

## Open-World Instance Segmentation: Exploiting Pseudo Ground Truth From Learned Pairwise Affinity

Weiyao Wang<sup>†</sup>, Matt Feiszli<sup>†</sup>, Heng Wang<sup>†</sup>, Jitendra Malik<sup>†§</sup>, Du Tran<sup>†</sup>

<sup>†</sup> Meta AI Research <sup>§</sup> UC Berkeley

{weiyawang,mdf,hengwang,trandu}@fb.com, malik@berkeley.edu

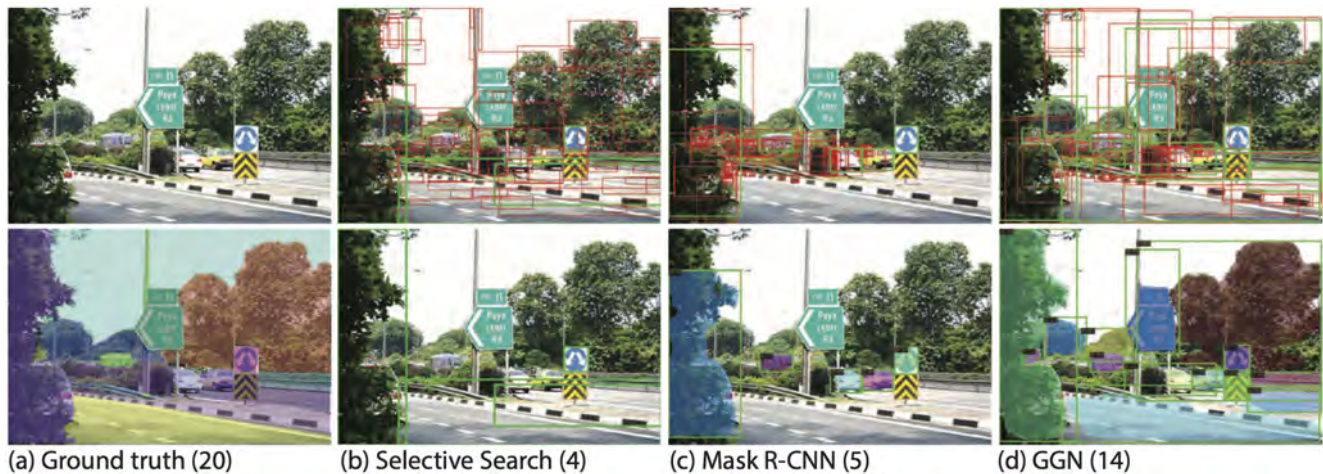


Figure 1. **Comparison of GGN with different baselines.** (a) an input image from ADE20K [61] (upper row) with the ground truth object masks overlaid (lower row). Three different approaches: (b) Selective Search (SS) [51], (c) Mask R-CNN [22], and (d) our Generic Grouping Network (GGN) are applied on the same image to predict the top 100 proposals. The upper images provide all top-100 proposals predicted by three approaches with false and true positive proposals visualized in red and green boxes, respectively. The lower images provide only true positive proposals. The number of true positive proposals or ground truth objects are denoted in parentheses. Among 20 ground truth objects, SS recalls only 4, Mask R-CNN detects 5, and GGN retrieves 14. SS is a bottom-up non-parametric approach, thus has no notion of objectness. Mask R-CNN can make whole object proposals; however it still fails to detect objects that are not seen during training. Our GGN can predict whole object proposals and generalize to unseen categories.

### Abstract

*Open-world instance segmentation is the task of grouping pixels into object instances without any pre-determined taxonomy. This is challenging, as state-of-the-art methods rely on explicit class semantics obtained from large labeled datasets, and out-of-domain evaluation performance drops significantly. Here we propose a novel approach for mask proposals, Generic Grouping Networks (GGNs), constructed without semantic supervision. Our approach combines a local measure of pixel affinity with instance-level mask supervision, producing a training regimen designed to make the model as generic as the data diversity allows. We introduce a method for predicting Pairwise Affinities (PA), a learned local relationship between pairs of pixels. PA generalizes very well to unseen categories. From PA*

*we construct a large set of pseudo-ground-truth instance masks; combined with human-annotated instance masks we train GGNs and significantly outperform the SOTA on open-world instance segmentation on various benchmarks including COCO, LVIS, ADE20K, and UVO.*

### 1. Introduction

Instance segmentation is the task of grouping pixels into object instances [22]. In the closed-world setup, the task is to detect and segment objects from a predefined taxonomy. In contrast, the open-world setting requires segmenting objects of arbitrary categories. For a model trained in a closed-world setup, this means segmenting not only the “seen” categories (those presented at training time) but also

the “unseen” categories (not seen during training) [27, 54].

There is generally a large performance gap between the seen and unseen domains. Leading computer vision systems today have tightly coupled recognition and segmentation; these systems are unable to segment out objects that they cannot recognize (e.g. Fig 1 (c)). Comparing Average Recall (AR@100) of Mask R-CNN [22] trained on 80 COCO [33] classes vs a subset of 20 classes, AR@100 of 60 classes out of training taxonomy drops from 49.6% to 19.9% when no mask of these classes is provided in training data. The “unseen” gap remains large if we train on larger taxonomy (e.g., 1,000+ classes in training data) Table 5). In contrast, humans can readily group and segment objects which they cannot categorize - few of us can identify the 6500 Passerine bird species, but we can readily segment out a perching bird from a tree branch. Or use another often-quoted example: our familiarity with a generic quadruped body plan enables us to segment out horses, donkeys and zebras, and even an okapi when first encountered.

On the other hand, models which were common in computer vision in 2000-2015 (e.g., [1, 4, 6, 18, 51, 58]), before deep learning for supervised object detection took off, were quite category-agnostic. They didn’t work as well as, say, Mask R-CNN on a category for which it has trained, but they worked across the board (e.g. Figure 1 (b)). The goal was to come up with a moderately sized set of object proposals which included the true objects. The emphasis was on recall; precision was secondary. MCG [4] is an illustrative example. It starts with local grouping which produces a set of elementary regions of coherent color and texture, “super-pixels”. These typically over-segment objects; e.g. a person might be broken up into a face, a torso, legs, parts of clothing, shadows, etc. MCG then assembles regions into objects by considering various groupings of regions, and ranks them on some “objectness” score. While some learning is involved in both edge detection and objectness ranking, the method works primarily with hand-crafted features and a small number of parameters, quite unlike the deep learning zeitgeist.

