

UniFS: Universal Few-shot Instance Perception with Point Representations

Sheng Jin^{1,4*}, Ruijie Yao^{2,4*}, Lumin Xu³, Wentao Liu⁴, Chen Qian⁴,
Ji Wu², and Ping Luo^{1,5}

¹ The University of Hong Kong ² Tsinghua University ³ The Chinese University
of Hong Kong ⁴ SenseTime Research and Tetras.AI ⁵ Shanghai AI Laboratory
js20@connect.hku.hk, yrj21@mails.tsinghua.edu.cn

Abstract. Instance perception tasks (object detection, instance segmentation, pose estimation, counting) play a key role in industrial applications of visual models. As supervised learning methods suffer from high labeling cost, few-shot learning methods which effectively learn from a limited number of labeled examples are desired. Existing few-shot learning methods primarily focus on a restricted set of tasks, presumably due to the challenges involved in designing a generic model capable of representing diverse tasks in a unified manner. In this paper, we propose UniFS, a universal few-shot instance perception model that unifies a wide range of instance perception tasks by reformulating them into a dynamic point representation learning framework. Additionally, we propose Structure-Aware Point Learning (SAPL) to exploit the higher-order structural relationship among points to further enhance representation learning. Our approach makes minimal assumptions about the tasks, yet it achieves competitive results compared to highly specialized and well optimized specialist models. Codes and data are available at <https://github.com/jin-s13/UniFS>.

Keywords: Few-Shot Learning · Unified Vision Model

1 Introduction

The field of instance perception, also known as object-centric learning, has made significant progress in recent years, with notable advancements. This paper primarily concentrates on four fundamental sub-tasks: object detection, instance segmentation, pose estimation¹, and object counting. Traditional approaches in this domain typically employ supervised learning methods, which necessitate enormous and costly annotated data, resulting in high costs. In contrast, there is a growing interest in few-shot learning methods that can leverage prior knowledge to handle novel concepts using only a small number of labeled samples.

Existing few-shot learning methods for instance perception are typically developed for a single sub-task or a subset of sub-tasks. Unifying different instance

* Equal contribution. ✉ Corresponding authors.

¹ In this paper, we use “keypoint localization” and “pose estimation” interchangeably.

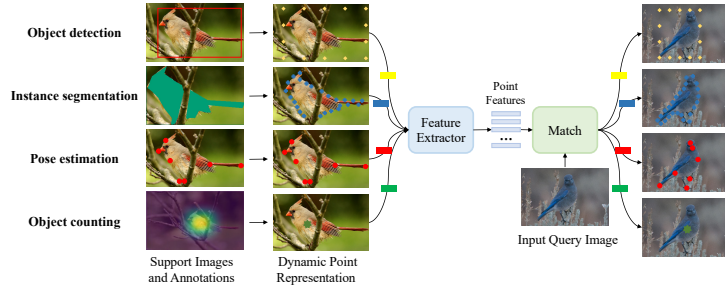


Fig. 1: UniFS utilizes a dynamic point representation learning paradigm to merge various task output spaces into a set of multiple points. With the provision of few-shot support images and annotated points from different tasks, UniFS can seamlessly produce corresponding point outputs on the query set.

perception tasks into a general model meets significant challenges due to variations in data sources, feature granularity, and output structures. (a) From a data perspective, images in different instance perception tasks and datasets exhibit diverse characteristics and inherent scale variances. For example, object detection datasets [37] often consist of scene images with multiple interacting objects, while pose estimation datasets [67] typically contain center-cropped images featuring a single object. (b) Regarding feature learning, different tasks require the extraction of relevant features at various levels of granularity. For instance, object detection prioritizes global semantic features, instance segmentation relies on fine-grained semantic features, and pose estimation necessitates both fine-grained semantic and localization features. (c) From the perspective of output, the outputs of different instance perception tasks have distinct structures. For example, the object detection task uses the coordinates, width and height of bounding box, while the pose estimation task employs Gaussian heatmaps for keypoint localization. These challenges hamper the development of multi-task few-shot learning methods that can comprehensively address a wide range of instance perception tasks in a unified manner.

Nevertheless, building a powerful generalist model that learns to solve multiple tasks [31] from only a few scraps of information is an ultimate goal in the AI community. Current approaches often adopt task-specific formulations, which involve separate learnable components and training objectives for each task. Such designs trend to be complicated and less scalable as the number of tasks increases. Although there have been recent works on generalist few-shot learners [1, 27, 48], they primarily focus on dense map prediction tasks, thus, cannot distinguish individual object instances. The problem of universal few-shot learning for instance perception tasks remains largely unexplored.

In this paper, we propose UniFS, that is a universal model with a task-agnostic structure designed to unify a wide range of few-shot instance perception tasks. As shown in Fig. 1, the key idea is to reformulate diverse instance perception tasks into a generalized point representation learning paradigm. By providing user-defined point annotations in the support set, the model is prompted to

learn task-specific points for each query image. For example, in object detection, points are placed along the bounding box edges, while for instance segmentation, points are distributed along the instance mask contour. This unified formulation of point learning brings several benefits: (1) **Task-agnostic architecture:** The model uses a unified representation that supports diverse task types and breaks the scaffolds of complex task-specific customization. (2) **Parameter efficiency:** Parameters are fully shared among tasks, reducing redundant computation when handling multiple tasks simultaneously. (3) **Task generalization:** The model can generalize to new tasks with minimal knowledge of the underlying mechanism. (4) **Knowledge sharing:** Joint training on diverse tasks enables knowledge sharing among tasks, enhancing overall performance.

Furthermore, we propose a novel structure-aware point learning objective to address the limitations of traditional L1/L2 losses in point learning. By considering the relationship between the center point and its neighboring points, our objective leverages contextual information and spatial dependencies within the local region. This leads to more accurate predictions by better capturing spatial relationships and structural patterns among points. Additionally, the inclusion of structural relationships provides regularization, enabling the model to handle noisy inputs and outliers more effectively, resulting in robust and stable training.

We address the lack of a comprehensive benchmark for evaluating multiple instance perception tasks by introducing the COCO-UniFS dataset. Derived and re-organized from existing datasets, COCO-UniFS encompasses tasks including object detection, instance segmentation, pose estimation, and object counting. This unified dataset serves as a platform for developing and assessing multi-task few-shot learning methods, facilitating model evaluation and generalization across various instance perception tasks. Our approach achieves competitive results on COCO-UniFS benchmark, with minimal task assumptions, compared to highly specialized task-specific models.

