

Towards Open Vocabulary Learning: A Survey

Jianzong Wu¹, Graduate Student Member, IEEE, Xiangtai Li², Member, IEEE, Shilin Xu³, Haobo Yuan⁴, Henghui Ding⁵, Yibo Yang⁶, Xia Li⁷, Jiangning Zhang⁸, Member, IEEE, Yunhai Tong⁹, Member, IEEE, Xudong Jiang¹⁰, Fellow, IEEE, Bernard Ghanem¹¹, Member, IEEE, and Dacheng Tao¹², Fellow, IEEE

(Survey Paper)

Abstract—In the field of visual scene understanding, deep neural networks have made impressive advancements in various core tasks like segmentation, tracking, and detection. However, most approaches operate on the close-set assumption, meaning that the model can only identify pre-defined categories that are present in the training set. Recently, open vocabulary settings were proposed due to the rapid progress of vision language pre-training. These new approaches seek to locate and recognize categories beyond the annotated label space. The open vocabulary approach is more general, practical, and effective than weakly supervised and zero-shot settings. This paper thoroughly reviews open vocabulary learning, summarizing and analyzing recent developments in the field. In particular, we begin by juxtaposing open vocabulary learning with analogous concepts such as zero-shot learning, open-set recognition, and out-of-distribution detection. Subsequently, we examine several pertinent tasks within the realms of segmentation and detection, encompassing long-tail problems, few-shot, and zero-shot settings. As a foundation for our method survey, we first elucidate the fundamental principles of detection and segmentation in close-set scenarios. Next, we examine various contexts where open vocabulary learning is employed, pinpointing recurring design elements and central themes. This is followed by a comparative analysis of recent detection and segmentation methodologies in commonly used datasets and benchmarks. Our

review culminates with a synthesis of insights, challenges, and discourse on prospective research trajectories. To our knowledge, this constitutes the inaugural exhaustive literature review on open vocabulary learning.

Index Terms—Open vocabulary, scene understanding, object detection, segmentation, survey.

I. INTRODUCTION

DEEP neural networks have revolutionized scene understanding tasks, including object detection, segmentation, and tracking [1], [2], [3], [4]. Nonetheless, using conventional approaches for real-world applications can be challenging due to limitations such as inadequate class annotations, close-set class definitions, and costly labeling expenses. These limitations can increase the difficulty and cost of implementing a deep model on new scenes, particularly when the number of categories or concepts in the scene is much larger than what is included in the training dataset.

For example, object detection is a core computer vision task involving scene understanding. It requires human annotations for each category and each object location, which can be costly and time-consuming. For instance, the COCO dataset [5], widely used for benchmarking object detection algorithms, only includes 80 categories. But in reality, natural scene images often have more than 80 different types of objects. We would need to incur significant annotation costs to extend object detectors to cover all these categories. Current research focuses on developing methods to train more flexible object detectors on the subset of COCO with base classes and let them identify new or unfamiliar objects without requiring additional annotations.

Several previous solutions adopt zero-shot learning (ZSL) [6], [7], [8]. These approaches extend a detector to generalize from annotated (seen) object classes to other (unseen) categories. The annotations of seen object classes are used during the training, while the annotations of unseen classes are *strictly* unavailable during training. Most approaches adopt word embedding projection to constitute the classifier for unseen class classification.

However, these approaches come with several limitations. ZSL, in particular, is highly restrictive. Typically, these methods lack examples of unseen objects and treat these objects as background objects during training. As a result, during inference, the model identifies novel classes solely based on their pre-defined word embeddings [9], [10], thereby limiting exploration of the visual information and relationships of those unseen classes.

Manuscript received 22 July 2023; revised 24 January 2024; accepted 29 January 2024. Date of publication 5 February 2024; date of current version 5 June 2024. This work was supported in part by the National Key Research and Development Program of China under Grant 2023YFC3807600 and in part by the interdisciplinary doctoral under Grant iDoc 2021-360 from the Personalized Health and Related Technologies (PHRT) of the ETH domain. Recommended for acceptance by V. Ferrari. (Jianzong Wu, Xiangtai Li, Shilin Xu, and Haobo Yuan contributed equally to this work.) (Corresponding author: Xiangtai Li.)

Jianzong Wu, Xiangtai Li, Shilin Xu, and Yunhai Tong are with the National Key Laboratory of General Artificial Intelligence, School of Intelligence Science and Technology, Peking University, Beijing 100871, China (e-mail: jzww@stu.pku.edu.cn; lxt@pku.edu.cn; xushilin@stu.pku.edu.cn; yhtong@pku.edu.cn).

Haobo Yuan is with Wuhan University, Wuhan 430072, China (e-mail: yuan-haobo@whu.edu.cn).

Henghui Ding and Xudong Jiang are with Nanyang Technological University, Singapore 639798 (e-mail: henghui.ding@gmail.com; exdjiang@ntu.edu.sg).

Yibo Yang and Bernard Ghanem are with the King Abdullah University of Science and Technology, Thuwal 23955, Saudi Arabia (e-mail: ibo@pku.edu.cn; bernard.ghanem@kaust.edu.sa).

Xia Li is with ETH Zurich, 8092 Zürich, Switzerland (e-mail: xia.li@inf.ethz.ch).

Jiangning Zhang is with Zhejiang University, Hangzhou 310027, China (e-mail: 186368@zju.edu.cn).

Dacheng Tao is with the University of Sydney, Sydney, NSW 2050, Australia (e-mail: dacheng.tao@gmail.com).

We keep tracing related works at <https://github.com/jianzongwu/Awesome-Open-Vocabulary>.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TPAMI.2024.3361862>, provided by the authors.

Digital Object Identifier 10.1109/TPAMI.2024.3361862

This is why ZSL approaches have been shown to yield unsatisfying results in novel classes.

Open vocabulary learning is proposed to handle the above issue and make the detector extendable. It has been successfully applied to multiple tasks, e.g., segmentation, detection, video understanding, scene understanding, etc. In particular, open vocabulary object detection [11], [12] is first proposed. Also, open vocabulary segmentation is proposed [13], [14]. Similarly, the model trains on base classes and inferences on both base and novel classes. The critical difference between zero-shot learning and open vocabulary is that one can use *visual-related language vocabulary data* like image captions as auxiliary supervision in open vocabulary settings. The motivations to use language data as auxiliary weak supervision are: 1) Language data requires less labeling effort and thus is more cost-friendly. The visual-related language data, like image captions, is widely available and more affordable than box or mask annotations. 2) Language data provides a more extensive vocabulary size and thus is more extendable and general. For example, words in captions are not limited to the pre-defined base categories. It may contain novel class names, attributes, and motions of objects. Incorporating captions in training has been proven to be extremely useful in helping improve the models' scalability.

Moreover, recently, visual language models (VLMs) [15], [16], which pre-train themselves on large-scale image-text pairs, show remarkable zero-shot performance on various vision tasks. The VLMs align images and language vocabularies into the same feature space, fulfilling the visual and language data gap. Many open vocabulary methods effectively eliminate the distinction between close-set and open-set scenarios by utilizing the alignment learned in VLMs, making them highly suitable for practical applications. For example, An open vocabulary object detector can be easily extended to other domains according to the demand without the need to gather relevant data or incur additional labeling expenses.

As open vocabulary models continue to advance rapidly and demonstrate impressive results, it is worthwhile to track and compare recent research on open vocabulary learning. Several surveys work on low-shot learning [17], [18], [19], [20], [21], including few-shot learning and zero-shot learning. There are also several surveys on multi-modal learning [22], [23], including using transformer for vision language tasks [22], [24] and vision language pre-training [25]. However, these works focus on learning with few examples, multi-modal fusion, or pre-training for better feature representation. As far as we know, there haven't been any surveys that thoroughly summarize the latest developments in open vocabulary learning, including methods, settings, benchmarks, and the use of vision foundation models. We aim to fulfill the blank with this work.

Contribution: In this survey, we systematically track and summarize recent literature on open vocabulary learning, including object detection, segmentation, video understanding, and 3D scene understanding. The survey covers the most representative works in each domain by extracting the common technical details. It also contains the background of open vocabulary learning and related concepts comparison, including zero-shot learning (ZSL) [7], open set recognition (OSR) [26], [27], and out-of-distribution detection (OOD) [28]. It also includes

large-scale visual language models and representative detection and segmentation works, which makes the survey self-contained. In addition, we present a comprehensive analysis and comparison of benchmarks and settings for each specific domain. As far as we know, we are the first to concentrate on the specific area. Finally, since the field of open vocabulary learning is rapidly evolving, we may not be able to keep up with all the latest developments. We welcome researchers to contact us and share their new findings in this area to keep us updated. Those new works will be included and discussed in the revised version.

Survey pipeline: In Section II, we will cover the background knowledge, including definition, datasets, metrics, and related research domains. Then, in Section III, we conduct main reviews on various methods according to different tasks. In particular, we will first include the preliminary knowledge of close-set detection and segmentation methods. Then, we present the method details of each direction, including detection, segmentation, video understanding, and 3D scene understanding. Next, we point out challenges and future directions in Section IV and conclude Section V. Finally, the appendix compares the results for different tasks and benchmarks.

II. BACKGROUND

Overview: In this section, we first present the concept definition of open vocabulary and related concepts comparison. Then, we present a historical review of open vocabulary learning and point out several representatives. Next, we present the standard datasets and metrics. We also present the unified notations for open vocabulary object detection and segmentation tasks. Finally, we review the related research domains.

A. Concept Definition

We take the classification task for concept illustration. The supervised methods assume that the training data and testing data share the same closed-set label space. However, the model trained under this assumption cannot be extended to new categories. To address this issue, researchers have introduced several new concepts like open-set learning and zero-shot learning, ultimately leading to open vocabulary learning.

These concepts have similar settings and notations. In particular, the examples are classified into base and novel classes (or called out-of-distribution examples). The base classes can be accessed for training, while the novel classes are not. We denote the base classes in the label set space C_B and the novel classes in the label set space C_N .

Due to the low performance on novel classes in earlier works for open-set, open world, and OOD tasks [11], [28], [30], [31] and the easier acquisition of image-text pairs, recent works [11], [12] propose the open vocabulary setting. It allows using additional low-cost training data [32] or pre-trained vision language models like CLIP [15], which have much larger language vocabularies C_L . It may contain the concepts of both C_B and C_N , but its main goal is to enable the model to generalize across more classes in the open domain. In mathematical terms and data view, **open vocabulary learning** can be defined as follows:

Training Data (D_{train}): The training dataset is a collection of data-label pairs, where each pair consists of an input x_i and its associated label y_i . The input x_i can be an image or a video

according to the task type, and in scene understanding tasks such as detection and segmentation, the label y_i also includes visual labels, such as bounding boxes or masks, in addition to the class labels. Alongside the standard input-label pairs, open vocabulary learning includes vision-aware language vocabulary data, represented as l_i . The language vocabulary data can be image-caption data or vision-aware class name embeddings in VLMs like CLIP [15]. The language augmented training set can be denoted as $D_{train} = \{(x_1, y_1, l_1), (x_2, y_2, l_2), \dots, (x_n, y_n, l_n)\}$, where n is the training data length, x_i is the visual data, y_i is the label from base classes C_B , and l_i is the associated language data from a large vocabulary space C_L . Note that C_L is not strictly required to contain C_B or C_N , as the language vocabulary may not cover all the class names in the vision data. On the contrary, C_L may also have words out of the pre-defined novel categories, which can further extend models' generalizability.

Evaluation Data (D_{eval}): The evaluation dataset is similarly a collection of data pairs. However, the labels for evaluation data include both base classes and novel classes. This can be represented as $D_{eval} = \{(x'_1, y'_1), (x'_2, y'_2), \dots, (x'_m, y'_m)\}$, where m is the evaluation data length, y'_i belongs to either C_B or C_N . During the evaluation, open vocabulary methods need to predict y'_i given x'_i in the realm $y'_i \in (C_B \cup C_N)$.

Here, we emphasize the concept differences between open vocabulary learning and analogous concepts as follows:

- **Open-Set Learning:** Open-set learning aims to classify known classes and reject unknown classes during testing [30], [33], [34]. Concretely, the training data is $D_{train} = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$, where y_i is from base classes C_B . These are also called known classes in open-set learning [30], [35], [36]. The evaluation data is $D_{eval} = \{(x'_1, y'_1), (x'_2, y'_2), \dots, (x'_m, y'_m)\}$, where $y'_i \in (C_B \cup \{u\})$. u is a single class that represents the 'unknown' class. Open-set learning tasks do not require further classifying the unknown classes.
- **Open World Learning:** Open world learning addresses real-world environments' dynamic and ever-evolving nature by recognizing and learning new categories incrementally over time, without the need for complete system retraining [37], [38]. This process includes classifying known objects, identifying unknowns, labeling the unknowns by humans, and incrementally learning new categories as they are labeled and added to the system. Suppose an open world learning process contains T learning steps. In the t th learning step, where $t \in \{1, 2, \dots, T\}$, the training data is $D_{train}^t = \{(x_1^t, y_1^t), (x_2^t, y_2^t), \dots, (x_n^t, y_n^t)\}$, where y_i^t is in the labels annotated in the t th step C_{label}^t , and x_i^t is the input data corresponding to the newly added labels. The evaluation data is $D_{eval}^t = \{(x'_1, y'_1), (x'_2, y'_2), \dots, (x'_m, y'_m)\}$, where $y'_i \in (C_B^t \cup \{u\})$. u is a single class that represents the 'unknown' class. C_B^t is the known class label set for timestamp t . It is an accumulation of labeled classes in the current and previous steps. $C_B^t = \bigcup_{k=0}^t C_{label}^k$. Note that evaluation input x'_i keeps the same across learning steps.
- **Out-of-Distribution Detection:** Out-of-distribution (OOD) detection focuses on the ability to detect data that is different in some way from the data used




Fig. 1. Concepts comparison between open-set/open world/out-of-distribution detection (OOD), zero-shot and open vocabulary. Different shapes represent different novel categories. Colors represent the predictions of the novel objects. (a), in the open-set/open World/OOD settings, the model only needs to identify novel classes and mark them as “unknown”. (b), in the zero-shot setting, a model must classify unknown classes into specific categories. (c), in the open vocabulary settings, the model can classify novel classes with the help of large language vocabulary knowledge C_L .

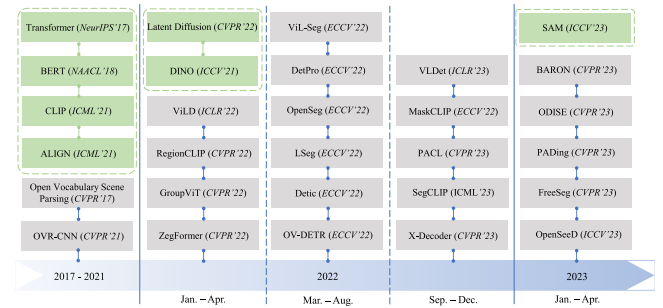


Fig. 2. Timeline of open vocabulary learning. The gray boxes indicate representative works. Green boxes indicate the foundation models and VLMs. In open vocabulary learning, many works exploit the knowledge learned by pre-trained vision foundation models like Swin [29] and VLMs like CLIP [15]. Recently, some works also explore the use of diffusion models in this setting.

during training [28]. The training data is $D_{train} = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ and is supposed to be sampled from a probability distribution $P_{train}(X, Y)$. During testing, the model encounters data x'_i that are sampled from a different distribution $P_{OOD}(X')$, which is not represented in the training dataset. The goal is to identify or appropriately handle these OOD samples. Metrics for OOD detection often involve measuring the model's certainty or confidence in its predictions, using scores like softmax probability, and checking whether it correctly identifies OOD samples [28].

- **Zero-Shot Learning:** Zero-shot learning aims to recognize objects or concepts not seen during training. The training data is $D_{train} = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$, where y is from base classes C_B , which is usually called seen classes in zero-shot learning [7], [39]. The evaluation data is $D_{eval} = \{(x'_1, y'_1), (x'_2, y'_2), \dots, (x'_m, y'_m)\}$, where $y'_i \in C_N$. C_N is novel classes or unseen classes. Zero-shot learning tasks require models to make clear classifications among the new unseen classes.

We briefly compare these concepts in Fig. 1, including open vocabulary, open-set, open world, OOD, and zero-shot.