How do we get the best of both worlds? A modern instance segmentation system (eg. Mask R-CNN) would do well if given comprehensive training data containing a large number of examples from all visual categories. While we have a practically infinite supply of raw natural images, obtaining mask annotation is very expensive. Multiple approaches have emerged to handle this data problem. Self-supervised learning [10, 11, 20, 42] is the most well-known; self-learning [46, 47, 59, 62] is another approach, based on the classical idea of adding high-confidence guessed labels to previously unlabeled data, and then combining this “pseudo ground truth” data with real ground truth. We exploit this second strategy.

Our approach begins with a learned pairwise affinity pre-

dictor (Figure 2a), followed by a module which extracts and ranks segments (Figure 2b, essentially a very simplified version of MCG [4]). We can run this on any image dataset without using annotations; we extract the highest ranked segments as “pseudo ground truth” candidate objects. This is a large and category-agnostic set; we add it to our (much smaller) datasets of curated annotations, to train a Mask R-CNN instance segmentation module. Ideally this model should become more generic and class-agnostic (Figure 2c). Indeed, this simple approach produces impressive gains compared to closed-world training on the same backbone (Mask R-CNN) (Table 5, Table 6, Table 7): **+11%** on VOC to Non-VOC cross-category evaluation, **+3.9%** on COCO to LVIS cross-category evaluation, **+5.8%** on COCO to ADE20K and **+5.2%** on COCO to UVO.

Our contributions in this paper include:

- A novel approach, **Generic Grouping Networks (GGNs)**, for open-world instance segmentation; GGN exploits additional pseudo ground truth supervision generated from learned pixel-level pairwise affinities.
- Comprehensive ablation experiments which provide insights about GGNs and the problem of open-world instance segmentation.
- GGNs achieve state-of-the-art performance in open-world instance segmentation on various benchmarks including COCO, LVIS, ADE20K, and UVO.

## 2. Related works

**Object and instance segmentation.** Before the success of deep learning, object segmentation approaches typically worked by grouping local regions into whole objects. Popular approaches include graph-based [15, 18], Normalized Cut [24], Graph Cut [7], Multiscale Combinatorial Grouping [4], and Selective Search [51]. Since deep learning, end-to-end approaches proved their success on problems such as semantic segmentation [37, 39], instance segmentation [22], panoptic segmentation [29, 52]. Despite sharing the common problem of segmentation, our approach is different in the open-world setup: instead of assuming a closed-world taxonomy, our work aims at detecting and segmenting both seen and unseen objects.

**Pairwise affinity based approaches.** Pairwise affinity is used in most graph-based segmentation methods [7, 23, 24] as an important term defining a relation graph of pixels for segmentation. The pixel-level pairwise affinity can be either hand-constructed [7, 23, 24] or learned [16, 17, 28, 34, 35, 40, 50]. Similar to pairwise affinity, object boundary detection [41, 57] is a dual problem but offers weaker supervision (sec 5) and cannot be trained as-is on non-exhaustive annotations. Different from previous approaches, instead of directly using the learned pairwise affinity for segmentation, we use it as an intermediate representation for pseudo-ground-truth generation which is later used to train our

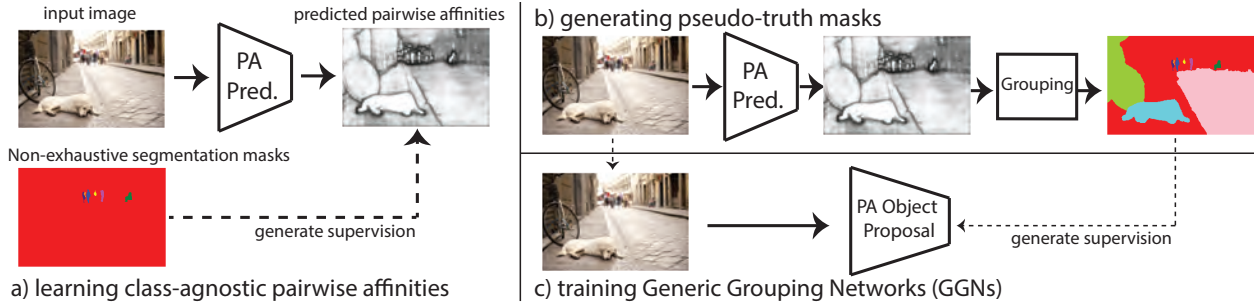


Figure 2. **Overview of our approach.** (a) First, a Pairwise-Affinity Predictor (*PA Pred.*) is trained to predict pairwise affinities using non-exhaustive segmentation masks as supervision. (b) Once trained, *PA Pred.* is used to predict pairwise affinities of the images. A grouping module is then applied on the predicted pairwise affinity maps to generate pseudo-ground-truth masks. (c) A class-agnostic generic object proposal network (e.g., class-agnostic Mask R-CNN) is trained end-to-end using a combination of GT and generated pseudo-GT masks.

generic grouping model. Another difference between our approach and previous learned pairwise affinity comes from the open-world setting of the problem.

**Open-world benchmarks and approaches.** Open-world setup [5, 43] has been recently (re)-introduced into various problems in computer vision such as recognition [30, 38], tracking [36], detection [25, 27, 53], and segmentation [54]. Among these, our work is mostly related to UVO [54] and OLN [27]. We share the same problem of interest of open-world instance segmentation with UVO [54]. However, [54] provides a new benchmark for the problem while our work provides an approach. Compared with the concurrent OLN work [27], which uses an objectness-based loss for generalization to unseen classes, our approach addresses generalization by combining pixel-level pairwise affinity with local grouping. Our work and OLN are orthogonal and complementary; as shown later, our approach alone is on par with OLN, and produces 4.5 – 5.7% improvements when combined with OLN (see Table 6 in section 5).

### 3. Learning pairwise pixel affinities

Grouping can be locally represented by pixel pairwise relationship: whether two neighboring pixels should be grouped together or not. Given a 3-channel RGB input image  $I \in \mathbb{R}^{3 \times H \times W}$ , we consider a pixel’s pairwise relationship in a  $3 \times 3$  neighborhood. This gives a pairwise affinities map of  $P \in \{0, 1\}^{8 \times H \times W}$  and  $P_{i,j} \in \{0, 1\}^8$  encodes the local pixel-level pairwise affinities of the pixel  $(i, j)$  with its 8-neighboring pixels in the image  $I$ . Figure 4 (c) illustrates the pairwise affinity encodings of the two pixels at the centers of the two image patches marked with the pink and yellow squares in Figure 4 (a) and (b). We use pixel-level-prediction convolutional neural networks to predict  $P$  (*PA Pred.*, Figure 2 a), such as FCN [39] and UPerNet [56]. We remark that this is a dual problem to binary object boundary detection ([41, 57]), and techniques used for training *PA Pred.* can also be adopted to binary object boundary detectors to serve as the local representation for our framework.



Figure 3. **Visualization of predicted pairwise affinities and generated pseudo masks, trained only on Person class in COCO.** Despite only seeing masks of Person during training, the PA predictor correctly captures pairwise relationships of other types of objects (top-row). By grouping pixels based on predicted PA, we can generate pseudo masks of other categories (bottom-row).

**Training from non-exhaustive segmentation masks.** Ideally, if all pixels in the image are exhaustively annotated with instance segmentation masks, e.g., all object boundaries are labelled, then all pairs of neighboring pixels can provide good supervision signal for learning pairwise affinity. However, exhaustive annotations for instance segmentation are expensive and time-consuming to obtain, so most datasets come with non-exhaustive segmentation masks (eg. COCO [33]). In particular, non-annotated out-of-taxonomy objects cannot be distinguished from background pixels. To address this, we only use neighboring pixels with object-object or object-background relations for training pairwise affinities; we ignore the unreliable background-background pairs. In addition, training of pairwise affinities is unbalanced: only pixels at an object’s boundary has zero-valued affinities; all other pixels have affinity 1 to all their neighbors. We weight the positive affinities by computing the ratio between the positive affinities and negative affinities on a subset of training data (e.g., 0.05 for positive).

**What does a Pairwise Affinity predictor learn?** Intuitively, an ideal PA measure should discriminate between instance boundaries and instance interiors: i.e. whether two neighboring pixels cross an object boundary or not. Ide-

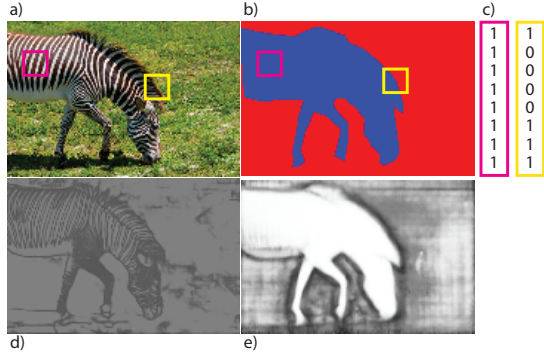


Figure 4. **Pairwise affinity encoding and prediction.** Visualization of an example input image (a) and its corresponding ground truth mask annotation (b) and two pairwise-affinity encoding vectors (c) of the two center pixels of the two image patches marked with pink and yellow squares. The pink-patch’s center pixel belongs to the same instance with all of its 8-neighbor pixels, thus it is encoded with a vector of all ones. The yellow-patch’s center pixel lies at the object boundary and has 4 neighbor pixels belonging to background, thus it is encoded with a binary vector with four 0s and four 1s. (d) Edge prediction of the image in (a) using off-the-shelf edge detector [49]. (e) Pairwise affinity prediction of the image in (a) using our Pairwise Affinity predictor. Our predictor is trained using only person category masks from COCO. Best viewed in color.

ally, our PA predictor should be robust to boundaries of all objects, not just the categories seen during training; this is a key requirement for success in the open-world setting. Indeed, this is the case: our PA predictor learns instance boundaries and generalizes well to unseen classes. Figure 3 top-row and Figure 4 (e) visualize pairwise affinity predictions from our PA predictor, trained using ground truth masks of **only** the person category from COCO. Our PA predictor learns to generalize to unseen classes such as zebra, bird, temple, and cooking pan. We note that ‘person’ masks are particularly diverse owing to clothing and accessories. This is helpful for models to learn and generalize the notion of semantic boundaries. We show quantitative results on PA’s generalization in section 5.2 and exploit this generalization behavior to improve segmentation in section 4. Finally, we point out that pairwise affinity should capture the semantics of instance boundaries; this is quite different from visual edge maps because many visual edges are not instance boundaries. Figure 4 (d) presents the edge prediction by an off-the-shelf edge detector [49]. Many visual edges on the back of the zebra are clearly not object boundaries, but are still detected by an edge detector.

#### 4. Augmenting with pseudo ground truth

Existing state-of-the-art detectors and instance segmentation models, such as Mask R-CNN, often fail to detect and segment novel objects unseen during training. This can be caused by difficulties to group pixels into unknown entities

due to lack of supervision signals during training. In addition, even a novel region is grouped and proposed, a generic concept of objectness is missing and such out-taxonomy proposals are suppressed. We kill two birds with one stone, by using pseudo-GT masks generated from PA to train these detectors. The pseudo-masks benefit from pixel diversity to provide novel segments not seen during training, and therefore enhance supervision signals for both novel grouping and a more inclusive concept of objectness.

**Grouping pixels into regions.** Based on predicted pairwise affinities, we leverage class-agnostic local grouping algorithms to group pixels into instances. One may use the *Connected Component* (CC) algorithm for grouping. CC treats all affinities independently using a hard cut-off threshold to decide pixel connections which may be a sensitive parameter to tune. Alternatively one may use *graph-based hierarchical* grouping (GBH) [15] which is a variant of agglomerative clustering. In segmentation literature [3, 60], Oriented Watershed Transform (OWT), globalized contour through Normalized Cut (gPb) [48] and Ultrametric Contour Map (UCM) [2] are also used for grouping from the edge map of an image. Following [2, 4], we first aggregate the image pairwise affinity map into a semantic edge map using pooling along the channel dimension, e.g., reducing from 8-channels to 1-channel. This semantic edge map is passed to OWT to generate initial segments, whose edges are then globalized through normalized cut. We take the average of semantic edge map and its globalized version as input to UCM for grouping. We acknowledge that a different linear combination of these two might work better upon further study. We provide the ablation comparing CC, GBH, and different components of OWT+gPb+UCM in section 5.

**Computing objectness.** Objectness [13] measures grouping qualities; in our framework, it is critical to decide which pseudo-GT masks to select for training detectors. An ideal objectness score should reveal over-segmentation and under-segmentation. In previous literature, objectness can be modeled by low-level features such as shape and contours such as MCG [4] or learned directly from annotated data as classification (Region Proposal Network [45]) or regression (Object Localization Network [27]). We consider both types of objectness. For low-level features, we use predicted pairwise affinities to define objectness score of each region  $R$  by total affinities  $\mathcal{O}_{PA}(R)$ :

$$\mathcal{O}_{PA}(R) = \frac{Inner(R)}{R_{inner}} - \frac{Outer(R)}{R_{boundary}} \quad (1)$$

where  $Inner(R)$  and  $Outer(R)$  are the inner and outer affinities of  $R$  defined by the sum of pairwise affinities of pixels inside or crossing-boundary of  $R$ , respectively.  $R_{inner}$  and  $R_{boundary}$  denote the number of pixels inside  $R$  and on the boundary of  $R$ . Intuitively, we want to rank high for the region with strong inner pairwise affinities and weak

affinities at the boundary (a.k.a strong cut). For learned objectness, we consider scoring from OLN [27]  $\mathcal{O}_{OLN}(R)$  :

$$\mathcal{O}_{OLN}(R) = \sqrt{\text{centerness}(R) * \text{IoUness}(R)} \quad (2)$$

where  $\text{centerness}(R)$  and  $\text{IoUness}(R)$  are the centerness and IoU predictions of the bounding box of  $R$ . We can optionally combine  $\mathcal{O}_{PA}$  and  $\mathcal{O}_{OLN}$  by taking average.

**Generic Grouping Networks (GGNs).** We generate class-agnostic masks from PA predictor and grouping module and use the objectness score to rank regions provided by grouping methods (Figure 2 b). We then select top ranked regions from each image as pseudo ground-truth (GT) masks for training our generic object proposal network (Figure 2 c). Since the whole approach to generate pseudo-GT masks are designed in a class-agnostic grouping fashion, we expect the pseudo-GT masks to cover a diverse set of objects and parts, and more importantly most of them are from unseen categories, as shown in Figure 3 bottom-row. Since our GGN is trained on a large and diverse set of masks, it is expected to generalize to unseen classes, thus providing a good solution for open-world instance segmentation. GGN is generic in the sense of both pixels and models: it can work on different domains of images, labeled or unlabeled, and it can work on any architecture for object detection or segmentation, such as Faster R-CNN [45], Mask R-CNN [22], YOLO [44], or Swin Transformer [37]. The adoption is as simple as making the multi-class classification prediction head into a binary foreground vs. background classification head.

## 5. Experiments

### 5.1. Implementation Details

**Datasets.** We conduct experiments on COCO17 [33], LVIS [19], ADE20K [61], and UVO [54]. **COCO** is a standard benchmark for instance segmentation with 80 object categories annotated on 164k images. **LVIS** is an instance segmentation dataset with 1203 classes in a long-tail distribution. It is labeled as a federated dataset and does not include exhaustive label for its categories. We adopt LVIS to study cross-category generalization when a large taxonomy is provided. **ADE20k** is a semantic segmentation dataset with all pixels exhaustively annotated by object instances or stuff. **UVO** is a video instance segmentation dataset of YouTube videos (Kinetics400 [26]) with object masks exhaustively labeled. We use validation set of ADE20K (2000 images) and UVO sparse (7356 frames) to evaluate open-world segmentation in the wild in section 5.4. In all setups, we use only mask annotation (without class labels) for training and evaluation in the open-world, class-agnostic. We note that PA predictor, baseline methods (e.g., Mask R-CNN), and GGN has access to the **same** labeled masks.

**Backbone architectures and loss function.** We adopt UperNet [56] for our PA predictor to learn pairwise affinities. For training, our generic grouping networks, we use Mask R-CNN [22] with a ResNet-50 backbone as a default setup. Unless specified otherwise, models are initialized by ImageNet [12] pre-training. We use Binary Cross Entropy loss to train pairwise affinities. We ignore background-background affinity as in section 3. We note that back-propagating losses that include background-background affinities leads to very poor cross-category generalization (e.g., -15% Average Recall).

**Ranking and selecting pseudo-GT masks.** Unless otherwise specified, we use  $\mathcal{O}_{PA}$  (Eq. 1) to rank pseudo-GT masks. We pick top-k pseudo-GT masks per image ( $k \in [1, 3]$ ), where k is selected to improve unseen categories performance while minimally impact seen performance.

**Training and evaluation.** We build model training and inference on MMDet [9] platform; all training are done with the default 1x schedule. Following previous object proposal literature [43,51], we use **average recall** (AR) over multiple IoU thresholds (0.5:0.95) to evaluate model performance.

**Cross-category evaluation.** Cross-category generalization is a major challenge for open-world: how do we detect and segment objects whose categories are outside training data. We split existing datasets by their categories to construct controlled environment for ablations (Table 1). In each setup, we train PA and baseline methods with the same splits of categories (no additional supervision for PA).

On COCO dataset, we follow common practice [27,43] to split COCO into 20 classes overlapped with Pascal VOC [14] for training (seen) and use the rest of 60 COCO-exclusive classes for evaluation (unseen). We further include an extreme case by using only *person* class for training and the rest 79 classes for evaluation.

On LVIS dataset, some categories are highly overlapped: for example, clothes (“jacket”, “wet suit”) are highly overlapped with “person” when the person is wearing the clothes. In a class-agnostic setup, a detector trained with person masks can detect clothes as person and vice versa. Other examples of high-overlapped category pairs are “ball” with “tennis ball”, “alcohol” with “beer bottle”, or “computer monitor” with “television set”. COCO and LVIS share the same set of images, but with different annotations. LVIS covers 1203 categories which include all 80 categories from COCO. COCO also exhaustively annotates all masks of objects that belong to its 80 categories while LVIS is annotated so as to maintain a similar number of masks across categories. This means that some object instances, even if they belong to LVIS categories, are not annotated. As COCO masks are more exhaustively annotated, we use COCO masks and validate cross-category overlap with LVIS masks. We find that there are 67k LVIS masks outside COCO taxonomy having  $>0.5$  IoU with COCO

Dataset	Train	Eval	Image	Mask
COCO	Person(1)	non-Person	64k	161k
	VOC(20)	non-VOC	95k	493k
LVIS	COCO(80)	non-COCO	100k	455k
	non-COCO(1122)	COCO	85k	749k
	+Person(1123)	non-Person	86k	775k

Table 1. **Cross-category generalization evaluation setups.** We split categories in COCO and LVIS to evaluate.

Grouping	CC	GBH	WT+UCM	+OWT	+gPb
Recall@all	14.4	17.1	23.6	23.8	<b>24.2</b>

Table 2. **Comparing different grouping methods.** Methods are applied on the same affinity maps and output roughly a similar number of proposals. OWT+gPb+UCM gives the best recall.

Aggregate	Min	Max	Mean
8-channel	22.8/ <b>19.3</b>	18.1/16.7	22.1/18.9
1-channel	19.9/18.4	NA/NA	<b>23.1</b> /18.5

Table 3. **The effect of different PA aggregation.** Evaluated by AR@100 on both seen and unseen categories (VOC/non-VOC) and results are separated by /. 1-channel prediction is not applicable with max-pooling since the prediction target is all 1s (all pixels have at least one neighbor connected). 8-channel PA prediction with min pooling provides the best AR.

masks. To ensure a clear distinction between seen and unseen categories, we remove those masks in both training and validation for cross-category generalization evaluation. We study transfer performance by training on COCO categories and evaluate on LVIS non-COCO categories and vice-versa.

## 5.2. Learning Pairwise Affinities: Ablation Study

**Grouping mechanisms.** We revisit different grouping methods to construct segment masks from pairwise affinities: Connected Component (CC), Graph-Based Hierarchical [18] (GBH) and methods based on Ultrametric Contour Map [2] (UCM). In UCM, we ablate the effect of orientation in watershed transform (OWT vs. WT), and the effect of including globalized edge (gPb). To evaluate, we generate mask outputs from each of the method and directly evaluate their AR. We tune parameters for each method so that each has roughly the same number of output segments on average. Since CC gives a single non-overlapping output instead of a hierarchical structure like GBH or UCM, we use multiple thresholds and use all segments from each threshold. We found that UCM-based methods significantly outperforms other two methods (Table 2): whereas CC and GBH make decision on a merge based on a single pairwise relationship, UCM uses all relations between two segments, and is more robust. In addition, adding orientation and gPb further improves grouping results.

**PA aggregation.** In UCM, PA (8 neighbors) need to be aggregated into one value to feed to WT. The aggregation can

be implemented by a pooling operation and can be applied before or after PA prediction. Specifically, we can either: (i) train a PA predictor to predict a 8-channel PA map then apply aggregation on the prediction output of the PA predictor; or (ii) train a PA predictor to predict a 1-channel PA map which is the aggregated version of ground truth. We compare different methods for aggregating pairwise affinity values (Table 4) and found min aggregation works the best. Alternatively, we can directly predict the aggregated pairwise affinities. We found that a single-channel prediction of mean of pairwise affinities works comparably with 8-neighbor predictions (Table 4). We remark that “1-channel, min” is equivalent to adopting a binary boundary detector trainer in our framework (e.g., HED [57]), which offers weaker supervision signals than PA.

## 5.3. Cross-category evaluation of GGN

We use the pseudo-GT masks from PA+Grouping to train detectors for open-world segmentation. We optionally use additional ground truth masks when they are available. Since PA generalize well in the open-world, the pseudo masks offer more diversities to the training data and therefore improve generalization of downstream detectors (GGNs). We begin by comparing PA with other candidate representations for open-world segmentation.

**Pairwise affinity is a strong representation for open-world.** Besides pairwise affinities, we consider a few other types of mid-level representations to encode grouping and generalize in the open-world:

- **Edge maps** are strong alternative to PA to encode grouping. We take SOTA edge detector DexiNed [49].
- **Depth maps** by pretrained Mannequin network [32].
- **Feature affinities** computed on semantic features, self-supervised trained on ImageNet (MoCoV2 [11]).

Most of the features here, except edge map, are not proper to run UCM to construct grouping. Therefore, we consider replacing the RGB input with the proposed representation (e.g., depth map or PA) to understand how well the representation can generalize in the open-world compared to RGB in cross-category evaluation.

Surprisingly, all representations have regularizing effects compared to RGB to improve generalization to unseen classes when only training on person (Table 4). Pairwise affinities outperform all other types of representations regardless of application methodologies (replacing input or adding pseudo masks). In particular, using UCM to generate pseudo-GT masks for edge does not benefit much, since without semantics, edge map can over-segment entities.

**GGN significantly outperforms baselines on cross-category generalization.** We take top-scoring pseudo-GT masks and use them in addition with the ground truth masks; we remove pseudo-GT masks whose IoU overlap with in-taxonomy GT masks are greater than 0.5.

Method	replace RGB					UCM mask	
	RGB	depth	edge	MoCo	PA	edge	PA
nonPerson	4.9	10.9	10.5	10.7	<u>14.1</u>	7.9	<b>20.9</b>
nonVOC	19.9	17.8	21.3	21.8	<u>26.5</u>	19.7	<b>28.7</b>

Table 4. **Compare Pairwise Affinity with other types of mid-level representations.** All mid-level representations can serve to help generalizing to unseen categories in certain scenario to certain degrees. The significant improvement of pairwise affinities over edge map shows the importance for boundaries to contain semantics. Methods are evaluated by AR@100. Pairwise Affinities provide strongest generalization signals for open-world grouping.

Train (# classes)	Mask R-CNN	PA+ Grouping	GGN	Upper Bound
COCO-Dataset				
Person (1)	4.9	14.6	<b>20.9</b>	49.2
VOC (20)	19.9	22.0	<b>28.7</b>	49.6
LVIS-Dataset				
COCO (80)	16.5	17.1	<b>20.4</b>	36.1
non-COCO (1123)	21.7	16.2	<b>23.6</b>	35.1
+Person (1124)	27.3	18.4	<b>29.1</b>	44.2

Table 5. **GGN generalizes to out-taxonomy categories significantly better than baselines.** GGN also outperforms the pseudo-GT masks generated by pairwise affinities (denoted as PA+Grouping), proving the benefit of the instance-level training with additional pseudo-GT supervision. Upper bound indicates AR@100 achieved by training on entire taxonomy (all classes).

GGN has significantly stronger cross-category generalization compared to baseline Mask R-CNN (Table 5). On low to medium-sized taxonomy on COCO dataset, GGN achieves **+16%** AR@100 gain and **+8.8%** AR@100 gain when training on Person-only and training on 20 VOC classes, respectively. In the large-taxonomy setup, GGN achieves 1.8% to 3.9% gain on AR@100 in different setups. The gain is smaller when training on non-COCO categories; we believe that this is caused by the fine-grained taxonomy of LVIS: many classes in LVIS are objects parts or parts of other classes. Training on non-COCO categories on LVIS makes pairwise affinities closer to edge maps.

Additionally, we evaluate pseudo-GT masks generated by pairwise affinities with UCM. This is equivalent to our GGN but without using top-down instance-level training (as in Figure 2 b), denoted as *PA+Grouping*. Comparing with Mask R-CNN baseline, local grouping using learned pairwise affinities offer stronger performance in low to medium sized taxonomy. Finally, GGN significantly outperforms the *PA+Grouping* baseline which suggests the benefits of instance-level end-to-end training on pseudo-GT masks.

**GGN is comparable and complementary to state-of-the-arts object proposal method.** Object Localization Network (OLN) [27] is a concurrent work to tackle open-world object proposal. OLN proposes to replace classifi-

Backbone	Base	OLN	GGN	GGN+ $\mathcal{O}_{OLN}$	GGN+ OLN
Faster R-CNN	24.9	33.0	31.5	34.7	<b>37.2</b>
Mask R-CNN	19.9	26.9	28.7	30.9	<b>33.7</b>

Table 6. **GGN is competitive yet complementary to OLN.** Trained on VOC and evaluated on non-VOC using AR@100. Adopting  $\mathcal{O}_{OLN}$  improves ranking of pseudo-masks and thus improves GGN; adopting OLN backbone further improves.

Method	Ranking	ADE20K		UVO	
		AR	AP	AR	AP
Selective Search		3.8	-	4.7	-
Mask R-CNN		14.7	6.4	40.1	18.5
GGN	$\mathcal{O}_{PA}$	18.3	7.9	42.6	19.4
	$\mathcal{O}_{PA} + \mathcal{O}_{OLN}$	<b>21.0</b>	<b>9.7</b>	43.4	<b>20.3</b>
GGN, pseudo-GT pre-training		<b>21.5</b>	<b>9.3</b>	<b>45.3</b>	<b>21.0</b>

Table 7. **Open-world segmentation in the wild on ADE20K and UVO.** GGN significantly outperforms the baseline Mask R-CNN when using the same amount of training data and annotation. In addition, having stronger objectness by combining  $\mathcal{O}_{PA}$  and  $\mathcal{O}_{OLN}$  further improves model performance. Finally, replacing ImageNet label pre-training with ImageNet pseudo-GT masks pre-training (sec 5.5) offers additional improvement.

cation with localization quality prediction to avoid overfitting to annotated objects, which is similar to how we train pairwise affinities by not backpropagating the loss in un-annotated relationships. Different from OLN, pseudo masks generated bring more diversity to training data, and therefore help to generalize better. We compare GGN with OLN in Table 6 and find that GGN achieves similar performance as OLN ( $-1.5\%$  on box AR@100,  $+1.8\%$  on mask AR@100). When adding  $\mathcal{O}_{OLN}$  (Eq. 2) to rank and select pseudo-GT masks, GGN improves by 2.2%. When adopting OLN as backbone, GGN sets new state-of-the-arts for cross-category generalization of VOC to COCO.

#### 5.4. Evaluate open-world segmentation in the wild

Ablations in section 5.3 focus on cross-category generalization in a controlled version of open-world. A more practical question is: how well detectors can generalize across datasets in the wild? It is difficult to evaluate since common datasets, e.g., COCO and LVIS, are only *partially-annotated*. Evaluating open-world segmentation on such datasets may fail to capture performance difference across methods due to punishing precision and not rewarding recall [8]. To handle this, we adopt ADE20K [61] and UVO [54] for evaluating generic proposals in the wild open-world. Specifically, we treat each segmentation mask in ADE20K or UVO as a ground-truth semantic entity and evaluate Average Recall (AR) and Average Precision (AP). This setup evaluates both in-taxonomy and out-taxonomy segments. While UVO contains only objects, ADE20K also includes stuff masks. We emphasize that this is truly **in the**

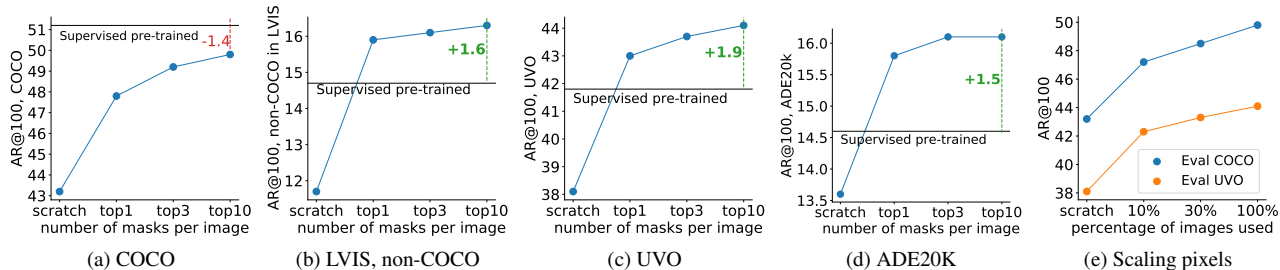


Figure 5. **GGN outperforms ImageNet supervised pre-training for open-world instance segmentation and demonstrates promising scaling behaviors.** We compare pseudo-GT mask pre-training by GGN with ImageNet label pre-training on closed-world (COCO, a), cross-category (non-COCO in LVIS, b) and open-world (UVO, ADE20K, c,d). Except closed-world setting, pseudo-GT masks provide stronger pre-training signals than ImageNet annotated labels (b-d). In addition, performance improves when more pseudo-GT masks selected per image (a-d) or more pixels (unlabeled images) (e) are used.

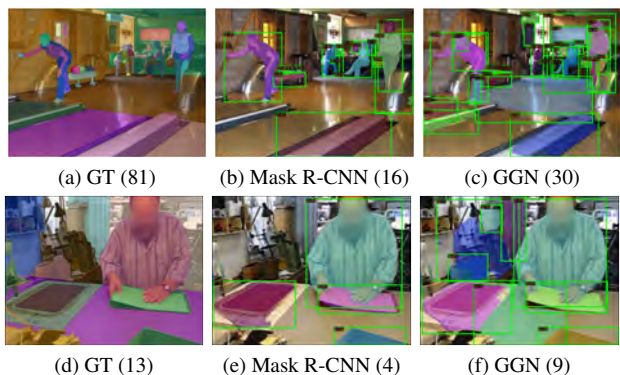


Figure 6. **Visualization of predictions of GGN and Mask R-CNN on ADE20K and UVO.** GGN retrieves more instances correctly (numbers denoted in parentheses) and covers a more diverse set of object categories.

**wild test** as no fine-tuning is done on ADE20K or UVO.

We compare GGN with Selective Search [51] and Mask R-CNN baselines trained with the GT masks of all 80 COCO classes from COCO dataset (Table 7). GGN (enhanced by pseudo masks) significantly outperforms the baseline on both ADE20K and UVO dataset, both AR and AP. In addition, better ranking by combining  $\mathcal{O}_{PA}$  and  $\mathcal{O}_{OLN}$  further improves model performance. Qualitative results comparing Mask R-CNN and GGN on UVO and ADE20K are showed in Figure 6 (more in supplementary).

### 5.5. Pre-training on unlabeled images with GGN

Since PA-based bottom-up grouping can generate masks for any unlabeled image, we hypothesize that training with pseudo-GT masks from additional pixels may help open-world segmentation. Masks from PA generalizes well to new categories on new pixels, and thus benefit from pixel diversity. We study the effect of training GGN on pseudo-GT masks from unlabeled images from ImageNet [12].

Specifically, we use PA trained on 80 COCO categories from COCO with random initialization. We generate pseudo-GT masks on ImageNet images and pre-train a ran-

domly initialized Mask R-CNN on pseudo-GT masks (18 epochs). We then finetune the model on COCO annotated masks (80 categories) for standard 1x schedule. Similar to previous training from random initialization setup [21], we use GroupNorm [55] for long training with small batch size.

Results are summarized in Figure 5. When evaluated on COCO categories (same as training, closed-world), pre-training by pseudo-GT masks performs slightly worse than supervised label pre-training (-1.4%). On open-world setup, however, pseudo-GT pre-training consistently outperforms supervised training. We note that different from closed-world [21], ImageNet supervised pre-training is a strong initialization for open-world (see supplementary). In addition, we observe two promising scaling behaviors of pseudo-GT pretraining: a. using more masks per image, despite some being noisy, improves performance; using more images/ pixels improves performance for both closed-world and open-world instance segmentation. We show similar results of Pre-training on images from OpenImages [31] in supplementary materials. Finetuning on both COCO annotated masks and PA generated Pseudo-masks on COCO images provides additional gain (last row in Table 7).

## 6. Conclusion

We have presented GGN, a novel approach for open-world instance segmentation which combines learned semantic boundaries with grouping to generate additional pseudo ground truth for instance-level training. GGNs significantly outperform baselines on various benchmarks. GGN is on par with state-of-the-art approaches, e.g., OLN [27], and when combined with OLN, GGN obtains an additional 6.8%, establishing new state-of-the-art results for open-world instance segmentation. Finally, we showed that GGN is robust when evaluated “in the wild” and benefits from training on additional unlabeled data.

**Acknowledgement.** We thank Ross Girshick for the discussion on baselines and grouping methods and Abhijit Ogale for the discussion about open-world setting.



## References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels, 2010. 2
- [2] Pablo Arbelaez. Boundary extraction in natural images using ultrametric contour maps. In *CVPR Workshops*, 2006. 4, 6
- [3] Pablo Arbelaez, Michael Maire, Charless C. Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):898–916, 2011. 4
- [4] Pablo Andrés Arbeláez, Jordi Pont-Tuset, Jonathan T. Barron, Ferran Marqués, and Jitendra Malik. Multiscale combinatorial grouping. In *CVPR*, 2014. 2, 4
- [5] Abhijit Bendale and Terrance Boulton. Towards open world recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1893–1902, 2015. 3
- [6] M. V. D. Bergh, G. Roig, X. Boix, S. Manen, and L. V. Gool. Online video seeds for temporal window objectness. In *2013 IEEE International Conference on Computer Vision*, pages 377–384, 2013. 2
- [7] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(11):1222–1239, 2001. 2
- [8] Neelima Chavali, Harsh Agrawal, Aroma Mahendru, and Dhruv Batra. Object-proposal evaluation protocol is ‘gameable’. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 835–844, 2016. 7
- [9] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 5
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 2
- [11] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 2, 6
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 5, 8
- [13] Ian Endres and Derek Hoiem. Category independent object proposals. In *ECCV*, 2010. 4
- [14] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, Jan. 2015. 5
- [15] Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Efficient graph-based image segmentation. *IJCV*, 59(2):167–181, 2004. 2, 4
- [16] Charless C. Fowlkes, David R. Martin, and Jitendra Malik. Learning affinity functions for image segmentation: Combining patch-based and gradient-based approaches. In *CVPR*, 2003. 2
- [17] Naiyu Gao, Yanhu Shan, Yupei Wang, Xin Zhao, Yinan Yu, Ming Yang, and Kaiqi Huang. Ssap: Single-shot instance segmentation with affinity pyramid. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 642–651, 2019. 2
- [18] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph-based video segmentation. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2141–2148, 2010. 2, 6
- [19] A. Gupta, P. Dollár, and R. Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5351–5359, 2019. 5
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735, 2020. 2
- [21] Kaiming He, Ross B. Girshick, and Piotr Dollár. Rethinking imagenet pre-training. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4917–4926, 2019. 8
- [22] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *ICCV*, pages 2980–2988, 2017. 1, 2, 5
- [23] Xuming He, Richard S. Zemel, and Miguel Á. Carreira-Perpiñán. Multiscale conditional random fields for image labeling. In *CVPR*, 2004. 2
- [24] Jianbo Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000. 2
- [25] K J Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *CVPR*, 2021. 3
- [26] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *CoRR*, abs/1705.06950, 2017. 5
- [27] Dahun Kim, Tsung-Yi Lin, Anelia Angelova, In So Kweon, and Weicheng Kuo. Learning open-world object proposals without learning to classify. *CoRR*, abs/2108.06753, 2021. 2, 3, 4, 5, 7, 8
- [28] Tae Hoon Kim, Kyoung Mu Lee, and Sang Uk Lee. Learning full pairwise affinities for spectral segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(7):1690–1703, 2013. 2
- [29] Alexander Kirillov, Kaiming He, Ross B. Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *CVPR*, 2019. 2
- [30] Shu Kong and Deva Ramanan. Opendata: Open-set recognition via open data generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 813–822, October 2021. 3
- [31] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan

- Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020. 8
- [32] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T Freeman. Learning the depths of moving people by watching frozen people. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2019. 6
- [33] Tsung-Yi Lin, M. Maire, Serge J. Belongie, James Hays, P. Perona, D. Ramanan, Piotr Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 3, 5
- [34] Sifei Liu, Shalini De Mello, Jinwei Gu, Guangyu Zhong, Ming-Hsuan Yang, and Jan Kautz. Learning affinity via spatial propagation networks. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 2017. 2
- [35] Yiding Liu, Si Cheng, Yang, Bin Li, Wen gang Zhou, Jizheng Xu, Houqiang Li, and Yan Lu. Affinity derivation and graph merge for instance segmentation. In *ECCV*, 2018. 2
- [36] Yang Liu, Idil Esen Zulfikar, Jonathon Luiten, Achal Dave, Aljosa Osep, Deva Ramanan, Bastian Leibe, and Laura Leal-Taixé. Opening up open-world tracking. *CoRR*, abs/2104.11221, 2021. 3
- [37] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 2, 5
- [38] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019. 3
- [39] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 2, 3
- [40] Michael Maire, Takuya Narihira, and Stella X. Yu. Affinity cnn: Learning pixel-centric pairwise relations for figure/ground embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [41] Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Arbeláez, and Luc Van Gool. Convolutional oriented boundaries. In *European Conference on Computer Vision (ECCV)*, 2016. 2, 3
- [42] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2536–2544, 2016. 2
- [43] Pedro O. Pinheiro, Ronan Collobert, and Piotr Dollár. Learning to segment object candidates. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS'15*, page 1990–1998, Cambridge, MA, USA, 2015. MIT Press. 3, 5
- [44] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, 2016. 5
- [45] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28, pages 91–99. Curran Associates, Inc., 2015. 4, 5
- [46] Ellen Riloff. Automatically generating extraction patterns from untagged text. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 2, AAAI'96*, page 1044–1049. AAAI Press, 1996. 2
- [47] H. Scudder. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371, 1965. 2
- [48] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, 2000. 4
- [49] Xavier Soria, Edgar Riba, and Angel Sappa. Dense extreme inception network: Towards a robust cnn model for edge detection. In *The IEEE Winter Conference on Applications of Computer Vision (WACV '20)*, 2020. 4, 6
- [50] Srinivas C. Turaga, Kevin L. Briggman, Moritz Helmstaedter, Winfried Denk, and H. Sebastian Seung. Maximin affinity learning of image segmentation. In *Neural Information Processing Systems*, 2009. 2
- [51] K. E. A. van de Sande, J. R. R. Uijlings, T. Gevers, and A. W. M. Smeulders. Segmentation as selective search for object recognition. In *ICCV*, pages 1879–1886, 2011. 1, 2, 5, 8
- [52] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan L. Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 2021. 2
- [53] Rui Wang, Dhruv Kumar Mahajan, and Vignesh Ramanan. What leads to generalization of object proposals? In *ECCV Workshops*, 2020. 3
- [54] Weiyao Wang, Matt Feiszli, Heng Wang, and Du Tran. Unidentified video objects: A benchmark for dense, open-world segmentation. In *ICCV*, 2021. 2, 3, 5, 7
- [55] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, 2018. 8
- [56] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, 2018. 3, 5
- [57] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1395–1403, 2015. 2, 3, 6
- [58] Chenliang Xu, Caiming Xiong, and Jason J. Corso. Streaming hierarchical video segmentation. In Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid, editors, *Computer Vision – ECCV 2012*, pages 626–639, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. 2
- [59] David Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. In *Proceedings of the 33rd Annual Meeting on Association for Computational Linguistics*,

- ACL '95, page 189–196, USA, 1995. Association for Computational Linguistics. [2](#)
- [60] Xiao Zhang and Michael Maire. Self-supervised visual representation learning from hierarchical grouping. In *NeurIPS*, 2020. [4](#)
- [61] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba. Scene parsing through ade20k dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5122–5130, 2017. [1](#), [5](#), [7](#)
- [62] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin Dogus Cubuk, and Quoc V. Le. Rethinking pre-training and self-training. *ArXiv*, abs/2006.06882, 2020. [2](#)