Our work makes the following key contributions:

- **Pioneering Universal Few-Shot Instance Perception:** To the best of our knowledge, we are the first to tackle the problem of universal few-shot learning for instance perception tasks. By introducing the general and unified vision problem and benchmark, we hope to inspire the research community to explore this promising direction.
- **A Simple and Effective Baseline:** We propose UniFS, a simple yet effective baseline approach that unifies multiple instance perception tasks under a few-shot point representation learning paradigm. It employs a unified architecture for all tasks and makes minimal assumptions about the specific characteristics of each task.
- **Structure-Aware Point Learning:** We introduce a novel structure-aware point learning objective that captures the higher-order structural relationships among points, enhancing the model’s ability to understand and exploit the spatial dependencies among points.

2 Related Works

2.1 Few-Shot Learning

Few-Shot Learning Methods. Few-shot learning aims to imitate humans’ generalization ability that can efficiently learn novel concepts from a few examples. It has gained attention in recent years, with approaches falling into meta-learning, metric-learning, and transfer learning (or finetuning) paradigms. Meta-learning based methods [10, 15, 20, 26, 33, 34, 59, 63, 69, 79] intent to learn efficient parameter updating rules or better parameter initialization strategies. Metric-learning based methods [28, 46, 50, 54] focus on obtaining a generalizable embedding metric space for pairwise similarity comparisons. Transfer learning based methods [2, 13, 42, 47, 57, 61] leverage pre-training on base classes and finetuning on novel classes. Our UniFS belongs to the transfer learning category.

Few-Shot Learning Tasks. Few-shot learning has extended beyond classification to applications like object detection [2, 10, 26, 33, 34, 42, 47, 59, 63, 79], instance segmentation [11–13, 39, 69], pose estimation [18, 67], and object counting [43, 73]. However, existing approaches tend to be task-specific, resulting in diverse models and datasets. These specialized methods hinder integration into a unified framework. For example, DeFRCN [42] uses a two-stage detector for few-shot object detection, while SAFECOUNT [73] employs a single-stage approach for few-shot counting. Although some methods handle pairs of tasks like object detection and instance segmentation [12, 13, 69], they rely on task-specific components. In contrast, our UniFS model minimizes task-specific designs through a unified architecture. While recent efforts focus on generalist few-shot learners for tasks like semantic segmentation [1, 27, 48], they overlook instance perception tasks. To the best of our knowledge, universal few-shot learning for instance perception tasks remains unexplored, which is the main focus of this paper.

Few-Shot Learning Benchmarks. Most existing few-shot learning benchmarks are designed for specific computer vision tasks, such as image classification [30, 54], semantic segmentation [36], object detection [37], pose estimation [67], and object counting [43]. Due to the lack of a unified benchmark dataset, it is hard to evaluate and develop generalist few-shot learners capable of solving diverse vision tasks. In this paper, we introduce the COCO-UniFS benchmark dataset, which extends existing datasets and covers multi-task few-shot learning of instance perception tasks. Existing datasets for multi-task few-shot learning, such as Taskonomy [27, 74] and Meta Omnium [1], do not target at instance perception tasks like object detection and instance segmentation. Moreover, these datasets are often combinations of several sub-datasets, making it challenging to analyze the impact factors involved.

2.2 Unified Vision Models

Unified Learning Paradigms. Several works have focused on general representation learning for foundation vision models [6, 7, 14, 19, 45, 51, 55, 75, 78, 80]. However, these models typically require costly task-specific fine-tuning to achieve

optimal performance on specific downstream tasks. In contrast, UniFS differs from them, as it aims to address multi-task few-shot learning.

Unified Task Representations. Efforts have been made to unify task output representations for multiple vision tasks. **Dense representations.** Some studies [27, 58] employ dense image-like representations, as seen in Painter [58], facilitating tasks like panoptic segmentation and depth estimation. However, they lack versatility in addressing core vision tasks such as object detection and instance segmentation. **Text representations.** Other approaches [4, 5, 38, 56, 71] leverage text representations inspired by NLP’s sequence-to-sequence modeling. Despite their potential, they suffer from slow inference speeds and overlook category-agnostic multi-instance pose estimation. **Point representations** are prevalent in various vision tasks, including object detection [8, 32], segmentation [64], and pose estimation [40, 52, 60]. While some works attempt to unify tasks using point representations, they often train and evaluate models in a single-task manner. In this paper, we propose a universal generalist model unifying few-shot instance perception tasks through point representations.

Unified Model Architectures. Designing a unified model capable of handling different tasks in a unified manner is considered a step toward general intelligence. Some attempts [16, 22, 25, 66, 68] have been made to design unified model architectures for visual tasks, but they often rely on task-specific heads to perform different tasks, and the interactions among different tasks are less studied. UniHCP [7] minimizes task-specific parameters for knowledge sharing but still requires task-specific queries and a task-guided interpreter. Additionally, it is designed specifically for human-centric tasks. In our work, we take a further step by unifying few-shot instance perception tasks with a fully shared model architecture, without any task-specific designs.

3 COCO-UniFS Benchmark

This work targets at universal few-shot instance perception that uniformly handles multiple tasks for arbitrary categories by providing a few visual examples and the corresponding annotations from that category. In order to address the task of universal few-shot instance perception, we introduce the COCO-UniFS benchmark. This benchmark is built upon several existing datasets, including MSCOCO [37] and MISC [49], aiming to provide a high-quality and high-diversity dataset for model training and evaluating this challenging task. We have made efforts to ensure unification for dataset annotation, dataset split, and evaluation protocols, enabling fair comparisons between models and methods.

3.1 Task Definition

We define the task of universal few-shot instance perception as follows: given a few reference images of a novel object category and the corresponding annotations (*e.g.* box, mask, keypoint, or center point), this task requires the model to find corresponding labels of all the instances belonging to this category in the

query image. The reference image shows a cropped single object instance of the category that is to be analyzed, while the query image shows an entire visual scene containing many objects. Note that we assume that the task information (*e.g.* detection or segmentation) is implicitly embedded in the given annotations, and is not explicitly provided. This makes it possible to define your own “novel” task by offering varying annotations in the support images.

3.2 Data Annotation

The COCO-UniFS benchmark provides dense annotations for four fundamental few-shot computer vision tasks: object detection, instance segmentation, pose estimation, and object counting. The annotations for object detection and instance segmentation are directly taken from the MSCOCO dataset [37], which provides bounding box and per-instance segmentation mask annotations for 80 object categories. For pose estimation, we extend the MSCOCO dataset by adding instance-level keypoint annotations for 34 object categories from the MISC dataset [49]. The MISC dataset was originally designed for multi-instance semantic correspondence, and we adapted it to fit the few-shot pose estimation task. The keypoints are defined on the 3D prototype model, encompassing contour points (*e.g.* head top), skeletal joints (*e.g.* wrists), and other distinctive feature points (*e.g.* eyes). We manually add keypoint semantic descriptions with their symmetry relationship for each object category. Objects with ambiguous or inconsistent keypoint definitions, insufficient keypoint numbers, or large morphology variations are filtered out to ensure annotation accuracy. Note that the keypoint definition (even the number of keypoints) varies from category to category. For object counting, each object is annotated with a point at the center of its bounding box, following the method in prior datasets [43].

3.3 Dataset Split

To evaluate few-shot generalization to unseen categories, the available object categories should be split into a training set and a non-overlapping test set. Previously, due to the independent development of various few-shot learning tasks, different conventions exist for dataset splits across research areas. To facilitate multi-task few-shot learning and ensure consistency, we unify the dataset split for all tasks in the COCO-UniFS benchmark. The dataset split follows the convention used in COCO-inst [37] for few-shot object detection and instance segmentation. The 20 overlapping categories with PASCAL VOC 2012 are novel classes, while the remaining 60 categories are base classes.

3.4 Evaluation Scenarios and Metrics

Evaluation Scenarios. We establish two evaluation scenarios for COCO-UniFS: (1) **Seen-task evaluation:** In the seen-task few-shot evaluation, the evaluation tasks are seen and trained during the base training phase. This evaluation scenario focuses on the tasks of object detection, instance segmentation, and pose

estimation. (2) **Unseen-task evaluation:** This evaluation scenario is designed to evaluate novel task generalization, specifically for the task of object counting. Note that the evaluation categories are unseen in both evaluation scenarios.

Evaluation Metric. For object detection, instance segmentation, and pose estimation, we use the COCO official metric and report the Average Precision (AP) with IoU (Intersection over Union) threshold ranging from 0.5 to 0.95 on the novel classes, where box IoU, mask IoU, and Object Keypoint Similarity (OKS) based IoU are applied respectively. Note that for pose estimation, we only consider the keypoints that exist both in at least one support image and in the query image during metric calculation. For object counting, we use Mean Squared Error (MSE) as the main evaluation metric. Previous works [43, 70] calculate the average over the images, instead of categories. To make the evaluation of object counting consistent with other tasks, we opt to calculate the mean MSE by averaging the MSE values over all the evaluation classes.

3.5 Uniqueness

The COCO-UniFS dataset stands out due to several unique properties: (1) **Comprehensiveness:** The COCO-UniFS dataset covers four fundamental instance perception tasks: object detection, instance segmentation, pose estimation, and object counting. With 100,000+ valid images, it features diverse lighting, resolutions, and object categories. (2) **Annotation Completeness:** The COCO-UniFS dataset differs from previous benchmarks by offering complete annotations for all tasks, eliminating the need for multiple sub-datasets. This comprehensive annotation simplifies algorithm development for handling multiple tasks concurrently. (3) **Domain Uniformity:** Existing multi-task benchmark datasets such as Meta Omnium [1] and Taskonomy [27, 74] are composed of multiple sub-datasets from different visual domains, ranging from natural images to medical and industrial images. This introduces additional domain factors that affect model learning and evaluation. In contrast, the COCO-UniFS dataset maintains domain uniformity, where all images and tasks are within the same data domain. This simplifies the development and analysis of multi-task few-shot learners, allowing researchers to concentrate on instance perception challenges without the interference of domain shifts. (4) **Unseen Task Generalization:** The COCO-UniFS benchmark introduces an unseen-task evaluation scenario for object counting, enabling researchers to assess model generalization to new tasks unseen during training. These features make the COCO-UniFS dataset a valuable asset for advancing multi-task few-shot instance perception research.

4 Methodology

4.1 Overview

Our goal is to develop a single unified model capable of addressing a diverse set of few-shot learning tasks without the need for task-specific algorithm customizations (*e.g.* output representations and loss functions) or model parameters

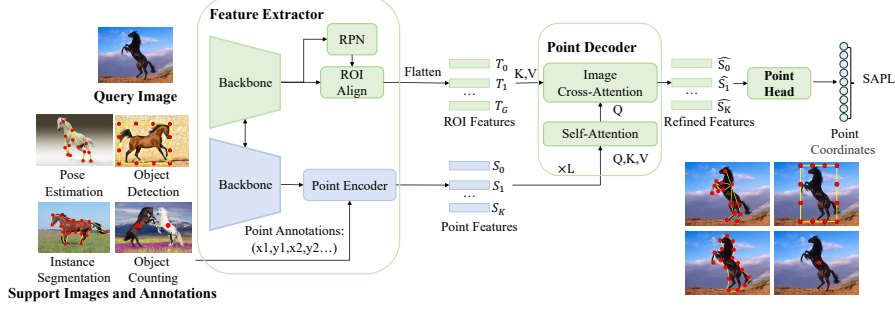


Fig. 2: Overview of UniFS. UniFS adopts a dynamic point representation learning paradigm to unify different task output spaces into a set of multiple points.

(*e.g.* separate architectures and task heads). The overall pipeline of our proposed UniFS is presented in Fig. 2. Our focus lies in unifying task representation, model architecture, and learning paradigms.

4.2 Unified Task Representation

Traditionally, instance perception tasks, such as image classification [30, 54, 72], object detection [77, 78], instance segmentation [16, 35, 64], and pose estimation [21, 23, 24, 62, 65, 66], each have their distinct representations, posing challenges for simultaneous handling. To achieve unified instance perception, we propose a point-based formulation for various tasks, offering a unified representation as follows: **Object Detection.** Object detection task is a fundamental computer vision task in which the goal is to obtain the bounding box level location of objects. we utilize 16 points to depict a bounding box, which are uniformly sampled across the box. **Instance Segmentation.** Instance segmentation requires pixel-wise masks, so we employ 32 points to outline the boundary of each target instance. In the COCO dataset, segmentation masks have been annotated using either clockwise or counterclockwise polygons. Following the methodology outlined in Deep Snake [41], we uniformly add or remove points from the initial contour until the desired number of points is reached. We ensure a consistent clockwise orientation and designate the point situated at the leftmost edge of the contour as the starting point. **Pose Estimation.** Pose estimation detects the coordinates of semantic keypoints. Our point-based representation naturally aligns with this task, with each point corresponding to a target keypoint. This facilitates estimating keypoints based on definitions in the support image, accommodating various object categories. **Object Counting.** UniFS adopts a center point-based approach by predicting the center point of each target object. The center point is defined as the midpoint of the object’s bounding box. By counting the center points for a specific category, we can determine the number of objects in that category. A detailed ablation of the number of points is provided in supplementary Table S1.

4.3 Unified Model Architecture

The UniFS model consists of three main components: the feature extractor, the point decoder, and the point head. These components work together to handle different instance perception tasks within a unified framework.

Feature Extractor. The feature extractor is proposed to extract features from the support and query images. We utilize an ImageNet [29] pre-trained ResNet-101 [17] backbone as the feature extractor, which is shared between both the support image and the query image. Given the support image I_S , the feature extractor \mathcal{F} generates the support image features $F_S = \mathcal{F}(I_S)$. Sinusoidal Position embeddings [53] are then added to the support image features. Given ground-truth points $[(x_1, y_1), (x_2, y_2), \dots, (x_K, y_K)]$ for the support sample, we obtain the point features $\{S_i\}_{i \in [1, K]}$ by extracting the features via bi-linear interpolation from the corresponding location of the support image features F_S . For the query image I_Q , the shared feature extractor is applied to extract the query image features $F_Q = \mathcal{F}(I_Q)$. Then, a Region Proposal Network (RPN) generates candidate regions, and RoI Align is applied to pool each region proposal feature into a sequence of flattened query object features $\{T_i\}_{i \in [1, G]}$, where G is the number of flattened RoI features for each given region proposal. Following Faster R-CNN [44], the MLP-based object classification head is applied to the proposal features that are obtained by global pooling of region proposal features to predict categories of the objects.

Point Decoder. To capture the correlative dependence between the support point set and the query point set, we introduce transformer-based attention mechanisms to enhance the representations of point features $\{S_i\}_{i \in [1, K]}$. The point decoder module consists of L transformer decoder layers. Each layer contains two main components: Self-Attention and Cross-Attention. *Self-Attention.* The self-attention layer [53] enables the exchange of information among the query points, and obtains the refined point features $\{S'_i\}_{i \in [1, K]}$. By allowing the query point features to interact with each other and aggregating these interactions using attention weights, the contextual meaning of the current point sequence, particularly the information relevant to the task, is better modeled. *Cross-Attention.* Then point features interact with the query image features to align the feature representations and mitigate the representation gap. In this case, a cross-attention layer [3] is used to aggregate useful information between the support point and query image. The support point features $\{S'_i\}_{i \in [1, K]}$ are used as queries, while the flattened query RoI features $\{T_i\}_{i \in [1, G]}$ of each region proposal serve as the keys and values. The channel dimension of the query RoI features is reduced to match the channel dimension of the support point features, and sinusoidal position embeddings [53] are added to the query RoI features. A feed-forward network (FFN) and layer normalization (LN) are also included following common practice. Finally, this process yields an enhanced group of point features $\{\hat{S}_i\}_{i \in [1, K]}$ for each set of RoI features. The computation in point decoder can be formulated as follows:

$$S'_i = \text{LN}(\text{SelfAttn}(S_i)), \quad \hat{S}_i = \text{LN}(\text{FFN}(\text{LN}(\text{CrossAttn}(S'_i, T_i)))) \quad (1)$$

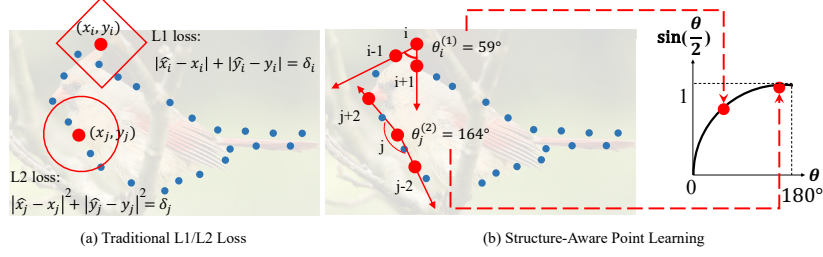


Fig. 3: Structure-Aware Point Learning (SAPL): (a) Traditional L1/L2 loss focuses on individual point error. (b) SAPL integrates structural relationships among points by supervising the angle between each point and its neighboring points.

Point Head. Given the refined point embeddings $\{\hat{S}_i\}$, we utilize a multi-layer perceptron (MLP) to transform them into the point coordinates. Rather than directly predicting absolute coordinates, we infer the offset $(\Delta x_i, \Delta y_i)$ from the center of RPN anchors. Finally, the scattered point coordinates (P_{xi}, P_{yi}) are generated by adding the offset prediction to the proposal anchor bounding box $A : (A_{cx}, A_{cy}, A_w, A_h)$:

$$P_{xi} = A_{cx} + \Delta x_i \times A_w, \quad P_{yi} = A_{cy} + \Delta y_i \times A_h. \quad (2)$$

4.4 Unified Learning Paradigm

In contrast to previous approaches that use different training objectives and loss functions for different tasks, we propose a unified learning paradigm for all available tasks. This approach enables us to train the model jointly on multiple tasks, leading to improved efficiency and performance.

Structure-Aware Point Learning (SAPL). Fig. 3 illustrates the differentiation between conventional point-level loss and our structure-aware point learning loss. While L1/L2 loss functions focus solely on individual point-level error, our proposed Structure-Aware Point Learning (SAPL) incorporates the inter-dependency between a specific point and its neighboring points, thus capturing the intrinsic structural information within the point sequence. As depicted in Fig. 3(a), while predicted points on the diamond shape result in identical L1 losses, those on the circle yield the same L2 losses. This introduces ambiguity in the model’s learning process. However, integrating SAPL helps mitigate such ambiguity, thus enhancing the representation learning.

Specifically, we additionally impose constraints on the structural information by supervising the angles $\theta_i^{(n)}$ between each point and its adjacent n -hop points. Here, $\theta_i^{(n)}$ represents the angle formed by points $i-n$, i , and $i+n$. As illustrated in Fig. 3(b), the angle $\theta_i^{(1)}$ formed by points $i-1$, i , and $i+1$ for each point i , while $\theta_j^{(2)}$ demonstrates the 2-hop formulation. Incorporating higher-order information supervision helps reduce errors from point annotation noise, thereby enhancing

robustness. Additionally, we integrate a sine function $\sin(\frac{\theta}{2})$ into the loss function to amplify gradients at sharp peaks and attenuate gradients in flat regions. This enables the effective capture of detailed shape information. Formally, the SAPL loss, integrating angles from 1-hop to N-hop, is defined as:

$$L_{SAPL}(\hat{P}_i, P_i) = \frac{1}{N} \sum_{n=1}^N L_1(\sin(\frac{\hat{\theta}_i^{(n)}}{2}), \sin(\frac{\theta_i^{(n)}}{2})) \quad (3)$$

where \hat{P}_i and P_i represent the predicted and ground-truth points, respectively. $\hat{\theta}_i^{(n)}$ and $\theta_i^{(n)}$ denote their respective n-hop angles. The final loss L_{point} is computed as the sum of the L1 loss and SAPL loss for each point, given by:

$$L_{point}(\hat{P}_i, P_i) = |\hat{P}_i - P_i| + L_{SAPL}(\hat{P}_i, P_i) \quad (4)$$

5 Experiments

5.1 Implementation Details

Experiments are conducted using Detectron2 (Wu et al., 2019). We employ the standard Faster R-CNN (Ren et al., 2015) with ResNet-101 (He et al., 2016) backbone. We set the number of point decoders L to 2 and configure the SAPL with an N-hop value of $N = 2$. For model training, we employed a two-stage transfer-learning approach: base class training and novel class fine-tuning. During base class training, models were jointly trained on object detection, instance segmentation, and pose estimation tasks simultaneously. Notably, the task of object counting was held out and not trained during this stage. Each batch randomly sampled different tasks. Due to the stochastic nature of task selection within each batch, employing a larger batch size was crucial for ensuring stable training. Thus, we set the batch size to 32 (4 per GPU * 8 GPUs), with a learning rate of 0.028, and utilized Stochastic Gradient Descent (SGD) to optimize our network, with a maximum of 55,000 iterations. For novel class fine-tuning, our model was fine-tuned with K-shots for every class, and the learning rate was reduced to 0.01 during few-shot fine-tuning. Following the previous work [13], all experimental results are averaged over 10 seeds.

5.2 Baseline Methods.

We compare our method to task-specific and multi-task few-shot baseline models. For few-shot object detection, we compare with several representative methods [2, 15, 57, 69, 76] including meta-learning based and metric-learning based methods. For few-shot instance segmentation, we compare with a number of state-of-the-art models [12, 12, 13, 42, 69]. Especially, we also compare with the DETR-based models (Meta-DETR [76]) beyond the Faster R-CNN framework. For few-shot object counting, we compare with both heatmap-based methods (CFOCNet [70]) and detection-based methods (MPSR [61] and FSDetView [63]).

Please note that these models are trained with object counting tasks in the base-training stage, while our model performs unseen-task generalization. As no prior methods are developed for universal few-shot instance-perception tasks, we also adapt Mask-RCNN [16], a well-known multi-task model, to our setup. Specifically, we add a heatmap based pose estimation head at the ROI of the Mask-RCNN. The baseline model initially trains on a mix of all base categories from the training set. During testing, it fine-tunes on support images of the new category. To address varying numbers of keypoints, the model outputs the maximum number across all categories, such as 52 on the COCO-UniFS dataset. Only valid keypoints for each category are supervised during training.

5.3 Experiments on COCO-UniFS benchmark

Results on COCO-UniFS. We compare our method to task-specific and multi-task few-shot models on the COCO-UniFS dataset in Table 1. All compared models adopt ResNet-101 [17] as the backbone. We find that Our model significantly outperforms the multi-task generalist baseline, *i.e.* Mask-RCNN-ft. Our approach achieves competitive results against state-of-the-art task-specific models, with only a slight disadvantage in 5-shot segmentation. It is worth noting that our method has a better advantage, especially in extremely low-shot ($K=1$ shot), and thus it is very suitable for a few-shot scenario. **Unseen Task Generalization.** Our model also shows strong generalization capability when evaluated on the task of few-shot object counting (or center point localization). Please note that these baseline models are trained with object counting tasks in the base-training stage, while our model performs unseen-task generalization. Nevertheless, our UniFS outperforms them.

5.4 Experimental Analysis

Experiments on PASCAL-5ⁱ benchmark. PASCAL-5ⁱ dataset [9] is a well-known benchmark dataset for evaluating few-shot object detection. The dataset contains 20 categories, which are randomly divided into three different splits (Novel Set #1, #2, #3), each consisting of 15 base classes and 5 novel classes. Following previous works [13, 42], training is conducted on the PASCAL VOC 2007+2012 train/val sets, and evaluation is conducted on the PASCAL VOC 2007 test set. The Average Precision results for novel classes (AP_{50}) in K-shot settings are reported. Table 2 summarises our results and compares them with the current state-of-the-art few-shot object detectors on PASCAL VOC. We observe that the proposed method outperforms state-of-the-art specialist models.

Effect of multi-task learning. In Table 3, we observed that co-learning with multiple tasks leads to improved overall performance. This enhancement can be attributed to the inter-task synergy that arises from jointly training different instance perception tasks.

Table 1: Comparisons with task-specific and multi-task few-shot learners on the COCO-UniFS val set. We report AP for object detection (Det.), instance segmentation (Seg.), and pose estimation (Kpt.), and report MSE for object counting (Cnt.). The results are averaged over all 10 seeds for fair comparisons. \uparrow means the higher the better, while \downarrow means the lower the better. \dagger means seen-task evaluation, while \ddagger means unseen-task generalization.

Model	Det. AP \uparrow		Seg. AP \uparrow		Kpt. AP \uparrow		Cnt. MSE \downarrow	
	K=1	K=5	K=1	K=5	K=1	K=5	K=1	K=5
FRCN-ft [69]	1.0	4.0	\times	\times	\times	\times	\times	\times
TFA [57]	4.4	7.7	\times	\times	\times	\times	\times	\times
FADI [2]	5.7	10.1	\times	\times	\times	\times	\times	\times
FCT [15]	5.1	12.0	\times	\times	\times	\times	\times	\times
Meta R-CNN [69]	-	3.5	-	2.8	\times	\times	\times	\times
MTFA [12]	2.5	6.6	2.7	6.6	\times	\times	\times	\times
iMTFA [12]	3.3	6.2	2.8	5.2	\times	\times	\times	\times
Mask-DeFRCN [42]	7.5	15.4	6.7	12.7	\times	\times	\times	\times
DCFS [13]	8.1	16.4	7.2	13.5	\times	\times	\times	\times
Meta-DETR [76]	-	15.3	-	8.1	\times	\times	\times	\times
DCFS-pose [13]	11.0	17.1	\times	\times	9.9	21.1	\times	\times
CFOCNet [70]	\times	\times	\times	\times	\times	\times	4.43 \dagger	2.26 \dagger
MPSR [61]	5.1	8.7	\times	\times	\times	\times	1.42 \dagger	1.40 \dagger
FSDetView [63]	4.5	10.7	\times	\times	\times	\times	1.42 \dagger	1.42 \dagger
Mask-RCNN-ft [16]	2.4	6.9	2.0	5.5	2.3	6.7	1.48 \dagger	1.45 \dagger
UniFS (Ours)	12.7	18.2	8.6	11.5	12.2	22.1	1.38\ddagger	1.32\ddagger

Table 2: Few-shot object detection results on PASCAL-5ⁱ. We evaluate model performance (AP_{50}) on three different splits. The results are averaged over 10 seeds.

Method	Shot	YOLO-ft [26]	Meta R-CNN [69]	TFA [57]	TIP [33]	DeFRCN [42]	DCFS [13]	UniFS
Novel Set #1	K=1	6.6	19.9	39.8	27.7	46.2	46.2	47.2
	K=5	24.8	45.7	55.7	50.2	62.4	62.9	63.6
Novel Set #2	K=1	12.5	10.4	23.5	22.7	32.6	32.6	33.4
	K=5	16.1	34.8	35.1	40.9	48.3	47.9	49.3
Novel Set #3	K=1	13	14.3	30.8	21.7	39.8	40.3	41.2
	K=5	32.2	41.2	49.5	44.5	56.1	56.9	58.5

Effect of Structure-Aware Point Learning (SAPL). In Table 4, we analyze the effect of SAPL loss, and compare with using L1/L2 loss only. Interestingly, L1 loss is superior to L2 loss on segmentation and keypoint, and achieves similar performance on detection. We also find that adopting SAPL could bring increased performance with accelerated learning convergence by considering the structural information. For example, the segmentation AP increases from 7.2 AP to 8.4 AP. As 1-hop SAPL is easily affected by point noises, incorporating 2-hop SAPL is also beneficial, further increasing the segmentation AP to 8.6 AP. However, 3-hop and 4-hop structural cues are too smooth to capture detailed shape information. Therefore, we choose to use $N = 2$ in our implementation.


Qualitative Analysis. In Fig. 4, we qualitatively evaluate the generalization ability of UniFS to novel categories. Our model can follow instructions to perform different tasks with different point annotations as the input.

Table 3: Effect of multi-task learning on the COCO-UniFS val set.

Model	Det. AP \uparrow		Seg. AP \uparrow		Kpt. AP \uparrow	
	K=1	K=5	K=1	K=5	K=1	K=5
Det.	12.2	17.9	\times	\times	\times	\times
Det.+Seg.	12.6	17.6	8.5	11.2	\times	\times
Det.+Seg.+Kpt.	12.7	18.2	8.6	11.5	12.2	22.1

Table 4: Effect of Structure-Aware Point Learning (SAPL).

Model	Det. AP \uparrow		Seg. AP \uparrow		Kpt. AP \uparrow	
	K=1	K=5	K=1	K=5	K=1	K=5
L2 only	10.9	16.1	6.1	7.7	9.1	19.5
L1 only	10.6	16.2	7.2	8.7	12.0	21.0
L1+1-hop SAPL	12.6	17.9	8.4	11.3	12.3	21.8
L1+2-hop SAPL	12.7	18.2	8.6	11.5	12.2	22.1
L1+3-hop SAPL	12.6	17.7	8.2	11.0	12.4	21.6
L1+4-hop SAPL	12.7	17.8	8.4	11.3	12.2	20.9

**Fig. 4:** Visualization of tasks in a 5-shot scenario with support images of new categories and point annotations, leading to corresponding point outputs on the query set.

6 Conclusion

Few-shot learning has been extensively studied to mimic the ability of humans to generalize to new tasks with only a few demonstrations. However, few-shot learning of multi-task instance perception is under explored. This paper introduces UniFS, a universal few-shot instance perception model that handles a diverse set of instance perception tasks in a unified point representation learning paradigm. To facilitate model training and evaluation, we also extend existing datasets and construct the COCO-UniFS benchmark for multi-task few-shot instance perception. It is worth noting that our proposed generalist model is achieved with a fully shared model architecture and minimum task-specific designs, which is widely considered to be a unique feature for general artificial intelligence. We hope for broad applicability and further research on more unified vision models.

Limitations and future work. Since point-based method may present challenges for segmentation due to sampling errors, there is room for improvement through further optimization. This work primarily focuses on point representation for task unification, and the image-level classification aspect of few-shot learner is not specifically improved. While our work focuses on 2D tasks, we remain optimistic about its potential for tasks such as 3D or sequential input. Future research endeavors should focus on exploring more comprehensive multi-task perception in few-shot scenarios.

Acknowledgement. This paper is partially supported by the National Key R&D Program of China No.2022ZD0161000 and the General Research Fund of Hong Kong No.17200622 and 17209324.

References

1. Bohdal, O., Tian, Y., Zong, Y., Chavhan, R., Li, D., Gouk, H., Guo, L., Hospedales, T.: Meta omnium: A benchmark for general-purpose learning-to-learn. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 7693–7703 (2023)
2. Cao, Y., Wang, J., Jin, Y., Wu, T., Chen, K., Liu, Z., Lin, D.: Few-shot object detection via association and discrimination. *Adv. Neural Inform. Process. Syst.* **34**, 16570–16581 (2021)
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *Eur. Conf. Comput. Vis.* pp. 213–229. Springer (2020)
4. Chen, T., Saxena, S., Li, L., Fleet, D.J., Hinton, G.: Pix2seq: A language modeling framework for object detection. *Int. Conf. Learn. Represent.* (2022)
5. Chen, T., Saxena, S., Li, L., Lin, T.Y., Fleet, D.J., Hinton, G.E.: A unified sequence interface for vision tasks. *Adv. Neural Inform. Process. Syst.* **35**, 31333–31346 (2022)
6. Chen, W., Xu, X., Jia, J., Luo, H., Wang, Y., Wang, F., Jin, R., Sun, X.: Beyond appearance: a semantic controllable self-supervised learning framework for human-centric visual tasks. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 15050–15061 (2023)
7. Ci, Y., Wang, Y., Chen, M., Tang, S., Bai, L., Zhu, F., Zhao, R., Yu, F., Qi, D., Ouyang, W.: Unihcp: A unified model for human-centric perceptions. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 17840–17852 (2023)
8. Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., Tian, Q.: Centernet: Keypoint triplets for object detection. In: *Int. Conf. Comput. Vis.* pp. 6569–6578 (2019)
9. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.* **88**(2), 303–338 (2010)
10. Fan, Q., Zhuo, W., Tang, C.K., Tai, Y.W.: Few-shot object detection with attention-RPN and multi-relation detector. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2020)
11. Fan, Z., Yu, J.G., Liang, Z., Ou, J., Gao, C., Xia, G.S., Li, Y.: Fgn: Fully guided network for few-shot instance segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 9172–9181 (2020)
12. Ganea, D.A., Boom, B., Poppe, R.: Incremental few-shot instance segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1185–1194 (2021)
13. Gao, B.B., Chen, X., Huang, Z., Nie, C., Liu, J., Lai, J., Jiang, G., Wang, X., Wang, C.: Decoupling classifier for boosting few-shot object detection and instance segmentation. *Adv. Neural Inform. Process. Syst.* **35**, 18640–18652 (2022)
14. Ghiasi, G., Zoph, B., Cubuk, E.D., Le, Q.V., Lin, T.Y.: Multi-task self-training for learning general representations. In: *Int. Conf. Comput. Vis.* pp. 8856–8865 (2021)
15. Han, G., Ma, J., Huang, S., Chen, L., Chang, S.F.: Few-shot object detection with fully cross-transformer. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5321–5330 (2022)
16. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *Int. Conf. Comput. Vis.* pp. 2961–2969 (2017)

17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conf. Comput. Vis. Pattern Recog. (2016)
18. Hirschorn, O., Avidan, S.: Pose anything: A graph-based approach for category-agnostic pose estimation. arXiv preprint arXiv:2311.17891 (2023)
19. Hong, F., Pan, L., Cai, Z., Liu, Z.: Versatile multi-modal pre-training for human-centric perception. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 16156–16166 (2022)
20. Hu, H., Bai, S., Li, A., Cui, J., Wang, L.: Dense relation distillation with context-aware aggregation for few-shot object detection. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 10185–10194 (2021)
21. Jiang, W., Jin, S., Liu, W., Qian, C., Luo, P., Liu, S.: Posetrans: A simple yet effective pose transformation augmentation for human pose estimation. In: Eur. Conf. Comput. Vis. pp. 643–659 (2022)
22. Jin, S., Li, S., Li, T., Liu, W., Qian, C., Luo, P.: You only learn one query: Learning unified human query for single-stage multi-person multi-task human-centric perception. Eur. Conf. Comput. Vis. (2024)
23. Jin, S., Liu, W., Ouyang, W., Qian, C.: Multi-person articulated tracking with spatial and temporal embeddings. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5664–5673 (2019)
24. Jin, S., Liu, W., Xie, E., Wang, W., Qian, C., Ouyang, W., Luo, P.: Differentiable hierarchical graph grouping for multi-person pose estimation. In: Eur. Conf. Comput. Vis. pp. 718–734. Springer (2020)
25. Jin, S., Xu, L., Xu, J., Wang, C., Liu, W., Qian, C., Ouyang, W., Luo, P.: Whole-body human pose estimation in the wild. In: Eur. Conf. Comput. Vis. (2020)
26. Kang, B., Liu, Z., Wang, X., Yu, F., Feng, J., Darrell, T.: Few-shot object detection via feature reweighting. In: Int. Conf. Comput. Vis. (2019)
27. Kim, D., Kim, J., Cho, S., Luo, C., Hong, S.: Universal few-shot learning of dense prediction tasks with visual token matching. Int. Conf. Learn. Represent. (2023)
28. Koch, G., Zemel, R., Salakhutdinov, R., et al.: Siamese neural networks for one-shot image recognition. In: ICML deep learning workshop (2015)
29. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. Communications of the ACM (2017)
30. Lake, B.M., Salakhutdinov, R., Tenenbaum, J.B.: Human-level concept learning through probabilistic program induction. Science **350**(6266), 1332–1338 (2015)
31. Lake, B.M., Salakhutdinov, R.R., Tenenbaum, J.: One-shot learning by inverting a compositional causal process. In: Adv. Neural Inform. Process. Syst. (2013)
32. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: Eur. Conf. Comput. Vis. pp. 734–750 (2018)
33. Li, A., Li, Z.: Transformation invariant few-shot object detection. In: IEEE Conf. Comput. Vis. Pattern Recog. (2021)
34. Li, B., Yang, B., Liu, C., Liu, F., Ji, R., Ye, Q.: Beyond max-margin: Class margin equilibrium for few-shot object detection. In: IEEE Conf. Comput. Vis. Pattern Recog. (2021)
35. Li, F., Zhang, H., Xu, H., Liu, S., Zhang, L., Ni, L.M., Shum, H.Y.: Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3041–3050 (2023)
36. Li, X., Wei, T., Chen, Y.P., Tai, Y.W., Tang, C.K.: Fss-1000: A 1000-class dataset for few-shot segmentation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 2869–2878 (2020)

37. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: *Eur. Conf. Comput. Vis.* (2014)
38. Lu, J., Clark, C., Zellers, R., Mottaghi, R., Kembhavi, A.: Unified-io: A unified model for vision, language, and multi-modal tasks. *Int. Conf. Learn. Represent.* (2022)
39. Michaelis, C., Ustyuzhaninov, I., Bethge, M., Ecker, A.S.: One-shot instance segmentation. *arXiv preprint arXiv:1811.11507* (2018)
40. Nie, X., Feng, J., Zhang, J., Yan, S.: Single-stage multi-person pose machines. In: *Int. Conf. Comput. Vis.* pp. 6951–6960 (2019)
41. Peng, S., Jiang, W., Pi, H., Li, X., Bao, H., Zhou, X.: Deep snake for real-time instance segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 8533–8542 (2020)
42. Qiao, L., Zhao, Y., Li, Z., Qiu, X., Wu, J., Zhang, C.: DeFRCN: Decoupled Faster R-CNN for few-shot object detection. In: *Int. Conf. Comput. Vis.* (2021)
43. Ranjan, V., Sharma, U., Nguyen, T., Hoai, M.: Learning to count everything. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 3394–3403 (2021)
44. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. *Adv. Neural Inform. Process. Syst.* (2015)
45. Shao, J., Chen, S., Li, Y., Wang, K., Yin, Z., He, Y., Teng, J., Sun, Q., Gao, M., Liu, J., et al.: Intern: A new learning paradigm towards general vision. *arXiv preprint arXiv:2111.08687* (2021)
46. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. *Adv. Neural Inform. Process. Syst.* **30** (2017)
47. Sun, B., Li, B., Cai, S., Yuan, Y., Zhang, C.: FSCE: Few-shot object detection via contrastive proposal encoding. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2021)
48. Sun, M., Zhao, Z., Chai, W., Luo, H., Cao, S., Zhang, Y., Hwang, J.N., Wang, G.: Uniap: Towards universal animal perception in vision via few-shot learning. In: *AAAI Conf. Artif. Intell.* vol. 38, pp. 5008–5016 (2024)
49. Sun, Y., Huang, Y., Guo, H., Zhao, Y., Wu, R., Yu, Y., Ge, W., Zhang, W.: Misc210k: A large-scale dataset for multi-instance semantic correspondence. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 7121–7130 (2023)
50. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1199–1208 (2018)
51. Tang, S., Chen, C., Xie, Q., Chen, M., Wang, Y., Ci, Y., Bai, L., Zhu, F., Yang, H., Yi, L., et al.: Humanbench: Towards general human-centric perception with projector assisted pretraining. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 21970–21982 (2023)
52. Tian, Z., Chen, H., Shen, C.: Directpose: Direct end-to-end multi-person pose estimation. *arXiv preprint arXiv:1911.07451* (2019)
53. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Adv. Neural Inform. Process. Syst.* (2017)
54. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. *Adv. Neural Inform. Process. Syst.* **29** (2016)
55. Wang, J., Chen, D., Wu, Z., Luo, C., Zhou, L., Zhao, Y., Xie, Y., Liu, C., Jiang, Y.G., Yuan, L.: Omnivl: One foundation model for image-language and video-language tasks. *Adv. Neural Inform. Process. Syst.* **35**, 5696–5710 (2022)

56. Wang, P., Yang, A., Men, R., Lin, J., Bai, S., Li, Z., Ma, J., Zhou, C., Zhou, J., Yang, H.: Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In: *Int. Conf. Mach. Learn.* pp. 23318–23340 (2022)
57. Wang, X., Huang, T.E., Darrell, T., Gonzalez, J.E., Yu, F.: Frustratingly simple few-shot object detection. In: *Int. Conf. Mach. Learn.* (2020)
58. Wang, X., Wang, W., Cao, Y., Shen, C., Huang, T.: Images speak in images: A generalist painter for in-context visual learning. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 6830–6839 (2023)
59. Wang, Y.X., Ramanan, D., Hebert, M.: Meta-learning to detect rare objects. In: *Int. Conf. Comput. Vis.* (2019)
60. Wei, F., Sun, X., Li, H., Wang, J., Lin, S.: Point-set anchors for object detection, instance segmentation and pose estimation. In: *Eur. Conf. Comput. Vis.* pp. 527–544 (2020)
61. Wu, J., Liu, S., Huang, D., Wang, Y.: Multi-scale positive sample refinement for few-shot object detection. In: *Eur. Conf. Comput. Vis.* pp. 456–472 (2020)
62. Xiao, B., Wu, H., Wei, Y.: Simple baselines for human pose estimation and tracking. In: *Eur. Conf. Comput. Vis.* pp. 466–481 (2018)
63. Xiao, Y., Marlet, R.: Few-shot object detection and viewpoint estimation for objects in the wild. In: *Eur. Conf. Comput. Vis.* (2020)
64. Xie, E., Sun, P., Song, X., Wang, W., Liu, X., Liang, D., Shen, C., Luo, P.: Polar-mask: Single shot instance segmentation with polar representation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 12193–12202 (2020)
65. Xu, L., Guan, Y., Jin, S., Liu, W., Qian, C., Luo, P., Ouyang, W., Wang, X.: Vipnas: Efficient video pose estimation via neural architecture search. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 16072–16081 (2021)
66. Xu, L., Jin, S., Liu, W., Qian, C., Ouyang, W., Luo, P., Wang, X.: Zoomnas: searching for whole-body human pose estimation in the wild. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(4), 5296–5313 (2022)
67. Xu, L., Jin, S., Zeng, W., Liu, W., Qian, C., Ouyang, W., Luo, P., Wang, X.: Pose for everything: Towards category-agnostic pose estimation. In: *Eur. Conf. Comput. Vis.* (October 2022)
68. Yan, B., Jiang, Y., Sun, P., Wang, D., Yuan, Z., Luo, P., Lu, H.: Towards grand unification of object tracking. In: *Eur. Conf. Comput. Vis.* pp. 733–751 (2022)
69. Yan, X., Chen, Z., Xu, A., Wang, X., Liang, X., Lin, L.: Meta R-CNN: Towards general solver for instance-level low-shot learning. In: *Int. Conf. Comput. Vis.* (2019)
70. Yang, S.D., Su, H.T., Hsu, W.H., Chen, W.C.: Class-agnostic few-shot object counting. In: *IEEE Winter Conf. Appl. Comput. Vis.* pp. 870–878 (2021)
71. Yang, Z., Gan, Z., Wang, J., Hu, X., Ahmed, F., Liu, Z., Lu, Y., Wang, L.: Unitab: Unifying text and box outputs for grounded vision-language modeling. In: *Eur. Conf. Comput. Vis.* pp. 521–539 (2022)
72. Yao, R., Jin, S., Xu, L., Zeng, W., Liu, W., Qian, C., Luo, P., Wu, J.: Gkgnet: Group k-nearest neighbor based graph convolutional network for multi-label image recognition. In: *Eur. Conf. Comput. Vis.* (2024)
73. You, Z., Yang, K., Luo, W., Lu, X., Cui, L., Le, X.: Few-shot object counting with similarity-aware feature enhancement. In: *IEEE Winter Conf. Appl. Comput. Vis.* pp. 6315–6324 (2023)
74. Zamir, A.R., Sax, A., Shen, W., Guibas, L.J., Malik, J., Savarese, S.: Taskonomy: Disentangling task transfer learning. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 3712–3722 (2018)

75. Zeng, W., Jin, S., Liu, W., Qian, C., Luo, P., