shentu, Dian Chen, Pulkit Agrawal, Trevor Darrell, Sergey Levine, and Jitendra Malik. Learning instance segmentation by interaction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 2042–2045, 2018. [3](#)
- [42] Zhongzheng Ren and Yong Jae Lee. Cross-domain self-supervised multi-task feature learning using synthetic imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 762–771, 2018. [3](#)
- [43] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. [6](#)
- [44] Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge, 2008. [8](#)
- [45] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Mask-guided contrastive attention model for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1179–1188, 2018. [4](#)
- [46] Robert L Thorndike. Who belongs in the family? *Psychometrika*, 18(4):267–276, 1953. [8](#)
- [47] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016. [3](#)
- [48] Yanzhao Zhou, Xin Wang, Jianbin Jiao, Trevor Darrell, and Fisher Yu. Learning saliency propagation for semi-supervised instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10307–10316, 2020. [1](#), [2](#), [3](#)
- [49] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3791–3800, 2018. [1](#), [2](#)