B. History and Roadmap

Before introducing the open vocabulary setting, reviewing the progress of open vocabulary learning is necessary. In Fig. 2, we summarize the timeline of open vocabulary learning. Localization and classification of arbitrary objects in the wild




Fig. 3. Summarization on open vocabulary learning works. (a) The number of research works still increases per year. (b) Detection and segmentation have more papers than 3D and video. (c) indicates each direction per year. The results are obtained on 2024/11/15.

have been challenging problems due to the limitations of existing datasets. The concept of open vocabulary in scene understanding comes from the work [40], where the authors build a joint image pixel and word concept embedding framework. The concept is hierarchically divided. Then, multi-modal pre-training was well studied with the rise of BERT [9] in NLP. Motivated by the process of vision language pre-training, OVR-CNN [11] proposed the concept of open vocabulary object detection, where the caption data are used for connecting novel classes semantics and visual region. Later on, CLIP was presented and open-sourced. After that, ViLD [12] is the first work that uses the knowledge of CLIP to build open vocabulary object detection. Meanwhile, LSeg [14] first explored the CLIP knowledge of language-driven segmentation tasks. After these works, recently, there have been more and more works on improving the performance of open vocabulary detectors or building new benchmarks for various settings. SAM [41] is proposed to build the segmentation foundation model, which is trained by billion-level masks. Combined with CLIP, SAM can also achieve good zero-shot segmentation without fine-tuning. Recently, with the rapid process of large language model (LLM) [42], open vocabulary learning has become a more promising direction since more language knowledge can be embedded in multi-modal architecture. As shown in Fig. 3(a), the number of research works in open vocabulary learning has increased significantly since 2021. We also summarize the statistics of different directions in Fig. 3(b) and (c). The details of these directions can be found in Section III.

C. Tasks, Datasets, and Metrics

- **Tasks:** Open vocabulary learning has included a wide range of computer vision tasks, including object detection [11], [43], [44], segmentation [45], [46], [47], video understanding, and 3D scene understanding. The center goal of these tasks is similar, recognizing the novel classes with the aid of large vocabulary knowledge for their corresponding tasks. In this survey, we mainly focus on the methods of scene understanding tasks, including object detection, instance segmentation, semantic segmentation, and object tracking. Nonetheless, we also consider other closely related tasks, such as open vocabulary attribution prediction, video classification, and point cloud classification.

- **Datasets:** For object detection, the common datasets are COCO [5] and LVIS [48]. Recently, a more challenging dataset, v3Det [49] with more than 10,000 categories, is proposed. For image segmentation, the most commonly used datasets are COCO [5], ADE20k [50], PASCAL-VOC 2012 [51], PASCAL-Context [52], and Cityscapes [53]. For video segmentation and tracking, the frequently used datasets are VSPW [54], Youtube-VIS [55], LV-VIS [56], MOSE [57], and TAO [58].
- **Metrics:** For detection tasks, the commonly used metrics are mean average precision (mAP) and mean average recall (mAR) for both base and novel classes. Among segmentation tasks, the commonly used metrics are mean intersection over union (mIoU) for semantic segmentation, mask-based mAP for instance segmentation, and panoptic quality (PQ) for panoptic segmentation.

D. Related Research Domains

- **Open-Set Recognition:** The concept of Open-Set Recognition (OSR) addresses the challenge of identifying unknown classes during classification tasks. In traditional classification systems, models are trained to recognize a finite set of known classes, but in real-world scenarios, they may encounter data that doesn't belong to these predefined categories. During testing, OSR aims to classify known classes seen during the training and reject unknown classes that are unseen [30]. The main methods used for OSR can be broadly categorized into two types: discriminative models [33], [34], [35], [36], [59] and generative models [60], [61], [62], [63], [64], [65]. The discriminative models focus on differentiating between known classes and identifying unknown classes by enhancing the boundary or margin between these classes. The generative models are either instance generation-based [60], [61], [62], [63], [64] or non-instance generation-based [65]. They emphasize generating new instances or features to improve the ability of the system to recognize new, unseen classes during testing. Each category employs specific techniques and approaches to address the challenges inherent in open-set recognition, which involves dealing with classes not seen during the training phase. In [66] extends the OSR task to require further identifying novel classes.

- *Open World Learning*: Open world learning involves identifying and labeling new, unknown categories (novel unknowns) [37]. This process includes incrementally learning new categories as they are labeled and added to the system. The goal is to create a system that remains robust to unknown categories and adapts continually to include new information, balancing the risks associated with open spaces in the learning model. In [37] first propose open word recognition. Recent works also explore the open world object detection [31], [67] task.
- *Out-of-Distribution Detection*: Out-of-distribution (OOD) detection methods can be structured into several categories [28]. Classification-Based Methods [68], [69], [70], [71]: These include output-based methods [68], [69], label space redesign [70], and OOD data generation techniques [71]. Density-Based Methods [72], [73], [74]: This category involves methods that detect OOD by modeling data density. Distance-Based Methods [75], [76]: These methods use distance metrics, typically in the feature space, to identify OOD instances. Reconstruction-Based Methods [77], [78], [79]: This approach achieves OOD detection that features reconstruction capabilities.
- *Zero-Shot Detection and Segmentation*: This task aims to segment classes that have not been encountered during training. Two streams of work have emerged: discriminative methods [80], [81], [82], [83], [84] and generative methods [85], [86], [87], [88]. Representative works in this field include SPNet [80] and ZS3Net [86]. SPNet [80] maps each pixel to a semantic word embedding space and projects pixel features onto class probabilities using a fixed semantic word embedding [89], [90] projection matrix. On the other hand, ZS3Net [86] first trains a generative model to produce pixel-wise features for unseen classes based on word embeddings. With these synthetic features, the model can be trained in a supervised manner. Both of these works treat zero-shot detection and segmentation as a pixel-level zero-shot classification problem. However, this formulation is not robust for zero-shot learning, as text embeddings are typically used to describe objects/segments rather than individual pixels. Subsequent works [85], [87], [88], [91], [92], [93] follow this formulation to address different challenges in zero-shot learning. In a weaker assumption where unlabeled pixels from unseen classes are available in the training images, self-training [86] is commonly employed. Despite promising results, self-training often requires model retraining whenever a new class appears. ZSI [39] also employs region-level classification for bounding boxes but focuses on instance segmentation rather than semantic segmentation. More recently, PAD-ing [85] proposes a unified framework to tackle zero-shot semantic segmentation, zero-shot instance segmentation, and zero-shot panoptic segmentation.

Most approaches in open vocabulary learning are based on zero-shot learning settings, such as replacing the fixed classifier with language embeddings. However, these methods struggle to generalize well to novel classes due to the absence of novel class knowledge. Consequently, their

performance is limited, and they are not practical for real-world applications.

- *Long-tail Object Detection and Instance Segmentation*: This task addresses the challenge of class imbalance in instance segmentation. Many approaches tackle this issue through techniques such as data re-sampling [48], [94], loss re-weighting [95], and decoupled training [96]. Specifically, some studies [94] employ image-level re-sampling. However, these methods tend to introduce bias in instance co-occurrence. To address this issue, other works [97] focus on more refined re-sampling techniques at the instance or feature level. Regarding loss re-weighting, most studies [95] rebalance the ratio of positive and negative samples during training. Additionally, decoupled training methods [96], [98] introduce different calibration frameworks to enhance classification results. Resolving long-tail object detection can lead to improved accuracy of rare classes. However, these methods currently cannot be applied to the detection of novel classes.
- *Few-Shot Detection and Segmentation*: Few-shot object detection aims to expand the detection capabilities of a model using only a few labeled samples. Several approaches [21], [99], [100], [101] have been proposed to advance this field. Notably, TFA [102] introduces a simple two-phase fine-tuning method, while DeFRCN [103] separates the training of RPN features and RoI classification. Moreover, SRR-FSD [104] combines multi-modal inputs while LVC [105] proposes a pipeline to generate additional examples for novel object detection and train a more robust detector. Few-shot segmentation comprises Few-Shot Semantic Segmentation (FSSS) and Few-Shot Instance Segmentation (FSIS). FSSS involves performing pixel-level classification on query images. Previous approaches [106], [107] typically build category prototypes from support images and segment the query image by computing the similarity distance between each prototype and query features. On the other hand, FSIS aims to detect and segment objects with only a few examples. FSIS methods can be categorized into single-branch and dual-branch methods. The former [108] primarily focuses on designing the classification head, while the latter [109], [110], [111] introduce an additional support branch to compute class prototypes or re-weighting vectors of support images. This support branch assists the segmenter in identifying target category features through feature aggregation. For example, Meta R-CNN [111] performs channel-wise multiplication on RoI features, and FGN [110] aggregates channel-wise features at three stages: RPN, detection head, and mask head. However, few-shot learning still requires examples of novel classes during training, and such data may not be available.

III. METHODS: A SURVEY

Overview: In this section, we first review preliminary knowledge and vision language modeling by extending close-set detectors into open vocabulary detectors via VLMs in

TABLE I
REPRESENTATIVE WORKS SUMMARIZATION AND COMPARISON IN SECTIONS III-C AND III-D

Method	Task	Text Training Data	Vision Training Annotations	Text Model	Vision Model	Highlight
Open Vocabulary Detection (Sec. 3.3)						
OVR-CNN [11]	OVOD	Captions	Bounding Boxes (Base)	BERT	Faster R-CNN	The first method that proposes OVOD and adopts grounded caption pre-training.
ViLD [12]	OVOD	None	Bounding Boxes (Base)	CLIP-text	CLIP-vision + Faster R-CNN	The first method that distills knowledge from the pre-trained CLIP model.
VL-PLM [215]	OVOD	None	Bounding Boxes (Base)	CLIP-text	CLIP-vision + Faster R-CNN	Generate pseudo-labels for novel classes using pre-trained OVOD.
RegionCLIP [216]	OVOD	Captions	Bounding Boxes (Base + Pseudo Novel)	CLIP-text	CLIP-vision + Faster R-CNN	Creates region-text pairs as pseudo labels using CLIP and pre-train the detector in the first stage.
Detic [134]	OVOD	ImageNet	Bounding Boxes (Base + Pseudo Novel)	CLIP-text	Centernet2	Propose a weakly supervised approach that is training rare classes with image-level annotations.
DetPro [217]	OVOD	ImageNet	Bounding Boxes (Base + Pseudo Novel)	CLIP-text	Centernet2	learn continuous prompt representations for open vocabulary object detection based on the pre-trained vision-language model.
OV-DETR [218]	OVOD	None	Bounding Boxes (Base)	CLIP-text	CLIP-vision + DETR	Introduces a conditional binary matching mechanism to let DETR model generalize to queries from unseen classes.
CORA [219]	OVOD	None	Bounding Boxes (Base)	CLIP-text	CLIP-vision + DETR	Propose Anchor Pre-Matching strategy to reduce both training and inference time for conditional binary matching.
VLDet [44]	OVOD	Captions	Bounding Boxes (Base)	CLIP-text	Faster R-CNN	Aligns image regions with words in captions by a set matching method.
F-VLM [137]	OVOD	None	Bounding Boxes (Base)	CLIP-text	CLIP-vision + Mask R-CNN	Train the detector with frozen VLMs and combine scores of joint detection and VLMs.
BARON [113]	OVOD	None	Bounding Boxes (Base)	CLIP-text	Faster R-CNN	Aligning Bag of Regions for Open Vocabulary Object Detection.
OWLv2 [135]	OVOD	None	Bounding Boxes (Base + Pseudo Novel)	CLIP-text	ViT + detection head	Generate pseudo-labels from WebLI dataset and train the detector with the generated datasets.
MaMMUT [220]	OVOD	Captions	Bounding Boxes (Base + Pseudo Novel)	CLIP-text	ViT + detection head	Joint pre-train with multi-modal tasks to benefit the novel object detection.
Open Vocabulary Segmentation (Sec. 3.4)						
LSeg [14]	OVSS	None	Segmentation Masks (Base)	CLIP-text	ViT	Aligns text embeddings from the VLM model with pixel features.
ZegFormer [114]	OVSS	None	Segmentation Masks (Base)	CLIP-text	CLIP-vision + Query-based Transformer Decoder	Decouple segmentation and classification by generating class-agnostic segmentation masks then classify each mask.
PADing [86]	OVSS	None	Segmentation Masks (Base)	CLIP-text	None	Introduce a generative model to synthesize features for unseen categories and achieve universal open-vocabulary segmentation.
OpenSeg [46]	OVSS	Captions	Segmentation Masks (Base)	ALIGN	EfficientNet-B7	Performs region-word grounding loss between mask features and word features.
MaskCLIP+ [120]	OVSS	None	Segmentation Masks (Pseudo All)	CLIP-text	CLIP-vision + DeepLabv2	Modifies CLIP so that it can output per-pixel feature maps.
OVSeg [121]	OVSS	Captions	Segmentation Masks (Base + Pseudo All)	CLIP-text	CLIP-vision + MaskFormer	Uses CLIP to match the proposed image regions with nouns in the captions to generate pseudo labels.
GroupViT [125]	OVSS	Captions	None	Transformer Encoder	ViT-S	Learns semantic segmentation only with caption data.
ViL-Seg [221]	OVSS	Captions	None	ViT-B	ViT-B	Learns semantic segmentation without pixel-level annotations using contrastive loss and clustering loss.
CGG [47]	OVIS	Captions	Segmentation Masks (Base)	BERT	Mask2Former	Fully exploits caption data using caption grounding and generation.
MaskCLIP [48]	OVPS	None	Segmentation Masks (Base)	CLIP-text	CLIP-vision + Mask2Former	Proposes Relative Mask Attention (RMA) modules to adapt cropped images to the pre-trained CLIP model.
ODISE [133]	OVPS	Captions	Segmentation Masks (Base)	CLIP-text	Stable Diffusion	Exploits the vision-language alignment learned by denoising diffusion models.
OVDiff [132]	OVSS	None	None	CLIP-text	Stable Diffusion	Proposes a prototype-based method. Use diffusion models to produce prototypes.
OpenSeed [128]	OVIS	None	Segmentation Masks & Boxes (Base)	UniCL	MaskDINO	Jointly learns from detection and segmentation data.
OVSegmentor [222]	OVSS	Captions	None	BERT	DINO	introduces masked entity completion and cross-image mask constituency objectives to improve training.

We list the training data, detectors or segmenters, VLMs, and highlighted features for comparison.

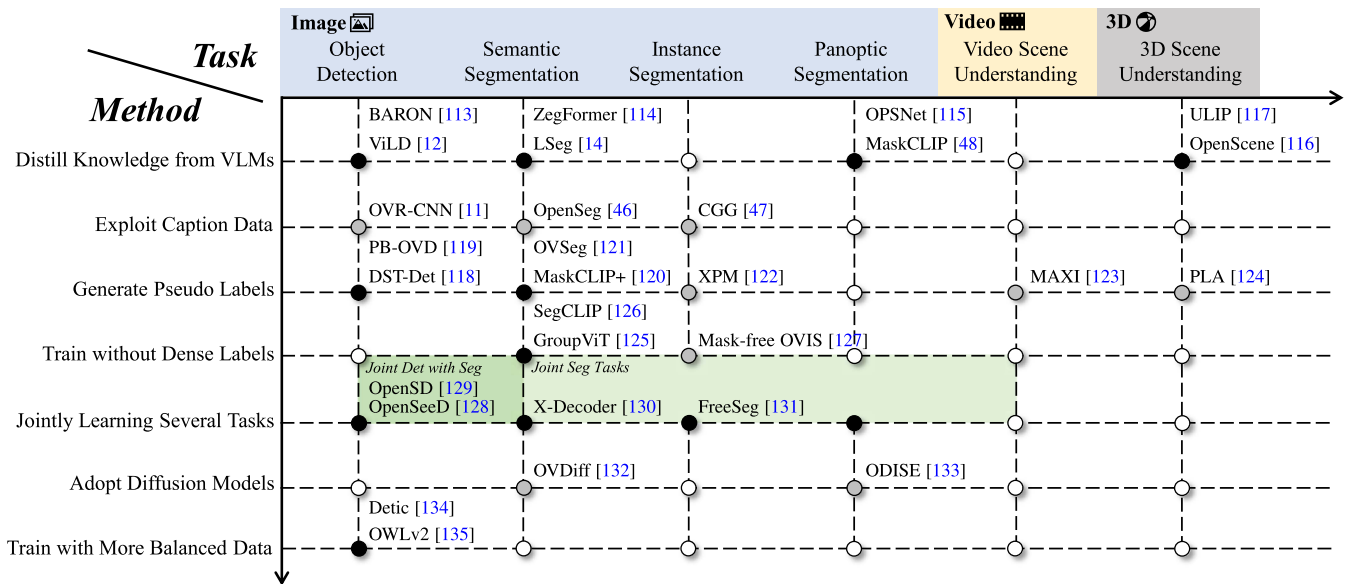


Fig. 4. Open vocabulary learning methods, organized by their tasks and approach types. We list several representative works here.

Sections III-A and III-B. Then we sequentially survey the six subsidiary tasks, including object detection (Section III-C), segmentation (Section III-D), video understanding (Section III-E), 3D scene understanding (Section III-F), and closely related tasks (Section III-G). Note that we only record and compare the most representative works. We list numerous works in Tables I, III, and IV. Moreover, since there are

several similar tasks (Section III-G), we also survey and compare other related tasks, including class agnostic detection and segmentation, open world object detection, and open-set panoptic segmentation.

Method taxonomy and relation of each subsection: We summarize the most commonly used methods among these different directions in Fig. 4. However, we cannot include all methods due

to the wide range of tasks. In each subsection, we summarize different methods according to different problems, segmentation, detection, video understanding, and closely related topics. We cluster methods according to common practice, such as knowledge distillation and region text pre-training. We start with object detection since it was first proposed. For segmentation tasks, we ignore the specific setting in Section III-D. We argue that despite these two directions, they may share similar ideas. However, most meta-architectures and tasks are different. Thus, we still review each direction individually. For video, 3D understanding, and other closely related topics, due to different task definitions and extra input cues, such as temporal information and multi-view inputs, we also survey two directions individually.

Motivation of survey organization: We argue that using the unified symbolic system to summarize all methods is hard. There are several reasons. (1), The task definitions and settings are different. In addition to different input and output formats, taking the open vocabulary segmentation as an example, there are at least three different settings [45], [46], [124] for open vocabulary segmentation. Thus, the design principles are different. For example, OpenSeg [45] uses supervised pretraining and contrastive losses while GroupViT [124] adopt unsupervised setting without mask annotations. (2), Different methods use different detectors or baseline models. Several works [11], [133] use R-CNN framework while several methods [46], [130] adopt query-based approaches. Several methods [135], [136] adopt the RPN's features to develop their methods, while query-based methods cannot perform these designs. (3), Different methods may also use different datasets during training. Several methods [133], [137], [138] mainly explore the dataset effect on open-vocabulary detection or segmentation. Since the supervision signals are different, it will hard to put all methods in one unified symbolic system. In summary, our survey focus on more comprehensive review to extract the common features of various domains in open vocabulary learning.

A. Preliminary

Pixel-based Object Detection and Segmentation: In general, it can be divided into two aspects: *semantic-level* tasks and *instance-level* tasks, where the difference lies in whether to distinguish each instance. For the former, we take semantic segmentation as an example. It was typically approached as a dense pixel classification problem, as initially proposed by FCN [139]. Then, the following works are all based on the FCN framework. These methods can be divided into the following categories, including better encoder-decoder frameworks [140], [141], larger kernels [142], multiscale pooling [143], [144], multiscale feature fusion [145], [146], [147], [148], non-local modeling [149], and better boundary delineation [150], [151], [152], [153]. After the transformer was proposed, with the goal of global context modeling, several works propose the variants of self-attention operators to replace the CNN prediction heads [154], [155].

For the latter, we take object detection and instance segmentation for illustration. Object detection aims to detect each

instance box and classify each instance. It mainly has two different categories: two-stage approaches and one-stage approaches. Two-stage approaches [2], [156] rely on an extra region proposal network (RPN) to recall foreground objects at the first stage. Then, the proposals (Region of Interests, RoI) with high scores are sent to the second-stage detection heads for further refinement. One-stage approaches [157], [158] directly output each box and label in a per-pixel manner. In particular, with the help of focal loss and feature pyramid networks, one-stage approaches can surpass the two-stage methods in several datasets. Both approaches need anchors for the regression of the bounding box. The anchors are the pixel locations where the objects are possibly emerging. Since two-stage approaches explicitly explore the foreground objects, where they have better recall than single-stage approaches. Two-stage approaches are more commonly used in open vocabulary object detection tasks.

Instance segmentation segmenting each object goes beyond object detection. Most instance segmentation approaches focus on how to represent instance masks beyond object detection, which can be divided into two categories: top-down approaches [2], [159] and bottom-up approaches [160], [161]. The former extends the object detector with an extra mask head. The designs of mask heads are various, including FCN heads [2], [162], diverse mask encodings [163], and dynamic kernels [159], [164]. The latter performs instance clustering from semantic segmentation maps to form instance masks. The performance of top-down approaches is closely related to the choice of detectors [165], while bottom-up approaches depend on both semantic segmentation results and clustering methods [166]. Besides, several approaches [167] use grid representation to learn instance masks directly.

Query-based Object Detection and Segmentation: With the rise of vision transformers [3], [4], recent works mainly use transformer-based approaches in segmentation, detection, and video understanding. Compared with previous pixel-based approaches, transformer-based approaches have more advantages in cases of flexibility, simplicity, and uniformity [4], [57], [168], [169], [170], [171], [172], [173], [174], [175], [176], [177]. One representative work is detection transformer (DETR) [4]. It contains a CNN backbone, a standard transformer encoder, and a standard transformer decoder. It also introduces the concepts of object query to replace the anchor design in pixel-based approaches. Object query is usually combined with bipartite matching [178] during training, uniquely assigning predictions with ground truth. This means each object query builds the one-to-one matching during training. Such matching is based on the matching cost between ground truth and predictions. The matching cost is defined as the distance between prediction and ground truth, including labels, boxes, and masks. By minimizing the cost with the Hungarian algorithm [178], each object query is assigned by its corresponding ground truth. For object detection, each object query is trained with classification and box regression loss [156]. For instance-wise segmentation, each object query is trained with classification loss and segmentation loss. The output masks are obtained via the inner product between object query and decoder features. Recently, mask transformers [179], [180], [181] further removed the box

head for segmentation tasks. Max-Deeplab [179] is the first to remove the box head and design a pure-mask-based segmenter. It combines a CNN-transformer hybrid encoder [182] and a query-based decoder as an extra path. Max-Deeplab still needs many auxiliary loss functions. Later, K-Net [180] uses mask pooling to group the mask features and designs a dynamic convolution to update the corresponding query. Meanwhile, MaskFormer [181] extends the original DETR by removing the box head and transferring the object query into the mask query via MLPs. It proves simple mask classification can work well enough for all three segmentation tasks. Then, Mask2Former [183] proposes masked cross-attention and replaces the cross-attention in MaskFormer. Masked cross-attention forces the object query only attends to the object area, guided by the mask outputs from previous stages. Mask2Former also adopts a stronger Deformable FPN backbone [184], stronger data augmentation [185], and multi-scale mask decoding. In summary, query-based approaches are stronger and simpler. They cannot directly detect novel classes but are widely used as base detectors and segmenters in open vocabulary settings.

B. Vision Language Modeling

Large Scale Visual Language Pre-training: Better visual language pre-training can lead to a better understanding of semantics on given visual inputs [169], [170], [186]. Previous works [187], [188], [189] focus on cross-modality research via learning the visual and sentence-dense connection between different modalities. Some works [189], [190] utilize two-stream neural networks based on the vision transformer model, while several works [191], [192] adopt the single-stream neural network, where the text embeddings are frozen. The two-stream neural networks process visual and language information and fuse them afterward by another transformer module. Recently, most approaches [15], [16] adopt pure transformer-based visual language pre-training. Both CLIP [15] and Align [16] are concurrent works that explore extremely large-scale pre-training on image text pairs. In particular, with such large-scale training, CLIP demonstrates that the simple pre-training task of predicting which caption goes with which image can already lead to stronger generalizable models. Several following works [193], [194], [195] aim to improve the CLIP training via mask image modeling, scaling up the training data and model size. These VLMs are the foundation of open vocabulary learning in different tasks, which means that the open vocabulary approaches aim to distill or utilize knowledge of VLMs into their corresponding tasks.

Moreover, in addition to achieving better zero-shot recognition, several works focus on designing better vision language models for language-related tasks, including visual question answering (VQA). Several works [196], [197], [198], [199], [200] explore how to better align caption loss and contrastive loss during the image-text pertaining. In particular, CoCa [197] adopts cross attention to connect caption generation part and contrastive learning. Recently, BLIP-2 [200] bootstraps

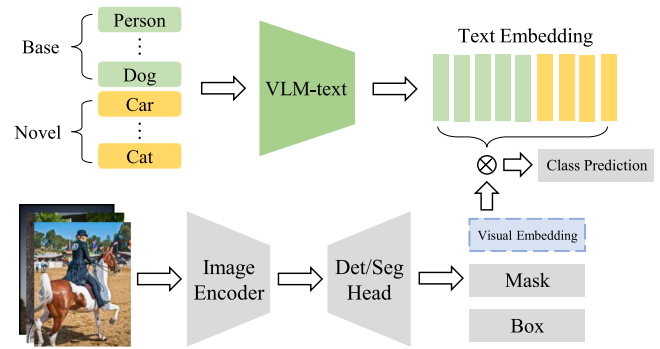


Fig. 5. An illustration of a common architecture in open vocabulary object detection and segmentation. The vision model predicts a class embedding for each box/mask. The embeddings are compared to a set of class embeddings generated by a VLM text model like CLIP or ALIGN, using dot products. The class with the highest score is chosen as the predicted label for the object. Note that while humans define the set of possible object classes, the system only has access to a limited set of “base” classes during training.

vision-language pre-training from off-the-shelf frozen pre-trained image encoders and frozen large language models via a lightweight Querying Transformer.

- *Visual Grounding Tasks:* Visual grounding tasks aim to localize the specific objects according to given text descriptions. Referring segmentation tasks [170], [201], [202], [203] segment the specific object. Referring expression comprehension [204], [205], [206] localizes the bounding boxes of given object texts. Recent research is mainly based on two-stream networks: one for visual encoder [1], [3] and the other for text encoder [9]. These works focus on how to better match text features and visual features via different architectures. For referring segmentation, previous works [186], [207], [208] adopt “decoder-fusion”, where they design a separate decoder at the end of the two-stream networks. Recent works [209], [210] explore the “encoder-fusion” to directly fuse the text features into a visual backbone before mask prediction. For referring expression comprehension, similar approaches [204], [211] are adopted such as encoder-fusion. Recently, several works [212], [213] aim to unify visual grounding and object detection. In particular, Grounding DINO [213] unifies detection and grounding in one framework. It uses stronger detectors with multiple dataset pre-training, which achieves strong results for many downstream tasks. Compared with open vocabulary learning, visual grounding tasks require visual text matching with specific text descriptions. Open vocabulary learning tasks require the model to automatically detect, segment and recognize new objects without the given text information, such as class names, which is more challenging.

Turning Close-set Detector and Segmenter Into Open Vocabulary Setting: A common way towards the open vocabulary setting is to replace the fixed classifier weights with the text embeddings from a VLM model. In Fig. 5, we present a meta-architecture. In particular, the vision model generates a visual embedding for each box/mask proposal and computes similarity scores by computing a dot product with the text embeddings

TABLE II
FEATURE SUMMARIZATION OF CURRENT OPEN VOCABULARY OBJECT DETECTION APPROACHES

Method	Common Features
Knowledge Distillation	Distil the aligned visual-text knowledge of VLM into object detection. Focus on distilling loss design and target proposal generation.
Region Text Pre-training	Use large-scale and easily available text-image pairs to pre-train a better and universal detector and finetune on the target datasets
Training with More Balanced Data	Adopt extra data or balanced datasets as an augmentation to train the classification head.
Prompting Modeling	Adopt different prompting modeling to better transfer the VLM knowledge into detector.
Region Text Alignment	Design a better region-text alignment method to better align the region visual features (RoIs) into text features of VLM.

from both base and novel classes. The classification scores are computed as follows:

$$p_{ij} = \frac{\exp \langle e_i^v, e_j^t \rangle}{1 + \sum_{j'=1}^{|C_B \cup C_N|} \exp \langle e_i^v, e_{j'}^t \rangle}, \quad (1)$$

where e_i^v is the i th vision embedding and e_j^t is the j th text embedding, and p_{ij} is the prediction score of the i th vision proposal predicted to the j th class. The $\langle \cdot, \cdot \rangle$ represents the dot product operation. The denominator is added with one because the background text embedding is set to all 0 or make it learnable. C_B and C_N are base and novel classes, respectively. Finally, the class with the highest prediction score is chosen as the prediction. In practice, the vision embeddings are trained with base annotations to fit the text embeddings. Therefore, the model combines the knowledge of VLM and the learned visual features, and the detector/segmenter can detect/segment novel classes via semantically related text embeddings.

C. Open Vocabulary Object Detection

In this section, we divide the methods into five categories: knowledge distillation, region text pre-training, training with more balanced data, prompting modeling, and region text alignment. Finally, we summarize the common features in Table II.

Knowledge Distillation: These techniques aim to distill the knowledge of Vision-and-Language Models (VLMs) into close-set detectors [2], [4], [156]. Since the knowledge of VLMs is much larger than close-set detectors, distilling novel classes into based classes trained detectors is a straightforward idea. Knowledge distillation aims to distill visual knowledge directly into close-set detectors since the visual features are aligned with text features during the VLM pre-training stage. One earlier method is ViLD [12], a two-stage detection approach that utilizes instance-level visual-to-visual knowledge distillation. ViLD consists of two branches: the ViLD-text branch and the ViLD-image branch. In the former branch, fixed text embeddings obtained from VLMs' text encoder output are treated as classifiers. Meanwhile, in the latter branch, pre-computed proposals are fed into a detector to obtain region embeddings using the RoIAlign [2] function. The cropped image is then sent to the VLMs' image encoder to generate image embeddings. Subsequently, ViLD proposes distilling this information onto each Region-of-Interest (RoI) via \mathcal{L}_1 Loss. LP-OVOD [222] extends the ViLD framework by making two main modifications. First, it replaces the softmax cross entropy loss with the sigmoid focal loss [158]. Second, LP-OVOD introduces a new classification branch that is supervised by pseudo labels. HierKD [223] uses

a single-stage detector and introduces a global-level language-to-visual knowledge distillation module. The module aims to narrow down performance gaps between one-stage and two-stage methods. This technique employs global-level knowledge distillation modules (GKD), which align global-level image representations with caption embeddings through contrastive loss. Both ViLD and HierKD use pixel-based detectors. However, Rasheed et al. [224] leverages query-based detector Deformable DETR [184] for its detection process. To ensure consistency between their detection region representations and CLIP's region representations, the authors employed inter-embedding relationship matching loss (IRM). Furthermore, the authors adopt mixed datasets pre-training to enhance the ability of novel class discovery. OADP [225] thinks previous works only distill object-level information from VLMs to downstream detectors and ignore the relation between different objects. To tackle this problem, OADP employs object-level distillation as well as global and block distillation methods. These supplementary techniques aim to compensate for the lack of relational information in object distillation by optimizing the L1 distance between the CLIP visual encoder and the detector backbone's global features or block features. Rather than using simple novel class names for text distillation, several works also utilize more fine-grained information, including attributes, captions, and relationships of objects. PCL [226] adopts an image captioning model to generate more comprehensive captions that describe object instances. OVRNet [227] simultaneously detects objects and their visual attributes in open vocabulary scenarios. By exploring the COCO attributes dataset [228] via a joint co-training strategy, the authors find that the recognition of fine-grained attributes works complementary for OVD. In summary, knowledge distillation is a common design, which effectively transfer VLM's knowledge into close set detectors. However, the recognition ability is still within the scope of teacher VLMs.

Region Text Pre-training: Another assumption of open vocabulary learning is the availability of large-scale image text pairs, which can be easily obtained in daily life. Since these pairs contain large enough knowledge to cover the most novel or unseen datasets for detection and segmentation. Most approaches adopt web-scale caption data for pre-training, which contains millions of image text pairs. The learning of region text alignment maps the novel classes of visual features and text features into an aligned feature space. Once trained for the alignment, it is nature to generalize the detector for novel class classification. OVR-CNN [11] first introduces the concept of open vocabulary object detection by using caption data for novel class detection. The model first trains a ResNet [1] and vision to language (V2L) layer using image-caption pairs via grounding,

masked language modeling, and image-text matching. Since captions are not constrained in language space, the V2L layer learns to map features from visual space into semantic space without limiting to closed-label space. During the next stage of training, Faster R-CNN [156] is used as a detection method with pre-trained ResNet as its backbone. OVR-CNN [11] replaces only learnable classifiers with fixed text embeddings from pre-trained language models [9]. For classification purposes, region visual features obtained from RoI-Align [2] are sent into the V2L layer and mapped into semantic space. Attribute-Sensitive OVR-CNN [229] proposes a different approach from OVR-CNN [11]. Instead of grounding vision regions to the input word embeddings of BERT [9], Attribute-Sensitive OVR-CNN aligns vision regions with contextualized word embeddings that are output from BERT [9]. Additionally, Attribute-Sensitive OVR-CNN proposes using an adjective-noun negative caption sampling strategy to enhance the model’s sensitivity to adjectives, verb phrases, and prepositional phrases other than object nouns in the caption. GLIP series [230], [231] unify the object detection and phrase grounding for pre-training. In particular, it leverages massive image-text pairs by generating grounding boxes in a self-training fashion, which sets strong results for both detection and grounding. RegionCLIP [215] learns visual region representation by matching image regions to region-level descriptions. It creates pseudo labels by CLIP for region-text pairs and then uses contrastive loss to match them before fine-tuning the visual encoder using human-annotated detection datasets. OWL-ViT [232] removes the final token pooling layer of pre-trained VLMs’ image encoder and attaches a lightweight classification head and box regression head to each transformer output token before fine-tuning it on standard detection datasets using bipartite matching loss. Meanwhile, MaMMUT [219] presents a simple text decoder and visual encoder for multimodal pre-training. Designing a two-pass text decoder combines both contrastive and generative learning in one framework. The former is for grounding text visual entities, while the latter learns to generate. DITO [233] presents a new image-level pretraining strategy to bridge the gap between image-level pretraining and open-vocabulary object detection. At the pertaining phase, DITO replaces the classification architecture used in CLIP with the detector architecture, which better serves the region-level recognition needs of detection by enabling the detector heads to learn from noisy image-text pairs. In summary, adopting more text-image pair can improve the performance on rare and novel classes. However, such process needs more computation cost for extra dataset training.

Training with More Balanced Data: Rare and unseen data are common in image classification datasets. Joint training can be used to address this issue. The core idea of these approaches is to leverage more balanced data, including image classification datasets, pseudo labels from image-text data, extra-related detection data, or even data generated by generation models. Detic [133] improves long-tail detection performance with image-level supervision. The classification head of Detic is trained using image-level data from ImageNet21K [234]. During training, the max area proposal from RPN is chosen as the RoI. Detic’s basic insight is that most classification data are object-centric.

Therefore, the maximum area proposal may completely cover an object represented by an image-level class. The mm-ovod [235] method improves upon Detic [133] by utilizing multi-modal text embeddings as the classifier. This approach employs a large language model to create a description of each class to generate text-based embeddings. In addition, mm-ovod also uses vision-based embeddings from image exemplars. By fusing the text-based and vision-based embeddings, the multi-modal text embeddings can significantly enhance Detic’s performance. To use more data, several methods generate pseudo bounding box annotations from large-scale image-caption pairs. PB-OVD [118] uses Grad-CAM [236] and generates the pseudo-bounding boxes combined with RPN’s proposals. The activation map of Grad-CAM is obtained from the alignment between region embeddings and word embeddings that come from pre-trained VLMs. Then, the boxes are generated from the activation map and jointly trained with existing box annotations. Meanwhile, several works leverage the rich semantics available in recent vision and language models to localize and classify objects in unlabeled images. VL-PLM [214] proposes to train Faster R-CNN as a two-stage class-agnostic proposal generator using a detection dataset without category information. LocOV [237] uses class-agnostic proposals in RPN to train Faster R-CNN by matching the region features and word embeddings from image and caption, respectively. From the data generation view, several works [137], [238], [239] adopt the diffusion model to generate the on-target data for effective training. In particular, X-Paste [238] generates the rare class data to improve the classification ability for existing approaches. It uses an extra segmentation model to provide the foreground object masks and adopt simple copy and paste [240] to augment training data. Recently, several works [134], [135], [241] have explored the self-training approaches to generate large pseudo labels for more balanced learning. OWLv2 [134] presents a self-training pipeline. It generates huge pseudo-box annotations on WebLI [242] dataset and pre-trains a model on such generated dataset. Finally, it fine-tunes this model on a specific OVD dataset. Since the WebLI contains a large vocabulary size, OWLv2 achieves significant gains on LVIS and COCO datasets. In summary, these approaches are more effective in unbalanced data setting. Designing a more efficient data augmentation still have room to explore.

Prompting Modeling: Prompt modeling is an effective technique for adapting foundation models to various domains, such as language modeling [243] and image classification [244], [245]. By incorporating learned prompts into the foundation model, the model can transfer its knowledge to downstream tasks more easily. To generate text embeddings of category names, prompts are fed to the text encoder of pre-trained VLMs. However, negative proposals do not belong to any specific category. To address this issue, DetPro [216] forces the negative proposal to be equally dissimilar to any object class instead of using a background class. PromptDet [44] introduces category descriptions into the prompt and explores the position of the category in the prompt. It also proposes to use cached web data to enhance the novel classes during training. Based on the DETR [4] framework, CORA [218] proposes region prompting and anchor pre-matching. The former reduces the gap between

the whole image and region distributions by prompting the region features of the CLIP-based region classifier, while the latter learns generalizable object localization via a class-aware matching mechanism. Prompt-OVD [246] follows the pipeline of OV-DETR [217]. It presents RoI-based masked attention and RoI pruning techniques by utilizing CLIP visual features to improve the novel object classification. Despite the effectiveness, without more data training, the performance of prompting is limited compared with other directions.

Region Text Alignment: Using language as supervision instead of a ground truth bounding box is an attractive alternative for open vocabulary object detection. However, obtaining enough object-language annotations is difficult and costly. Compared with region text pre-training, region text alignment aims at a better matching between region visual features and text features during the base class training without introducing extra data. OV-DETR [217] introduces a transformer-based detector for open vocabulary object detection by replacing the bipartite matching method with a conditional binary matching mechanism. VLDet [43] converts the image into a set of regions and the caption into a set of words, and solves the object-language alignment problem using a set matching method. It proposes a simple matching strategy to align caption and vision features. DetCLIPv2 [247] uses ATSS [248] as an object detector and trains it with three datasets: a standard detection dataset, a grounding dataset, and an image-text pairs dataset for word-text alignment. Recently, BARON [112] proposes to align the embedding in bags of different regions rather than only individual regions. It first groups contextually interrelated regions as a bag and treats each region in the bag as a word in a sentence. Then, it sends the bag of regions into the text encoder to get bag-of-regions embeddings. These bag-of-regions embeddings will align with cropped region embeddings from the image encoder of VLMs. CoDet [249] reformulates the region-word alignment as a co-occurring object discovery problem and aligns the co-occurring objects with the shared concept. F-VLM [136] finds that the origin CLIP features already have grouping effects. It is a two-branch method similar to ViLD-text [12]. F-VLM uses a CLIP vision encoder as the backbone and applies the VLM feature pooler on the region features from the backbone to get VLM predictions. The final result of F-VLM combines the detection scores and the VLM predictions. RO-ViT [250] builds upon F-VLM [136], which believes that the difference in position embeddings between image-level and region-level is responsible for the gap between vision-language pre-training and open vocabulary object detection. To address this issue, RO-ViT proposes a cropped positional embedding module in VLM that bridges the gap between vision-language pre-training and downstream open vocabulary object detection tasks. In summary, region-text alignment is a core research topic, and it still has room to explore when considering recent large language models [42].

D. Open Vocabulary Segmentation

Although segmentation problems can be defined in various ways, such as semantic segmentation or panoptic segmentation, we categorize the methods based on their technical aspects.

Utilizing VLMs to Leverage Recognition Capabilities: VLMs have shown remarkable performance in image classification by learning rich visual and linguistic representations. Therefore, it is natural to extend VLMs to semantic segmentation, which can be seen as a dense classification task. For instance, LSeg [14] aligns the text embeddings of category labels from a VLM language encoder with the dense embeddings of the input image. This enables LSeg to leverage the generalization ability of VLMs and segment objects that are not predefined but depend on the input texts. Following LSeg, several works propose methods to utilize the VLMs for open vocabulary segmentation tasks. Fusioner [251] utilizes self-attention operations to combine visual and language features in a transformer-based framework at an early stage. ZegFormer [113] decouples the problem into a class-agnostic segmentation task and a mask classification task. It uses the label embeddings from a VLM to classify the proposal masks and applies the CLIP-vision encoder to obtain language-aligned visual features for them. After that, more approaches are proposed to extract the rich knowledge in VLMs. SAN [252] attaches a lightweight side network to the pre-trained VLM to predict mask proposals and classification outputs. CAT-Seg [253] jointly aggregates the image and text embeddings of CLIP by fine-tuning the image encoder. Han et al. [254] develop an efficient framework that does not rely on the extra computational burden of the CLIP model. OP-SNet [114] proposes several modulation modules to enhance the information exchange between the segmentation model and the VLMs. The modulation modules fuse the knowledge of the learned vision model and VLM, which improves the zero-shot performance in novel classes. MaskCLIP [47] inserts Relative Mask Attention (RMA) modules into a pre-trained CLIP model. It can utilize the CLIP features more efficiently. TagCLIP [255] proposes a trusty token module to explicitly predict pixels that contain objects (both base and novel) before classifying each pixel, which avoids the problem that models tend to misidentify pixels into novel categories. Recently, several works [256], [257], [258] directly fuse the frozen CLIP visual encoder. They fuse the learned CLIP score and prediction score to achieve better close-set and open-set recognition ability trade-offs. Same as the open vocabulary detection, there are still several improve space when adopting more advanced VLMs.

Learning from Caption Data: Image caption data can provide weak supervision for identifying novel classes, as they expose potential novel category names. Like such approaches in open vocabulary object detection [11], image caption data is also explored in open vocabulary segmentation tasks also explore image caption data. OpenSeg [45] applies a region-word grounding loss to directly ground objects and nouns in the caption data. CGG [46] combines caption grounding and caption generation losses to fully exploit the knowledge in caption data. It leverages the role of object nouns in visual grounding and the mutual benefits of words in caption generation.

Generating Pseudo Labels: Intuitively, providing the model with more data on novel categories can improve classification performance. MaskCLIP+ [119] modifies the CLIP-vision model by replacing the last pooling layer with a convolution layer, which produces dense feature maps. These feature maps are then used to generate pseudo labels for training a

TABLE III
FEATURE SUMMARIZATION OF CURRENT OPEN VOCABULARY SEGMENTATION APPROACHES

Method	Common Feature
Utilizing VLMs to Leverage Recognition Capabilities.	Design fusion or alignment methods to better fuse VLM knowledge into the existing segmenters.
Learning from Caption Data.	Use extra caption data to ground the objects' visual features with caption data.
Generating Pseudo Labels.	Fully explore the potential of the VLM model to generate better masks to train the segmentation models
Training without Pixel-Level Annotations.	Combine the VLM and different unsupervised approaches to perform unsupervised mask generation or segmentation.
Jointly Learning Several Tasks.	Pre-train on large-scale and easily available text-region pairs datasets or detection using a unified detection and segmentation model.
Adopting Denoising Diffusion Models.	Explore the feature representation of text-to-image diffusion models or utilize the generation ability to augment masks.

OV-VC refers to open vocabulary video classification. OV-3D refers to open vocabulary 3D recognition.

segmentation model. OVSeg [120] matches the proposed image regions with nouns in captions using CLIP to generate pseudo annotations. It also proposes a mask prompt tuning module to help CLIP adapt to masked images without changing their weights. XPM [121] follows a similar pseudo data generation procedure by aligning words in captions with regions in images to generate pseudo instance mask labels. However, these approaches are limited by the data annotations, since both mask and region caption are hard to collection. Maybe more automatic data generation [41], [131] can be used in the future.

Training without Pixel-Level Annotations: Despite the absence of annotations on novel objects, most methods still require base mask annotations during training, which are costly and labor-intensive to obtain. To reduce the annotation burden, many recent works explore training segmentation models with merely weak supervision, such as image caption. GroupViT [124] proposes a semantic segmentation framework that leverages a grouping mechanism to automatically merge image patches with the same semantics. It trains the model with a contrastive loss on image-caption pairs and does not need pixel-level annotations. PACL [259] enhances the contrastive loss with a patch alignment objective that aligns the image patches and the CLS token of the captions. ViL-Seg [220] combines both contrastive loss and clustering loss. SegCLIP [125] further introduces a reconstruction loss and a superpixel-based KL loss. To compute the superpixel-based KL loss, it uses an unsupervised graph-based segmentation model [260] to generate pseudo labels. TCL [261] proposes a finer-grained contrastive loss, namely text-grounded contrastive loss, which explicitly aligns captions and regions. The model can segment the region that corresponds to a given text expression during inference. OVSegmentor [221] introduces masked entity completion and cross-image mask constituency tasks to improve the training efficiency. Mask-free OVIS [126] generates pseudo mask annotations with purely image-text pair data to assist the segmentation models. Since these methods are trained without mask supervision, the quality of segmentation results is low which make it hard to use in real application.

Jointly Learning Several Tasks: Open vocabulary segmentation encompasses different tasks, including OVSS, OVIS, and OVPS. How to jointly learn multiple segmentation tasks in one model becomes a practical problem. X-Decoder [129] proposes a framework that can handle various tasks, including open vocabulary semantic segmentation, open vocabulary instance segmentation, and open vocabulary panoptic segmentation. It utilized a query-based segmentation architecture, pre-trains the model on a mixture of segmentation data and image-text pairs, and then fine-tuned or applied in zero-shot settings for downstream tasks.

FreeSeg [130] also proposes a generic framework to tackle the three tasks in a unified manner. It designs an adaptive task prompt module and performs test time tuning on the learnable prompts to capture the task-specific features. POMP [262] first trains class prompts at a large vocabulary dataset (Imagenet-21 K) and then transfers the prompts into multiple open vocabulary tasks. Moreover, if the model can learn from multiple tasks, it raises the question of whether it can benefit from multi-sourced data, e.g., detection and segmentation data. To handle this question, OpenSeeD [127] and OpenSD [128] jointly learn from segmentation and detection datasets. To bridge the task gap, they propose decoupled decoding frameworks, which decode foreground and background masks separately and generate masks for the bounding box proposals. One shortcoming of these approaches is extra computation costs that brought other tasks.

Adopting Denoising Diffusion Models: Recently, diffusion-based generative models [263] have achieved remarkable success in text-based image generation suggests. There are mainly two ways to let diffusion models enhance open vocabulary tasks. First, The reality and diversity of the images generated by diffusion models suggest that the intermediate representations in the diffusion models may be highly aligned with natural language vocabularies. Inspired by this, ODISE [132] proposes a framework that leverages the middle representation of the diffusion model. The middle representations contain rich semantic information for both base and novel classes. Thus, ODISE can perform open vocabulary segmentation by training a decoder head on the representations. In [264] follows ODISE to use the intermediate representations of diffusion models. Another way to incorporate diffusion models into the open vocabulary setting is to take advantage of the image and mask generation ability [131], [138], [265]. OVDiff [131] proposes a prototype-based method to tackle open vocabulary semantic segmentation. It uses the diffusion model to generate images for various categories and treat them as prototypes. The method does not need training. During testing, the input images are compared with these generated prototypes, and the best-matching one is the predicted class. Meanwhile, several works [138], [265] use the diffusion model to generate images and masks for the rare classes to augment data. Thus, the segmenter can be trained in a more balanced manner or even totally using generated data. Despite the inference time of these methods is quite limited, it still has room to explore joint generation and segmentation in one framework.

E. Open Vocabulary Video Understanding

We also review several video open vocabulary tasks. Most works focus on designing VLM models' temporal fusion or

association in various settings, including action recognition and tracking.

Video Classification: Traditional video classification methods usually require large datasets specific to video (e.g., Kinetics [278]). However, annotating video datasets requires very high costs. Using semantic information of label texts from web data may alleviate this deficiency. Building on the success of CLIP in image open vocabulary recognition, ActionCLIP [266] uses knowledge from image-based vision-language pre-training. It adds a temporal fusion layer for the zero-shot capability in video action recognition. A concurrent work I-VL [267] follows a similar paradigm by training the transformer layer on top of a frozen CLIP image encoder and also has a strong zero-shot capability. Another concurrent work, EVL [279], also employs the frozen CLIP for efficient video action recognition. X-CLIP [280] uses a video encoder that leverages the temporal information in the encoder for video recognition. It aligns the feature from the video encoder with the text encoder that is pre-trained on the image-based vision-language pairs (e.g., CLIP). To take advantage of other modalities, MOV [281] further fuses the audio information and the pre-trained CLIP model to build a multi-modal open vocabulary video classification model. Recently, Open-VCLIP [282] formulates the CLIP-to-video knowledge transfer as a continual learning problem and proposes *Interpolated Weight Optimization* to address the issue. AIM [283] tackles the open vocabulary video classification problem by adding adapt layers on top of the CLIP image encoder. ViFi-CLIP [268] reveals that a simple fine-tuning baseline (image-level feature extraction with CLIP visual encoder following temporal pooling) instead of advanced fusion layers can have a strong performance and further investigates the influence of prompt tuning. ASU [284] explores the use of fine-grained language features extracted by semantic units (e.g., head, arms, balloon, knee bend posture of exercise in the video) to guide the training of video classification. Recent advancements in open vocabulary video classification have turned our attention to the language part of VLM modeling. VicTR [285] introduces *video-conditioned text representations* to optimize the visual and text information jointly. MAXI [122] leverages LLMs to build a text bag for video without annotation by text expansion. The verbs, which are lacking in the realm of image, have the opportunity to participate in video-language modeling.

Object Tracking and Video Instance Segmentation: The object tracking and video instance segmentation can also enjoy the rich knowledge of VLMs to build an open vocabulary tracker based on a close-set tracker. Going beyond the large vocabulary object tracking [286], to tackle the real-world multiple object tracking (MOT), OVTrack [269] first introduces large VLMs to the object tracking and tackles their proposed open vocabulary MOT (OV-MOT) task. Specifically, OVTrack extracts RoIs via an RPN and uses CLIP for knowledge distillation. A separate tracking head is used for tracking and supervised by pseudo-LVIS videos. As for video instance segmentation (VIS), to make the VIS model capable of generalizing to novel classes in the real world, MindVLT [56] adopts a frozen CLIP backbone and proposes an end-to-end method with an open vocabulary classifier for segmenting and tracking unseen categories. It collects

a large-vocabulary VIS dataset and tests their MindVLT on it. A concurrent work, OpenVIS [270], also tackles the open vocabulary video instance segmentation task but in a different manner. It generates the class-agnostic mask of instances and leverages the masks to crop raw images to feed into the CLIP visual encoder for calculating class scores. Beyond the open vocabulary video instance segmentation, DVIS++ [287] proposes the first open vocabulary universal video segmentation scheme supporting video semantic segmentation, video instance segmentation, and video panoptic segmentation. DVIS++ [287] leverages the knowledge from the CLIP backbone and adopts mask pooling on the CLIP backbone for enabling open vocabulary capability.

F. Open Vocabulary 3D Understanding

In this section, we review the open vocabulary approaches used in 3D scene understanding tasks. We mainly survey the works for point cloud classification and segmentation, which explores the knowledge of 2D VLMs into 3D.

3D Recognition: The VLM has shown great success in the zero-shot or few-shot learning of 2D images by leveraging a huge amount of image data. However, such internet-scale data is not available for the 3D point cloud. To this extent, extending the open vocabulary mechanism into 3D perception is not trivial since it is hard to collect enough data for point-language contrastive training like what CLIP does in 2D perception. To mitigate the gap, PointCLIP [271] takes the first step to 3D open vocabulary perception. The principle insight behind PointCLIP is that the 3D point cloud can be converted to CLIP-recognizable images. By projecting the 3D point cloud to a 2D plane and extracting visual features from the projected depth map via the CLIP visual encoder, the point cloud feature can be naturally aligned with the language feature extracted by the language encoder. Then, the whole framework is with 3D point cloud zero-shot capability, just like 2D images. However, simply projecting the point cloud to depth maps may yield inferior performance. To further improve the performance, rather than directly using the CLIP visual encoder for visual feature extraction of depth map, CLIP2Point [288] aligns the RGB image feature from the CLIP visual encoder and depth feature from a depth encoder through contrastive learning. Specifically, they collect image-depth pairs to align the from-scratch depth map encoder with the CLIP image encoder at training time and only use the depth map encoder at inference time. Then, the depth feature can be aligned with the language embedding, facilitating the 3D point cloud zero-shot capability with CLIP. Also aiming at improving the performance of PointCLIP [271], PointCLIPV2 [272] focuses on different perspectives and proposes a realistic shape projection scheme along with an LLM-based 3D prompting generation scheme. Although without extra annotations, PointCLIPV2 achieves a significant improvement over PointCLIP [271].

Though projecting the point cloud to depth maps can make it easy to leverage the existing 2D visual encoders pre-trained in the CLIP, it may not exhaustively use the 3D information in the point cloud. Accordingly, beyond the depth map-based methods, ULIP [116] first collects multi-modalities triplets (point cloud,

TABLE IV
 REPRESENTATIVE WORKS SUMMARIZATION AND COMPARISON IN SECTIONS III-E, III-F, AND III-G

Method	Task	Vision Training Annotations	Text Model	Vision Model	Highlight
Open Vocabulary Video Understanding (Sec. 3.5)					
ActionCLIP [268]	OV-VC	None	CLIP-text	CLIP-vision + Temporal Pooling	The first work proposes to use CLIP for video understanding. I-VL provides a simple baseline for adapting CLIP to video understanding. ViFi-CLIP indicates that a simple fine-tuning baseline is strong enough. OVTrack is the first open vocabulary multi-object tracking method. MindVLT is the first open vocabulary video instance segmentation method. OpenVIS conduct open vocabulary VIS in a two-stage manner.
I-VL [269]	OV-VC	Video Class Label (base)	CLIP-text	CLIP-vision + transformer (temporal fusion)	
ViFi-CLIP [270]	OV-VC	Video Class Label (base)	CLIP-text	CLIP-vision (learnable) + Temporal Pooling	
OVTrack [271]	OV-Tracking	Image Bounding Boxes	CLIP-text	CLIP-vision + Faster R-CNN	
MindVLT [57]	OV-VIS	Video Segmentation Masks (base)	CLIP-text	CLIP-vision + Mask2Former + SORT	
OpenVIS [272]	OV-VIS	Video Segmentation Masks (base)	CLIP-text	CLIP-vision + Mask2Former	
Open Vocabulary 3D Scene Understanding (Sec. 3.6)					
PointCLIP [273]	OV-3D	None	CLIP-text	CLIP-vision (projected depth map)	PointCLIP uses projected depth for OV 3D recognition. PointCLIPV2 generates prompts with LLMs to enhance performance.
PointCLIPV2 [274]	OV-3D	None	CLIP-text + GPT-3	CLIP-vision (projected depth map)	
ULIP [117]	OV-3D	Point Cloud Label	CLIP-text	Point Cloud Encoder + CLIP-vision	ULIP aligns the 3D encoder and 2D CLIP encoder feature by distillation. CLIP2 leverages real-world point cloud data and CLIP to train the 3D encoder for 3D open vocabulary segmentation.
CLIP2 [275]	OV-3D	None	CLIP-text	Point Cloud Encoder + CLIP-vision	
OpenShape [276]	OV-3D	Point Cloud Label	CLIP-text + GPT-4	Point Cloud Encoder + CLIP-vision	OpenShape builds a large-scale dataset that provides image, point cloud, and text triplets. OV-3DEITC exploits information from two modalities to achieve 3D open vocabulary object detection.
OV-3DEITC [277]	OV-3D-OD	Pseudo Labels from 2D Detector	CLIP-text	3DETR [278]	
PLA [124]	OV-3D-SS/IS	3D segmentation masks (base)	CLIP-text	sparse-conv UNet [279]	PLA first tackles the 3D open vocabulary scene understanding problem. OpenScene train a 3D Encoder yielding dense features co-embedded with text and image pixels for open vocabulary semantic segmentation.
OpenScene [116]	OV-3D-SS	None	CLIP-text	3D Encoder + LSeg [14]	

OV-VC refers to open vocabulary video classification. OV-3D refers to open vocabulary 3D recognition.

image, and text) for training the 3D backbone. Since CLIP already aligns the language and image encoders, it only needs to align the 3D backbone to the image-language feature space. With only small-scale data, ULIP [116] aligns the 3D feature extracted by the 3D backbone to the CLIP-aligned visual and text feature to enable the zero-shot capability and enhance the standard 3D recognition capability. The advantage of ULIP [116] also includes unifying the three modalities, which may bring more downstream applications, such as image-to-3D retrieval. Though ULIP [116] aligns the 3D and 2D and thus enables the native 3D open vocabulary capability, it trains on a small-scale dataset, which limits its performance. CLIP2 [273] resorts to finding training samples from the real world and proposes *Triplet Proxies Collection* scheme for finding instance-level 3D point cloud, 2D image crop, and text triplets for 3D recognition. OpenShape [274] scales up both the dataset and backbone. For the dataset, they build a text-3D shape dataset containing 876 k training shapes with over 1 k categories. They also explore adopting larger 3D backbones and their performance. As a result, OpenShape [274] drastically enhances the performance on 3D zero-shot capability. LidarCLIP [289] also aligns the 3D and 2D features with a similar approach, but LidarCLIP [289] focuses on the driving scenes. Besides the 3D zero-shot capability, LidarCLIP [289] also shows point cloud captioning and lidar-to-image generation capability with off-the-shelf foundation models.

3D Object Detection: In 3D object detection, the training data is often hard to get, and existing 3D detectors are often trained on very limited classes, which means they are hard to generalize to novel classes. Inspired by the success of 2D open vocabulary detection [133], aiming at taking full use of both image and point cloud modalities for 3D detection, OV-3DEITC [275] and OV-3DET [290] proposes to decouple the localization and recognition in point cloud object detection into localization and recognition. For localization, 2D pre-trained detectors can be adapted to train the 3D detector by back-projecting the 2D bounding box and further optimizing it by the point cloud. For recognition, they propose aligning the regional features of 3D and 2D encoders with the corresponding text feature. Therefore, during the inference, 3DEITC [275] and OV-3DET [290] can generalize to unseen classes with the knowledge in the CLIP.

While OV-3DET [290] achieves significant performance on the Open-Vocabulary 3D detection, it relies on the localization capability of pre-trained 2D detectors, which may limit the 3D object discovery capability. To solve this problem, CODA [291] proposes to discover objects with both 2D semantic prior and 3D geometry prior, and further proposes a cross-modal alignment module to align 3D and 2D features. Also trying to get rid of the 2D image detector, concurrent work Object2Scene [292] proposes to augment the existing 3D object detection datasets with large-scale 3D object datasets. It also introduces a new framework, L3Det [292], for 3D object-text alignment. As a result, both CODA [291] and L3Det [292] outperforms OV-3DET [290]. As driving scenarios may yield new challenges, the above-mentioned methods may have an inferior performance. OpenSight [293] focuses on the outdoor scenes and proposes to leverage the knowledge from 2D grounding DINO [213] to train the 3D detectors. In particular, it uses the size prior (e.g., cars are 1.96 m wide) generated by the LLM to re-calibrate the output of the model to make the model more suitable for the driving scene.

3D Scene Understanding: In the understanding of the 3D scene, the same obstacle, lacking training data, also limits the capability of generalization. Sharing a similar motivation that there is a lack of 3D-text pairs for point-language contrastive training, PLA [123] proposes to extract features from multi-view images sampled from a scene to generate descriptions with pre-trained language models. The descriptions can then be used to extract language features to train 3D backbone-language alignment. With the language-aligned 3D backbone, it is natural to conduct open vocabulary segmentation like with the 2D image. OpenScene [115] proposes to link the 3D features for each point with CLIP features for each pixel. By back-projecting the 2D pixels to 3D space, each point in the point cloud can be ensembled with features of several pixels in different views. The 3D-text co-embedding makes it possible for open vocabulary semantic segmentation. Inspired by MaskCLIP [119], a concurrent work CLIP-FO3D [294] also focuses on 3D semantic segmentation and has a similar approach. It leverages the feature map extracted by the CLIP visual encoder and directly transfers CLIP’s knowledge without any extra annotation. A

recent work, RegionPLC [295] on 3D open vocabulary semantic segmentation, considers the benefits of both PLA [123] and OpenScene [115]. It proposes to use the dense caption of any random region to extract more information in CLIP for 3D backbone distillation. Fine-grained dense supervision in RegionPLC [295] beyond view-level or instance-level further improves performance compared to PLA [123]. Different from training 3D backbones like in [115], [123], [294], [295], PartSLIP [296] and SATR [297] directly projecting the 3D point cloud and mesh to 2D plane and uses GLIP [230] to localize and segment. Meanwhile, OpenMask3D [298] mainly focuses on the 3D instance segmentation. It first generates class-agnostic instance mask proposals with the 3D point cloud. Then, OpenMask3D [298] selects views and projects 3D masks into 2D images. The 2D masks are further refined by SAM [41]. Finally, the 2D masks can be fed into the CLIP visual encoder to generate label prediction with the help of the CLIP language encoder. While OpenMask3D [298] has a strong performance, it requires point cloud-2D image pairs for training. On the contrary, OpenIns3D [299] does not require 2D images but only needs point clouds with colors. OpenIns3D [299] proposes a "Mask-Snap-Lookup" pipeline that first learns class-agnostic mask proposals, then generates synthetic scene-level images, and finally assigns categories for each proposal. Another work trying to improve OpenMask3D [298] is Open3DIS [300]. Open3DIS [300] improves the 3D mask proposal quality by employing a 3D instance network and a 2D-guide-3D Instance Proposal Module. After the mask proposal, it aggregates CLIP features in a multi-scale, multi-view manner. Open3DIS [300] achieves a stronger performance over OpenMask3D [298] on several datasets. In the driving scene, to mitigate the heavy dependence on the point cloud data annotation and fast generalization to the new scenes, CLIP2Scene [301] proposes to train a 3D network via semantic-driven cross-modal contrastive learning for 3D segmentation. CLIP2Scene [301] also considers spatial-temporal semantic consistency regularization to facilitate training.

G. Closely Related Tasks

- *Class Agnostic Detection and Segmentation:* The goal of class agnostic detection and segmentation is to learn a general region proposal system that can be used in different scenes. For detection, OLN [302] replaces the binary classification head in RPN by predicting the IoU score of foreground objects, which proves the generalization ability on cross-dataset testing. For example, the model trained on COCO shows the localization ability on the LVIS dataset. Open world instance segmentation [303] aims to correctly detect and segment instances whether their categories are in training taxonomy. GGN [304] leverages the bottom-up grouping idea, combining a local pixel affinity measure with instance-level mask supervision, producing a training regimen designed to make the model more generalizable. Meanwhile, from the data augmentation view, MViT [305] combines the Deformable DETR and CLIP text encoder for the open world class-agnostic detection, where the authors

build a large dataset by mixing existing detection datasets. For the segmentation domain, Entity segmentation [306], [307] aims to segment all visual entities without predicting their semantic labels. The goal of such segmentation is to obtain high-quality and generalized segmentation results. Fined-grained entity [307] presents a two-view image crop ensemble method named CropFormer to enhance the fine-grained details. Recently, SAM [41] proposes more generalized prompting methods, including masks, points, boxes, and texts. In particular, the authors build a larger dataset with 1 billion masks following the spirits of CLIP. The SAM achieves zero-shot testing in various segmentation datasets, which are highly generalizable.

- *Open World Object Detection:* Open world recognition [26] requires the model to identify novel categories and label them as "unknown." Then, the novel categories are progressively annotated. The model learns incrementally with new data and recognizes the newly-annotated categories. Inspired by this setting, open world object detection (OWOD) [31] expands the recognition task to object detection. The authors propose a novel method that utilizes contrastive clustering and an energy-based unknown identification module. Following them, OW-DETR [308] introduces an end-to-end transformer-based framework. It consists of three dedicated components, namely, attention-driven pseudo-labeling, novelty classification, and objectness scoring, to address the OWOD challenge explicitly. Meanwhile, open world DETR [309] proposes a two-stage training approach based on Deformable DETR. It focuses on alleviating catastrophic forgetting when the annotations of the unknown classes become available incrementally using knowledge distillation and exemplar replay technologies. PROB [310] proposes a probabilistic objectness head into the OW-DETR to better mine unknown background classes. Recently, researchers [311] propose unknown-classified open world object detection (UC-OWOD), which aims to detect and classify unknown instances into different unknown classes. This task is close to the open vocabulary object detection but requires incremental learning.
- *Open-Set Panoptic Segmentation:* Similar to other open-set tasks, open-set panoptic segmentation (OSPS) [312] requires the model to identify novel categories as 'unknown' in a panoptic segmentation task. The 'un' classes are chosen from the thing classes (foreground objects). The authors apply exemplar theory for novel class discovery. After that, Dual [313] proposes a divide-and-conquer scheme to develop a dual decision process for OSPS. The results indicate that by properly combining a known class discriminator with an additional class-agnostic object prediction head, the OSPS performance can be significantly improved.

IV. CHALLENGES AND OUTLOOK

A. Challenges

Base Classes Over-fitting Issues: Most approaches detect and segment novel objects by learning proposals from base class

annotations. Thus, there are natural gaps in shapes and semantic information between novel and base objects. VLM models can bridge such gaps via pre-trained visual-text knowledge. However, most detectors still easily overfit the base classes when the novel classes have similar shapes and semantics, since these classes are trained with higher confidence scores. More fine-grained feature discriminative modeling [314], [315], [316], including parts or attributes, is required to handle these problems.

Training Costs: Most state-of-the-art methods need *huge* data for pre-training to achieve good performances. However, the costs are expensive and even unavailable for many research groups to follow. Thus, with the aid of VLM, designing more efficient data learning pipelines or learning methods is more practical and affordable. One simple solution can be adopting a frozen backbone. However, this may limit the representation capacity.

Across Dataset Generation and Evaluation: As shown in the benchmark section, current state-of-the-art methods design specific models for each benchmark. Designing one shared model [213], [257], [317], [318], [319] across different datasets on open vocabulary detection and segmentation follows the origin spirits of open vocabulary learning. However, there are still performance gaps between unified models and dataset-specific models in several OVD benchmarks, such as LVIS datasets.

Better Benchmarks and Metrics: Since several classes contain overlapping concepts (for example, building and tower, person and woman), designing more new metrics [320] is also needed to better measure the open vocabulary methods. Moreover, current datasets are still small. To realize real open vocabulary settings, more datasets like SAM-1B [41] are needed.

B. Future Work

Explore Temporal Information: In practical applications, video data is readily available and used more frequently. Accurately segmenting and tracking objects that are not predetermined requires a great deal of attention, which is necessary for a wide range of real-world scenarios, such as short video clips and autonomous vehicles. However, there are only a few works [269] exploring open vocabulary learning on detection and tracking in video. Moreover, the input scenes [56] are simple. For example, several clips only contain a few instances, which makes the current tracking solution more trivial. Thus, a more dynamic, challenging video dataset is needed to fully explore the potential of vision language models for open vocabulary learning.

3D Open Vocabulary Scene Understanding: Compared with image and video, point cloud data are more expensive to annotate, in particular for dense prediction tasks. Thus, research on 3D open vocabulary scene understanding is more urgent. Current solutions for 3D open vocabulary scene understanding focus on designing projection functions for better usage of 2D VLMs. More new solutions for aligning 2D models knowledge into 3D models will be a future direction.

Explore Foundation Models With Specific Adapter For Custom Tasks: Vision foundation models [15], [41] can achieve good zero-shot performance on several standard classification and segmentation datasets. However, for several custom tasks,

such as medical image analysis, and aerial images, there are still many corner cases. Thus, designing a task-specific adapter is needed for these custom tasks. Such adapters can fully utilize the knowledge of pre-trained foundation models. One possible solution is to explore the in-context learning [321], [322] to fully explore or connect the knowledge of VLMs and LLMs.

Combining with Incremental Learning: In real scenarios, the data annotations are usually open-world and non-stationary, where novel classes may occur continuously and incrementally. However, directly turning into incremental learning may lead to catastrophic forgetting problems. On the other hand, current open world object detection [308] only focuses on novel class localization, rather than classification. How to handle both catastrophic forgetting problems and novel class detection in one framework is worth exploring in the future.

Combining with Large Language Models: Compared with VLMs, most LLMs contain more text concepts, which naturally have a broader scope than various dataset taxonomies, even larger than the recent V3Det [49] dataset. Thus, how to better align the LLMs knowledge [323] with visual detectors or segmenters to achieve stronger zero-shot results still needs exploration.

V. CONCLUSION

This survey offers a detailed examination of the latest developments in open vocabulary learning in computer vision, which appears to be a first of its kind. We provide an overview of the necessary background knowledge, which includes fundamental concepts and introductory knowledge of detection, segmentation, and vision language pre-training. Following that, we summarize more than 50 different models used for various scene understanding tasks. For each task, we categorize the methods based on their technical viewpoint. Additionally, we provide information regarding several closely related domains. In the experiment section, we provide a detailed description of the settings and compare results fairly. Finally, we summarize several challenges and also point out several future research directions for open vocabulary learning.

REFERENCES

- [1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [2] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2980–2988.
- [3] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Representations*, 2021, pp. 1–21.
- [4] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 213–229.
- [5] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 213–229.
- [6] B. Romera-Paredes and P. Torr, "An embarrassingly simple approach to zero-shot learning," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2152–2161.
- [7] A. Bansal, K. Sikka, G. Sharma, R. Chellappa, and A. Divakaran, "Zero-shot object detection," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 384–400.
- [8] Y. Xian, Z. Akata, G. Sharma, Q. Nguyen, M. Hein, and B. Schiele, "Latent embeddings for zero-shot classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 69–77.

- [9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2019, pp. 4171–4186.
- [10] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 3111–3119.
- [11] A. Zareian, K. D. Rosa, D. H. Hu, and S.-F. Chang, "Open-vocabulary object detection using captions," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 14388–14397.
- [12] X. Gu, T.-Y. Lin, W. Kuo, and Y. Cui, "Open-vocabulary object detection via vision and language knowledge distillation," in *Proc. Int. Conf. Learn. Representations*, 2022, pp. 1–20.
- [13] G. Ghiasi, X. Gu, Y. Cui, and T.-Y. Lin, "Open-vocabulary image segmentation," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 540–557.
- [14] B. Li, K. Q. Weinberger, S. Belongie, V. Koltun, and R. Ranftl, "Language-driven semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Learn. Representations*, 2023, pp. 11207–11216.
- [15] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.
- [16] C. Jia et al., "Scaling up visual and vision-language representation learning with noisy text supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 4904–4916.
- [17] Y. Song, T. Wang, S. K. Mondal, and J. P. Sahoo, "A comprehensive survey of few-shot learning: Evolution, applications, challenges, and opportunities," *ACM Comput. Surv.*, vol. 55, pp. 140, 2022.
- [18] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM Comput. Surv.*, vol. 53, pp. 1–34, 2020.
- [19] W. Wang, V. W. Zheng, H. Yu, and C. Miao, "A survey of zero-shot learning: Settings, methods, and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, pp. 1–37, 2019.
- [20] M. Köhler, M. Eisenbach, and H.-M. Gross, "Few-shot object detection: A comprehensive survey," *IEEE Trans. Neural Netw. Learn. Syst.*, to be published, doi: [10.1109/TNNLS.2023.3265051](https://doi.org/10.1109/TNNLS.2023.3265051).
- [21] Z. Hu, Y. Sun, and Y. Yang, "Suppressing the heterogeneity: A strong feature extractor for few-shot segmentation," in *Proc. Int. Conf. Learn. Representations*, 2023, pp. 1–15.
- [22] P. Xu, X. Zhu, and D. A. Clifton, "Multimodal learning with transformers: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 10, pp. 12113–12132, Oct. 2023.
- [23] T. Baltrušaitis, C. Ahuja, and L.-P. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 2, pp. 423–443, Feb. 2019.
- [24] X. Li et al., "Transformer-based visual segmentation: A survey," 2023, *arXiv:2304.09854*.
- [25] L. Ruan and Q. Jin, "Survey: Transformer based video-language pre-training," *AI Open*, vol. 3, pp. 1–13, 2022.
- [26] A. Bendale and T. E. Boulton, "Towards open world recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1893–1902.
- [27] X. Sun, H. Ding, C. Zhang, G. Lin, and K.-V. Ling, "M2IOSR: Maximal mutual information open set recognition," 2021, *arXiv:2108.02373*.
- [28] J. Yang, K. Zhou, Y. Li, and Z. Liu, "Generalized out-of-distribution detection: A survey," 2021, *arXiv:2110.11334*.
- [29] Z. Liu et al., "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9992–10002.
- [30] C. Geng, S.-J. Huang, and S. Chen, "Recent advances in open set recognition: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 10, pp. 3614–3631, Oct. 2021.
- [31] K. Joseph, S. Khan, F. S. Khan, and V. N. Balasubramanian, "Towards open world object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 5826–5836.
- [32] X. Chen et al., "Microsoft COCO captions: Data collection and evaluation server," 2015, *arXiv:1504.00325*.
- [33] A. R. Dhamija, M. Günther, and T. E. Boulton, "Reducing network agnostophobia," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2018.
- [34] A. Bendale and T. E. Boulton, "Towards open set deep networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 87–90.
- [35] M. Hassen and P. K. Chan, "Learning a neural-network-based representation for open set recognition," in *Proc. SIAM Int. Conf. Data Mining*, 2018, pp. 154–162.
- [36] L. Shu, H. Xu, and B. Liu, "DOC: Deep open classification of text documents," in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2017, pp. 2911–2916.
- [37] A. Bendale and T. E. Boulton, "Towards open world recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 1893–1902.
- [38] K. J. Joseph, S. H. Khan, F. S. Khan, and V. N. Balasubramanian, "Towards open world object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 5826–5836.
- [39] Y. Zheng, J. Wu, Y. Qin, F. Zhang, and L. Cui, "Zero-shot instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 2593–2602.
- [40] H. Zhao, X. Puig, B. Zhou, S. Fidler, and A. Torralba, "Open vocabulary scene parsing," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2021–2029.
- [41] A. Kirillov et al., "Segment anything," 2023, *arXiv:2304.02643*.
- [42] H. Touvron et al., "LLaMA: Open and efficient foundation language models," 2023, *arXiv:2302.13971*.
- [43] C. Lin et al., "Learning object-language alignments for open-vocabulary object detection," in *Proc. Int. Conf. Learn. Representations*, 2023, pp. 1–16.
- [44] C. Feng et al., "PromptDet: Towards open-vocabulary detection using uncurated images," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 701–717.
- [45] G. Ghiasi, X. Gu, Y. Cui, and T.-Y. Lin, "Scaling open-vocabulary image segmentation with image-level labels," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 540–557.
- [46] J. Wu et al., "Betrayed by captions: Joint caption grounding and generation for open vocabulary instance segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 21881–21891.
- [47] Z. Ding, J. Wang, and Z. Tu, "Open-vocabulary panoptic segmentation with maskclip," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 8090–8102.
- [48] A. Gupta, P. Dollar, and R. Girshick, "LVIS: A dataset for large vocabulary instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 5351–5359.
- [49] J. Wang et al., "V3Det: Vast vocabulary visual detection dataset," 2023, *arXiv:2304.03752*.
- [50] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, "Semantic understanding of scenes through the ADE20K dataset," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5122–5130.
- [51] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes challenge: A retrospective," *Int. J. Comput. Vis.*, vol. 11, pp. 98–136, 2015.
- [52] R. Mottaghi et al., "The role of context for object detection and semantic segmentation in the wild," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 891–898.
- [53] M. Cordts et al., "The CityScapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 3213–3223.
- [54] J. Miao, Y. Wei, Y. Wu, C. Liang, G. Li, and Y. Yang, "Vspw: A large-scale dataset for video scene parsing in the wild," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 4133–4143.
- [55] L. Yang, Y. Fan, and N. Xu, "Video instance segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 5188–5197.
- [56] H. Wang et al., "Towards open-vocabulary video instance segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 4057–4066.
- [57] H. Ding, C. Liu, S. He, X. Jiang, P. H. Torr, and S. Bai, "MOSE: A new dataset for video object segmentation in complex scenes," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 20167–20177.
- [58] A. Dave, T. Khurana, P. Tokmakov, C. Schmid, and D. Ramanan, "TAO: A large-scale benchmark for tracking any object," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 436–454.
- [59] N. Kardan and K. O. Stanley, "Mitigating fooling with competitive overcomplete output layer neural networks," in *Proc. Int. Joint Conf. Neural Netw.*, 2017, pp. 518–525.
- [60] Z. Ge, S. Demyanov, Z. Chen, and R. Garnavi, "Generative openmax for multi-class open set classification," 2017, *arXiv:1707.07418*.
- [61] L. Neal, M. L. Olson, X. Z. Fern, W.-K. Wong, and F. Li, "Open set learning with counterfactual images," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 613–628.
- [62] I. Jo, J. Kim, H. Kang, Y.-D. Kim, and S. Choi, "Open set recognition by regularising classifier with fake data generated by generative adversarial networks," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2018, pp. 2686–2690.
- [63] Y. Yu, W.-Y. Qu, N. Li, and Z. Guo, "Open category classification by adversarial sample generation," in *Proc. Int. Joint Conf. Artif. Intell.*, 2017, pp. 3357–3363.
- [64] Y. Yang, C. Hou, Y. Lang, D. Guan, D. Huang, and J. Xu, "Open-set human activity recognition based on micro-Doppler signatures," *Pattern Recognit.*, vol. 85, pp. 60–69, 2019.

- [65] C. Geng and S. Chen, "Collective decision for open set recognition," *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 1, pp. 192–204, Jan. 2022.
- [66] J. Zheng, W. Li, J. Hong, L. Petersson, and N. Barnes, "Towards open-set object detection and discovery," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 3961–3970.
- [67] Z. Chen, Y. Luo, Z. Wang, Z. Wang, X. Yu, and Z. Huang, "Towards open world active learning for 3D object detection," 2023, *arXiv:2310.10391*.
- [68] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," in *Proc. Int. Conf. Learn. Representations*, 2016, pp. 1–12.
- [69] Y.-C. Hsu, Y. Shen, H. Jin, and Z. Kira, "Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 10951–10960.
- [70] K. Lee, K. Lee, K. Min, Y. Zhang, J. Shin, and H. Lee, "Hierarchical novelty detection for visual object recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 1034–1042.
- [71] K. Lee, H. Lee, K. Lee, and J. Shin, "Training confidence-calibrated classifiers for detecting out-of-distribution samples," in *Proc. Int. Conf. Learn. Representations*, 2017, pp. 1–16.
- [72] B. Zong et al., "Deep autoencoding Gaussian mixture model for unsupervised anomaly detection," in *Proc. Int. Conf. Learn. Representations*, 2018, pp. 1–19.
- [73] D. Abati, A. Porrello, S. Calderara, and R. Cucchiara, "Latent space autoregression for novelty detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 481–490.
- [74] S. Pidhorskyi, R. Almhosen, D. A. Adjeroh, and G. Doretto, "Generative probabilistic novelty detection with adversarial autoencoders," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 6823–6834.
- [75] J. J. Ren, S. Fort, J. Z. Liu, A. G. Roy, S. Padhy, and B. Lakshminarayanan, "A simple fix to Mahalanobis distance for improving near-OOD detection," 2021, *arXiv:2106.09022*.
- [76] E. Techapanurak, M. Suganuma, and T. Okatani, "Hyperparameter-free out-of-distribution detection using cosine similarity," in *Proc. Asian Conf. Comput. Vis.*, 2020, pp. 53–69.
- [77] T. Denouden, R. Salay, K. Czarniecki, V. Abdelzad, B. Phan, and S. Vernekar, "Improving reconstruction autoencoder out-of-distribution detection with Mahalanobis distance," 2018, *arXiv: 1812.02765*.
- [78] Y. Zhou, "Rethinking reconstruction autoencoder-based out-of-distribution detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 7379–7387.
- [79] Y. Yang, R. Gao, and Q. Xu, "Out-of-distribution detection with semantic mismatch under masking," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 373–390.
- [80] Y. Xian, S. Choudhury, Y. He, B. Schiele, and Z. Akata, "Semantic projection network for zero-and few-label semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 8256–8265.
- [81] H. Zhang and H. Ding, "Prototypical matching and open set rejection for zero-shot semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 6974–6983.
- [82] S. He, H. Ding, and W. Jiang, "Semantic-promoted debiasing and background disambiguation for zero-shot instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 19498–19507.
- [83] S. He, X. Jiang, W. Jiang, and H. Ding, "Prototype adaption and projection for few- and zero-shot 3D point cloud semantic segmentation," *IEEE Trans. Image Process.*, vol. 32, pp. 3199–3211, 2023.
- [84] D. Baek, Y. Oh, and B. Ham, "Exploiting a joint embedding space for generalized zero-shot semantic segmentation," 2021, *arXiv:2108.06536*.
- [85] S. He, H. Ding, and W. Jiang, "Primitive generation and semantic-related alignment for universal zero-shot segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 11238–11247.
- [86] M. Bucher, T.-H. Vu, M. Cord, and P. Pérez, "Zero-shot semantic segmentation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 468–479.
- [87] Z. Gu, S. Zhou, L. Niu, Z. Zhao, and L. Zhang, "Context-aware feature generation for zero-shot semantic segmentation," in *Proc. 28th ACM Int. Conf. Multimedia*, 2020, pp. 1921–1929.
- [88] F. Shen, J. Liu, and P. Hu, "Counterfactual generative zero-shot semantic segmentation," 2021, *arXiv:2106.06360*.
- [89] A. Joulin, E. Grave, P. Bojanowski, M. Douze, H. Jégou, and T. Mikolov, "FastText.zip: Compressing text classification models," 2016, *arXiv:1612.03651*.
- [90] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 3111–3119.
- [91] P. Li, Y. Wei, and Y. Yang, "Consistent structural relation learning for zero-shot segmentation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 10317–10327.
- [92] Z. Gu, S. Zhou, L. Niu, Z. Zhao, and L. Zhang, "From pixel to patch: Synthesize context-aware features for zero-shot semantic segmentation," 2020, *arXiv: 2009.12232*.
- [93] G. Tian, S. Wang, J. Feng, L. Zhou, and Y. Mu, "Cap2seg: Inferring semantic and spatial context from captions for zero-shot image segmentation," in *Proc. 28th ACM Int. Conf. Multimedia*, 2020, pp. 4125–4134.
- [94] J. Liu, Y. Sun, C. Han, Z. Dou, and W. Li, "Deep representation learning on long-tailed data: A learnable embedding augmentation perspective," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2970–2979.
- [95] J. Ren et al., "Balanced meta-softmax for long-tailed visual recognition," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 4175–4186.
- [96] Y. Li et al., "Overcoming classifier imbalance for long-tail object detection with balanced group softmax," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 10991–11000.
- [97] X. Hu, Y. Jiang, K. Tang, J. Chen, C. Miao, and H. Zhang, "Learning to segment the tail," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 14045–14054.
- [98] T. Wang et al., "The devil is in classification: A simple framework for long-tail instance segmentation," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 728–744.
- [99] Y. Li et al., "Few-shot object detection via classification refinement and distractor retreatment," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 15395–15403.
- [100] H. Lee, M. Lee, and N. Kwak, "Few-shot object detection by attending to per-sample-prototype," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2022, pp. 2445–2454.
- [101] A. Wu, S. Zhao, C. Deng, and W. Liu, "Generalized and discriminative few-shot object detection via SVD-dictionary enhancement," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 6353–6364.
- [102] X. Wang, T. Huang, J. Gonzalez, T. Darrell, and F. Yu, "Frustratingly simple few-shot object detection," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 9919–9928.
- [103] L. Qiao, Y. Zhao, Z. Li, X. Qiu, J. Wu, and C. Zhang, "DeFRNC: Decoupled faster R-CNN for few-shot object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 8681–8690.
- [104] C. Zhu, F. Chen, U. Ahmed, Z. Shen, and M. Savvides, "Semantic relation reasoning for shot-stable few-shot object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 8782–8791.
- [105] P. Kaul, W. Xie, and A. Zisserman, "Label, verify, correct: A simple few-shot object detection method," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 14237–14247.
- [106] H. Ding, H. Zhang, and X. Jiang, "Self-regularized prototypical network for few-shot semantic segmentation," *Pattern Recognit.*, vol. 133, 2023, Art. no. 109018.
- [107] K. Rakelly, E. Shelhamer, T. Darrell, A. Efros, and S. Levine, "Conditional networks for few-shot semantic segmentation," in *Proc. Int. Conf. Learn. Representations*, 2018, pp. 1–4.
- [108] D. A. Ganea, B. Boom, and R. Poppe, "Incremental few-shot instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 1185–1194.
- [109] C. Michaelis, I. Ustyuzhaninov, M. Bethge, and A. S. Ecker, "One-shot instance segmentation," 2018, *arXiv: 1811.11507*.
- [110] Z. Fan et al., "FGN: Fully guided network for few-shot instance segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 9172–9181.
- [111] X. Yan, Z. Chen, A. Xu, X. Wang, X. Liang, and L. Lin, "Meta R-CNN: Towards general solver for instance-level low-shot learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 9577–9586.
- [112] S. Wu, W. Zhang, S. Jin, W. Liu, and C. C. Loy, "Aligning bag of regions for open-vocabulary object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 15254–15264.
- [113] J. Ding, N. Xue, G.-S. Xia, and D. Dai, "Decoupling zero-shot semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 11583–11592.
- [114] X. Chen, S. Li, S.-N. Lim, A. Torralba, and H. Zhao, "Open-vocabulary panoptic segmentation with embedding modulation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 1141–1150.
- [115] S. Peng, K. Genova, C. M. Jiang, A. Tagliasacchi, M. Pollefeys, and T. Funkhouser, "Openscene: 3D scene understanding with open vocabularies," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 815–824.

- [116] L. Xue et al., "Ulip: Learning unified representation of language, image and point cloud for 3D understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 1179–1189.
- [117] S. Xu et al., "DST-DET: Simple dynamic self-training for open-vocabulary object detection," 2023, *arXiv:2310.01393*.
- [118] M. Gao et al., "Open vocabulary object detection with pseudo bounding-box labels," in *Proc. Eur. Conf. Comput. Vis.*, 2021, pp. 266–282.
- [119] C. Zhou, C. C. Loy, and B. Dai, "Extract free dense labels from clip," in *Proc. Eur. Conf. Comput. Vis.*, 2021, pp. 696–712.
- [120] F. Liang et al., "Open-vocabulary semantic segmentation with mask-adapted clip," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 7061–7070.
- [121] D. Huynh, J. Kuen, Z. Lin, J. Gu, and E. Elhamifar, "Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 7020–7031.
- [122] W. Lin et al., "Match, expand and improve: Unsupervised finetuning for zero-shot action recognition with language knowledge," in *Proc. Int. Conf. Comput. Vis.*, 2023, pp. 2839–2850.
- [123] R. Ding, J. Yang, C. Xue, W. Zhang, S. Bai, and X. Qi, "PLA: Language-driven open-vocabulary 3D scene understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 7010–7019.
- [124] J. Xu et al., "GroupViT: Semantic segmentation emerges from text supervision," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 18134–18144.
- [125] H. Luo, J. Bao, Y. Wu, X. He, and T. Li, "Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 23033–23044.
- [126] V. Vs et al., "Mask-free OVIS: Open-vocabulary instance segmentation without manual mask annotations," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 23539–23549.
- [127] H. Zhang et al., "A simple framework for open-vocabulary segmentation and detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 1020–1031.
- [128] S. Li, M. H. Li, P. Wang, and L. Zhang, "OpenSD: Unified open-vocabulary segmentation and detection," 2023, *arXiv:2312.06703*.
- [129] X. Zou et al., "Generalized decoding for pixel, image, and language," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 15116–15127.
- [130] J. Qin et al., "FreeSeg: Unified, universal and open-vocabulary image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 19446–19455.
- [131] L. Karazija, I. Laina, A. Vedaldi, and C. Rupprecht, "Diffusion models for zero-shot open-vocabulary segmentation," 2023, *arXiv:2306.09316*.
- [132] J. Xu, S. Liu, A. Vahdat, W. Byeon, X. Wang, and S. De Mello, "Open-vocabulary panoptic segmentation with text-to-image diffusion models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 2955–2966.
- [133] X. Zhou, R. Girdhar, A. Joulin, P. Krähenbühl, and I. Misra, "Detecting twenty-thousand classes using image-level supervision," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 350–368.
- [134] M. Minderer, A. A. Gritsenko, and N. Houlsby, "Scaling open-vocabulary object detection," 2023, *arXiv:2306.09683*.
- [135] S. Xu et al., "DST-DET: Simple dynamic self-training for open-vocabulary object detection," 2023, *arXiv:2310.01393*.
- [136] W. Kuo, Y. Cui, X. Gu, A. J. Piergiovanni, and A. Angelova, "F-VLM: Open-vocabulary object detection upon frozen vision and language models," in *Proc. Int. Conf. Learn. Representations*, 2023, pp. 1–20.
- [137] J. Xie, W. Li, X. Li, Z. Liu, Y. S. Ong, and C. C. Loy, "MosaicFusion: Diffusion models as data augmenters for large vocabulary instance segmentation," 2023, *arXiv:2309.13042*.
- [138] W. Wu, Y. Zhao, M. Z. Shou, H. Zhou, and C. Shen, "DiffuMask: Synthesizing images with pixel-level annotations for semantic segmentation using diffusion models," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 1206–1217.
- [139] H. Ding, H. Zhang, J. Liu, J. Li, Z. Feng, and X. Jiang, "Interaction via bi-directional graph of semantic region affinity for scene parsing," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 15848–15858.
- [140] C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, and N. Sang, "Learning a discriminative feature network for semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 1857–1866.
- [141] H. Ding, X. Jiang, B. Shuai, A. Q. Liu, and G. Wang, "Semantic segmentation with context encoding and multi-path decoding," *IEEE Trans. Image Process.*, vol. 29, pp. 3520–3533, 2020.
- [142] C. Peng, X. Zhang, G. Yu, G. Luo, and J. Sun, "Large kernel matters—improve semantic segmentation by global convolutional network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 4353–4361.
- [143] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2881–2890.
- [144] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," 2017, *arXiv:1706.05587*.
- [145] H. Ding, X. Jiang, B. Shuai, A. Q. Liu, and G. Wang, "Context contrasted feature and gated multi-scale aggregation for scene segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 2393–2402.
- [146] X. Li, H. Zhao, L. Han, Y. Tong, S. Tan, and K. Yang, "Gated fully fusion for semantic segmentation," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 11418–11425.
- [147] X. Li et al., "Semantic flow for fast and accurate scene parsing," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 775–793.
- [148] X. Li et al., "SFNet: Faster, accurate, and domain agnostic semantic segmentation via semantic flow," 2022, *arXiv:2207.04415*.
- [149] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7794–7803.
- [150] A. Kirillov, Y. Wu, K. He, and R. Girshick, "PointRend: Image segmentation as rendering," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 9796–9805.
- [151] H. Ding, X. Jiang, A. Q. Liu, N. M. Thalmann, and G. Wang, "Boundary-aware feature propagation for scene segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 6819–6829.
- [152] X. Li et al., "Improving semantic segmentation via decoupled body and edge supervision," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 435–452.
- [153] H. He, X. Li, G. Cheng, Y. Tong, and L. Weng, "BoundarySqueeze: Image segmentation as boundary squeezing," 2021, *arXiv:2105.11668*.
- [154] J. Fu, J. Liu, H. Tian, Z. Fang, and H. Lu, "Dual attention network for scene segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3146–3154.
- [155] H. Hu, J. Gu, Z. Zhang, J. Dai, and Y. Wei, "Relation networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 3588–3597.
- [156] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *TPAMI*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [157] Z. Tian, C. Shen, H. Chen, and T. He, "FCOS: Fully convolutional one-stage object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 9627–9636.
- [158] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2980–2988.
- [159] Z. Tian, C. Shen, and H. Chen, "Conditional convolutions for instance segmentation," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 282–298.
- [160] D. Neven, B. D. Brabandere, M. Proesmans, and L. V. Gool, "Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 8837–8845.
- [161] B. De Brabandere, D. Neven, and L. Van Gool, "Semantic instance segmentation with a discriminative loss function," 2017, *arXiv:1708.02551*.
- [162] K. Chen et al., "Hybrid task cascade for instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4974–4983.
- [163] R. Zhang, Z. Tian, C. Shen, M. You, and Y. Yan, "Mask encoding for single shot instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 10226–10235.
- [164] D. Bolya, C. Zhou, F. Xiao, and Y. J. Lee, "Yolact: Real-time instance segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 9157–9166.
- [165] S. Qiao, L.-C. Chen, and A. Yuille, "Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 10213–10224.
- [166] B. Cheng et al., "Panoptic-Deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 12475–12485.
- [167] X. Wang, R. Zhang, T. Kong, L. Li, and C. Shen, "SOLOv2: Dynamic and fast instance segmentation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 17721–17732.
- [168] Q. Zhou et al., "TransVOD: End-to-end video object detection with spatial-temporal transformers," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 6, pp. 7853–7869, Jun. 2023.

- [169] H. Ding, C. Liu, S. Wang, and X. Jiang, "VLT: Vision-language transformer and query generation for referring segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 6, pp. 7900–7916, Jun. 2023.
- [170] C. Liu, H. Ding, and X. Jiang, "GRES: Generalized referring expression segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 23592–23601.
- [171] X. Li, S. Xu, Y. Yang, G. Cheng, Y. Tong, and D. Tao, "Panoptic-partformer: Learning a unified model for panoptic part segmentation," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 729–747.
- [172] X. Li et al., "Panoptic-PartFormer++: A unified and decoupled view for panoptic part segmentation," 2023, *arXiv:2301.00954*.
- [173] H. Yuan et al., "Polyphonicformer: Unified query learning for depth-aware video panoptic segmentation," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 582–599.
- [174] S. Xu, X. Li, J. Wang, G. Cheng, Y. Tong, and D. Tao, "Fashionformer: A simple, effective and unified baseline for human fashion segmentation and recognition," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 545–563.
- [175] Y. Xu et al., "Multi-task learning with multi-query transformer for dense prediction," 2022, *arXiv:2205.14354*.
- [176] X. Li, H. Yuan, W. Zhang, J. Pang, G. Cheng, and C. C. Loy, "Tube-link: A flexible cross tube baseline for universal video segmentation," 2023, *arXiv:2303.12782*.
- [177] Z. Hu, Y. Sun, J. Wang, and Y. Yang, "DAC-DETR: Divide the attention layers and conquer," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 1–12.
- [178] H. W. Kuhn, "The Hungarian method for the assignment problem," *Nav. Res. Logistics Quart.*, vol. 2, pp. 83–97, 1955.
- [179] H. Wang, Y. Zhu, H. Adam, A. Yuille, and L.-C. Chen, "Max-deeplab: End-to-end panoptic segmentation with mask transformers," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 5463–5474.
- [180] W. Zhang, J. Pang, K. Chen, and C. C. Loy, "K-Net: Towards unified image segmentation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 10326–10338.
- [181] B. Cheng, A. G. Schwing, and A. Kirillov, "Per-pixel classification is not all you need for semantic segmentation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 17864–17875.
- [182] H. Wang, Y. Zhu, B. Green, H. Adam, A. Yuille, and L.-C. Chen, "Axial-deeplab: Stand-alone axial-attention for panoptic segmentation," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 108–126.
- [183] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, "Masked-attention mask transformer for universal image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 1290–1299.
- [184] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable {DETR}: Deformable transformers for end-to-end object detection," in *Proc. Int. Conf. Learn. Representations*, 2021, pp. 1–16.
- [185] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, "Detectron2," 2019. [Online]. Available: <https://github.com/facebookresearch/detectron2>
- [186] H. Ding, C. Liu, S. Wang, and X. Jiang, "Vision-language transformer and query generation for referring segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 16321–16330.
- [187] L. H. Li, M. Yatskar, D. Yin, C.-J. Hsieh, and K.-W. Chang, "VisualBERT: A simple and performant baseline for vision and language," 2019, *arXiv:1908.03557*.
- [188] G. Li, N. Duan, Y. Fang, M. Gong, and D. Jiang, "Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training," in *Proc. AAAI Conf. Artif. Intell.*, 2021, pp. 11336–11344.
- [189] J. Lu, D. Batra, D. Parikh, and S. Lee, "ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 13–23.
- [190] H. Tan and M. Bansal, "LXMERT: Learning cross-modality encoder representations from transformers," in *Proc. Empirical Methods Natural Lang. Process.*, 2019, pp. 5100–5111.
- [191] D. Qi, L. Su, J. Song, E. Cui, T. Bharti, and A. Sacheti, "ImageBERT: Cross-modal pre-training with large-scale weak-supervised image-text data," 2020, *arXiv:2001.07966*.
- [192] Z. Huang, Z. Zeng, B. Liu, D. Fu, and J. Fu, "Pixel-BERT: Aligning image pixels with text by deep multi-modal transformers," 2020, *arXiv:2004.00849*.
- [193] Y. Li, H. Fan, R. Hu, C. Feichtenhofer, and K. He, "Scaling language-image pre-training via masking," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 23390–23400.
- [194] N. Mu, A. Kirillov, D. Wagner, and S. Xie, "Slip: Self-supervision meets language-image pre-training," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 529–544.
- [195] Q. Sun, Y. Fang, L. Wu, X. Wang, and Y. Cao, "EVA-CLIP: Improved training techniques for clip at scale," 2023, *arXiv:2303.15389*.
- [196] J. Li, R. R. Selvaraju, A. D. Gotmare, S. Joty, C. Xiong, and S. Hoi, "Align before fuse: Vision and language representation learning with momentum distillation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 9694–9705.
- [197] J. Yu, Z. Wang, V. Vasudevan, L. Yeung, M. Seyedhosseini, and Y. Wu, "COCA: Contrastive captioners are image-text foundation models," 2022, *arXiv:2205.01917*.
- [198] S. Yan et al., "Video-text modeling with zero-shot transfer from contrastive captioners," 2022, *arXiv:2212.04979*.
- [199] J. Li, D. Li, C. Xiong, and S. Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *Proc. Int. Conf. Mach. Learn.*, 2022, pp. 12888–12900.
- [200] J. Li, D. Li, S. Savarese, and S. Hoi, "BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 19730–19742.
- [201] R. Yeh, J. Xiong, W.-M. Hwu, M. Do, and A. Schwing, "Interpretable and globally optimal prediction for textual grounding using image concepts," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 1913–1923.
- [202] J. Wu, X. Li, X. Li, H. Ding, Y. Tong, and D. Tao, "Towards robust referring image segmentation," 2022, *arXiv:2209.09554*.
- [203] H. Ding, C. Liu, S. He, X. Jiang, and C. C. Loy, "MeViS: A large-scale benchmark for video segmentation with motion expressions," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 2694–2703.
- [204] J. Deng et al., "TransVG: End-to-end visual grounding with transformers," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 1749–1759.
- [205] M. Li et al., "Iterative robust visual grounding with masked reference based centerpoint supervision," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 4651–4656.
- [206] C. Wang et al., "OV-VG: A benchmark for open-vocabulary visual grounding," 2023, *arXiv:2310.14374*.
- [207] S. Liu, T. Hui, S. Huang, Y. Wei, B. Li, and G. Li, "Cross-Modal progressive comprehension for referring segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 9, pp. 4761–4775, Sep. 2022.
- [208] G. Luo et al., "Multi-task collaborative network for joint referring expression comprehension and segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 10034–10043.
- [209] G. Feng, Z. Hu, L. Zhang, and H. Lu, "Encoder fusion network with co-attention embedding for referring image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 15506–15515.
- [210] Z. Yang, J. Wang, Y. Tang, K. Chen, H. Zhao, and P. H. Torr, "LAVT: Language-aware vision transformer for referring image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 18155–18165.
- [211] L. Chen et al., "Ref-NMS: Breaking proposal bottlenecks in two-stage referring expression grounding," in *Proc. AAAI Conf. Artif. Intell.*, 2021, pp. 1036–1044.
- [212] A. Kamath et al., "MDETR - modulated detection for end-to-end multi-modal understanding," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 1780–1790.
- [213] S. Liu et al., "Grounding DINO: Marrying DINO with grounded pre-training for open-set object detection," 2023, *arXiv:2303.05499*.
- [214] S. Zhao et al., "Exploiting unlabeled data with vision and language models for object detection," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 159–175.
- [215] Y. Zhong et al., "RegionCLIP: Region-based language-image pretraining," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 16793–16803.
- [216] Y. Du, F. Wei, Z. Zhang, M. Shi, Y. Gao, and G. C. Li, "Learning to prompt for open-vocabulary object detection with vision-language model," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 14084–14093.
- [217] Y. Zang, W. Li, K. Zhou, C. Huang, and C. C. Loy, "Open-vocabulary DETR with conditional matching," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 106–122.
- [218] X. Wu, F. Zhu, R. Zhao, and H. Li, "CORA: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 7031–7040.
- [219] W. Kuo et al., "Mammut: A simple architecture for joint learning for multimodal tasks," 2023, *arXiv:2303.16839*.

- [220] Q. Liu, Y. Wen, J. Han, C. Xu, H. Xu, and X. Liang, "Open-world semantic segmentation via contrasting and clustering vision-language embedding," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 275–292.
- [221] J. Xu et al., "Learning open-vocabulary semantic segmentation models from natural language supervision," 2023, *arXiv:2301.09121*.
- [222] C. Pham, T. Vu, and K. Nguyen, "LP-OVOD: Open-vocabulary object detection by linear probing," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2024, pp. 779–788.
- [223] Z. Ma et al., "Open-vocabulary one-stage detection with hierarchical visual-language knowledge distillation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 14074–14083.
- [224] H. Rasheed, M. Maaz, M. U. Khattak, S. Khan, and F. S. Khan, "Bridging the gap between object and image-level representations for open-vocabulary detection," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 33781–33794.
- [225] L. Wang et al., "Object-aware distillation pyramid for open-vocabulary object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 11186–11196.
- [226] H.-C. Cho, W. Y. Jhoo, W. Kang, and B. Roh, "Open-vocabulary object detection using pseudo caption labels," 2023, *arXiv:2303.13040*.
- [227] K. Chen et al., "OvarNet: Towards open-vocabulary object attribute recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 23518–23527.
- [228] G. Patterson and J. Hays, "COCO attributes: Attributes for people, animals, and objects," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 85–100.
- [229] K. Buettner and A. Kovashka, "Enhancing the role of context in region-word alignment for object detection," 2023, *arXiv:2303.10093*.
- [230] L. H. Li et al., "Grounded language-image pre-training," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 10965–10975.
- [231] H. Zhang et al., "GLIPv2: Unifying localization and vision-language understanding," 2022, *arXiv:2206.05836*.
- [232] M. Minderer et al., "Simple open-vocabulary object detection with vision transformers," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 728–755.
- [233] D. Kim, A. Angelova, and W. Kuo, "Detection-oriented image-text pretraining for open-vocabulary detection," 2023, *arXiv:2310.00161*.
- [234] O. Russakovsky et al., "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, pp. 211–252, 2015.
- [235] P. Kaul, W. Xie, and A. Zisserman, "Multi-modal classifiers for open-vocabulary object detection," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 15946–15969.
- [236] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *Int. J. Comput. Vis.*, vol. 128, pp. 336–359, 2016.
- [237] M. A. Bravo, S. Mittal, and T. Brox, "Localized vision-language matching for open-vocabulary object detection," in *Proc. DAGM German Conf. Pattern Recognit.*, 2022, pp. 393–408.
- [238] H. Zhao et al., "X-paste: Revisit copy-paste at scale with clip and stablediffusion," 2022, *arXiv:2212.03863*.
- [239] Z. Li, Q. Zhou, X. Zhang, Y. Zhang, Y. Wang, and W. Xie, "Guiding text-to-image diffusion model towards grounded generation," 2023, *arXiv:2301.05221*.
- [240] G. Ghiasi et al., "Simple copy-paste is a strong data augmentation method for instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 2918–2928.
- [241] R. Arandjelović, A. Andonian, A. Mensch, O. J. Hénaff, J.-B. Alayrac, and A. Zisserman, "Three ways to improve feature alignment for open vocabulary detection," 2023, *arXiv:2303.13518*.
- [242] X. Chen et al., "PaLI: A jointly-scaled multilingual language-image model," 2022, *arXiv:2209.06794*.
- [243] T. Brown et al., "Language models are few-shot learners," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 1877–1901.
- [244] R. Zhang et al., "Tip-adaptor: Training-free clip-adaptor for better vision-language modeling," 2021, *arXiv:2111.03930*.
- [245] K. Zhou, J. Yang, C. C. Loy, and Z. Liu, "Learning to prompt for vision-language models," *Int. J. Comput. Vis.*, vol. 130, pp. 2337–2348, 2022.
- [246] H. Song and J. Bang, "Prompt-guided transformers for end-to-end open-vocabulary object detection," 2023, *arXiv:2303.14386*.
- [247] L. Yao et al., "DetCLIPv2: Scalable open-vocabulary object detection pre-training via word-region alignment," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 23497–23506.
- [248] S. Zhang, C. Chi, Y. Yao, Z. Lei, and S. Z. Li, "Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 9759–9768.
- [249] C. Ma, Y. Jiang, X. Wen, Z. Yuan, and X. Qi, "CoDet: Co-occurrence guided region-word alignment for open-vocabulary object detection," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 1–17.
- [250] D. Kim, A. Angelova, and W. Kuo, "Region-aware pretraining for open-vocabulary object detection with vision transformers," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 11144–11154.
- [251] C. Ma, Y. Yang, Y. Wang, Y. Zhang, and W. Xie, "Open-vocabulary semantic segmentation with frozen vision-language models," in *Proc. Brit. Mach. Vis. Conf.*, 2022, pp. 1–15.
- [252] M. Xu, Z. Zhang, F. Wei, H. Hu, and X. Bai, "Side adapter network for open-vocabulary semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 2945–2954.
- [253] S. Cho et al., "CAT-Seg: Cost aggregation for open-vocabulary semantic segmentation," 2023, *arXiv:2303.11797*.
- [254] K. Han et al., "Global knowledge calibration for fast open-vocabulary segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 797–807.
- [255] J. Li, P. Chen, S. Qian, and J. Jia, "TagCLIP: Improving discrimination ability of open-vocabulary semantic segmentation," 2023, *arXiv:2304.07547*.
- [256] Q. Yu, J. He, X. Deng, X. Shen, and L.-C. Chen, "Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 1–20.
- [257] X. Li, "OMG-Seg: Is one model good enough for all segmentation?" 2024, *arXiv:2401.10229*.
- [258] M. Wysockańska, O. Siméoni, M. Ramamonjisoa, A. Bursuc, T. Trzeciński, and P. P'erez, "Clip-dinoiser: Teaching clip a few dino tricks," 2023, *arXiv:2312.12359*.
- [259] J. Mukhoti et al., "Open vocabulary semantic segmentation with patch aligned contrastive learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 19413–19423.
- [260] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *Int. J. Comput. Vis.*, vol. 59, pp. 167–181, 2004.
- [261] J. Cha, J. Mun, and B. Roh, "Learning to generate text-grounded mask for open-world semantic segmentation from only image-text pairs," 2022, *arXiv:2212.00785*.
- [262] S. Ren et al., "Prompt pre-training with twenty-thousand classes for open-vocabulary visual recognition," 2023, *arXiv:2304.04704*.
- [263] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 10674–10685.
- [264] T.-A. Vu et al., "Leveraging open-vocabulary diffusion to camouflaged instance segmentation," 2023, *arXiv:2312.17505*.
- [265] Z. Li, Q. Zhou, X. Zhang, Y. Zhang, Y. Wang, and W. Xie, "Guiding text-to-image diffusion model towards grounded generation," in *Proc. Int. Conf. Comput. Vis.*, 2023.
- [266] M. Wang, J. Xing, and Y. Liu, "ActionCLIP: A new paradigm for video action recognition," 2021, *arXiv:2109.08472*.
- [267] C. Ju, T. Han, K. Zheng, Y. Zhang, and W. Xie, "Prompting visual-language models for efficient video understanding," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 105–124.
- [268] H. Rasheed, M. U. Khattak, M. Maaz, S. Khan, and F. S. Khan, "Fine-tuned clip models are efficient video learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 6545–6554.
- [269] S. Li, T. Fischer, L. Ke, H. Ding, M. Danelljan, and F. Yu, "OVTrack: Open-vocabulary multiple object tracking," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 5567–5577.
- [270] P. Guo et al., "OpenVIS: Open-vocabulary video instance segmentation," 2023, *arXiv:2305.16835*.
- [271] R. Zhang et al., "Pointclip: Point cloud understanding by clip," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 8552–8562.
- [272] X. Zhu, R. Zhang, B. He, Z. Zeng, S. Zhang, and P. Gao, "PointCLIP V2: Adapting clip for powerful 3D open-world learning," in *Proc. Int. Conf. Comput. Vis.*, 2023, pp. 2639–2650.
- [273] Y. Zeng et al., "CLIP2: Contrastive language-image-point pretraining from real-world point cloud data," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 15244–15253.

- [274] M. Liu et al., "OpenShape: Scaling up 3D shape representation towards open-world understanding," 2023, *arXiv:2305.10764*.
- [275] Y. Lu et al., "Open-vocabulary 3D detection via image-level class and debiased cross-modal contrastive learning," 2022, *arXiv:2207.01987*.
- [276] I. Misra, R. Girdhar, and A. Joulin, "An end-to-end transformer model for 3D object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 2906–2917.
- [277] B. Graham, M. Engelcke, and L. Van Der Maaten, "3D semantic segmentation with submanifold sparse convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 9224–9232.
- [278] J. Carreira and A. Zisserman, "Quo Vadis, action recognition? a new model and the kinetics dataset," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6299–6308.
- [279] Z. Lin et al., "Frozen clip models are efficient video learners," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 388–404.
- [280] B. Ni et al., "Expanding language-image pretrained models for general video recognition," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 1–18.
- [281] R. Qian, Y. Li, Z. Xu, M.-H. Yang, S. Belongie, and Y. Cui, "Multimodal open-vocabulary video classification via pre-trained vision and language models," 2022, *arXiv:2207.07646*.
- [282] Z. Weng, X. Yang, A. Li, Z. Wu, and Y.-G. Jiang, "Transforming CLIP to an open-vocabulary video model via interpolated weight optimization," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 36978–36989.
- [283] T. Yang, Y. Zhu, Y. Xie, A. Zhang, C. Chen, and M. Li, "AIM: Adapting image models for efficient video action recognition," in *Proc. Int. Conf. Learn. Representations*, 2023.
- [284] Y. Chen, D. Chen, R. Liu, H. Li, and W. Peng, "Video action recognition with attentive semantic units," 2023, *arXiv:2303.09756*.
- [285] K. Kahatapitiya, A. Arnab, A. Nagrani, and M. S. Ryoo, "VicTR: Video-conditioned text representations for activity recognition," 2023, *arXiv:2304.02560*.
- [286] S. Li, M. Danelljan, H. Ding, T. E. Huang, and F. Yu, "Tracking every thing in the wild," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 498–515.
- [287] T. Zhang et al., "DVIS++: Improved decoupled framework for universal video segmentation," 2023, *arXiv:2312.13305*.
- [288] T. Huang et al., "Clip2point: Transfer clip to point cloud classification with image-depth pre-training," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 22157–22167.
- [289] G. Hess, A. Tonderski, C. Petersson, K. Åström, and L. Svensson, "LiDARclip or: How I learned to talk to point clouds," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2024, pp. 7438–7447.
- [290] Y. Lu et al., "Open-vocabulary point-cloud object detection without 3D annotation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 1190–1199.
- [291] Y. Cao, Y. Zeng, H. Xu, and D. Xu, "CoDA: Collaborative novel box discovery and cross-modal alignment for open-vocabulary 3D object detection," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 1–12.
- [292] C. Zhu, W. Zhang, T. Wang, X. Liu, and K. Chen, "Object2scene: Putting objects in context for open-vocabulary 3D detection," 2023, *arXiv:2309.09456*.
- [293] H. Zhang et al., "OpenSight: A simple open-vocabulary framework for LiDAR-based object detection," 2023, *arXiv:2312.08876*.
- [294] J. Zhang, R. Dong, and K. Ma, "CLIP-FO3D: Learning free open-world 3D scene representations from 2D dense clip," 2023, *arXiv:2303.04748*.
- [295] J. Yang, R. Ding, Z. Wang, and X. Qi, "RegionPLC: Regional point-language contrastive learning for open-world 3D scene understanding," 2023, *arXiv:2304.00962*.
- [296] M. Liu et al., "PartSLIP: Low-shot part segmentation for 3D point clouds via pretrained image-language models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 21736–21746.
- [297] A. Abdelreheem, I. Skorokhodov, M. Ovsjanikov, and P. Wonka, "SATR: Zero-shot semantic segmentation of 3D shapes," 2023, *arXiv:2304.04909*.
- [298] A. Takmaz, E. Fedele, R. W. Sumner, M. Pollefeys, F. Tombari, and F. Engelmann, "OpenMask3D: Open-vocabulary 3D instance segmentation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 1–24.
- [299] Z. Huang, X. Wu, X. Chen, H. Zhao, L. Zhu, and J. Lasenby, "Openins3D: Snap and lookup for 3D open-vocabulary instance segmentation," 2023, *arXiv:2309.00616*.
- [300] P. D. Nguyen et al., "Open3DIS: Open-vocabulary 3D instance segmentation with 2D mask guidance," 2023, *arXiv:2312.10671*.
- [301] R. Chen et al., "CLIP2scene: Towards label-efficient 3D scene understanding by CLIP," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 7020–7030.
- [302] D. Kim, T.-Y. Lin, A. Angelova, I. S. Kweon, and W. Kuo, "Learning open-world object proposals without learning to classify," *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 5453–5460, Apr. 2022.
- [303] W. Wang, M. Feiszli, H. Wang, and D. Tran, "Unidentified video objects: A benchmark for dense, open-world segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 10776–10785.
- [304] W. Wang, M. Feiszli, H. Wang, J. Malik, and D. Tran, "Open-world instance segmentation: Exploiting pseudo ground truth from learned pairwise affinity," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 4422–4432.
- [305] M. Maaz, H. Rasheed, S. Khan, F. S. Khan, R. M. Anwer, and M.-H. Yang, "Class-agnostic object detection with multi-modal transformer," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 512–531.
- [306] L. Qi et al., "Open world entity segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 7, pp. 8743–8756, Jul. 2023.
- [307] L. Qi et al., "Fine-grained entity segmentation," in *Proc. Int. Conf. Comput. Vis.*, 2023, pp. 4047–4056.
- [308] A. Gupta, S. Narayan, K. Joseph, S. Khan, F. S. Khan, and M. Shah, "OW-DETR: Open-world detection transformer," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 9235–9244.
- [309] N. Dong, Y. Zhang, M. Ding, and G. H. Lee, "Open world DETR: Transformer based open world object detection," 2022, *arXiv:2212.02969*.
- [310] O. Zohar, K.-C. Wang, and S. Yeung, "PROB: Probabilistic objectness for open world object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 11444–11453.
- [311] Z. Wu, Y. Lu, X. Chen, Z. Wu, L. Kang, and J. Yu, "UC-OWOD: Unknown-classified open world object detection," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 193–210.
- [312] J. Hwang, S. W. Oh, J.-Y. Lee, and B. Han, "Exemplar-based open-set panoptic segmentation network," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 1175–1184.
- [313] H.-M. Xu, H. Chen, L. Liu, and Y. Yin, "Dual decision improves open-set panoptic segmentation," in *Proc. Brit. Mach. Vis. Conf.*, 2022, Art. no. 3.
- [314] P. Sun et al., "Going denser with open-vocabulary part segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 15407–15419.
- [315] L. Qi et al., "AIMS: All-inclusive multi-level segmentation," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 1–20.
- [316] T.-Y. Pan, Q. Liu, W.-L. Chao, and B. Price, "Towards open-world segmentation of parts," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 15392–15401.
- [317] X. Zhao et al., "An open and comprehensive pipeline for unified object grounding and detection," 2024, *arXiv:2401.02361*.
- [318] F. Li et al., "Semantic-SAM: Segment and recognize anything at any granularity," 2023, *arXiv:2307.04767*.
- [319] H. Yuan, X. Li, C. Zhou, Y. Li, K. Chen, and C. C. Loy, "Open-vocabulary SAM: Segment and recognize twenty-thousand classes interactively," 2024, *arXiv:2304.04704*.
- [320] H. Zhou et al., "Rethinking evaluation metrics of open-vocabulary segmentation," 2023, *arXiv:2311.03352*.
- [321] Z. Fang, X. Li, X. Li, J. M. Buhmann, C. C. Loy, and M. Liu, "Explore in-context learning for 3D point cloud understanding," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 1–14.
- [322] X. Wang, W. Wang, Y. Cao, C. Shen, and T. Huang, "Images speak in images: A generalist painter for in-context visual learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 6830–6839.
- [323] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 1–25.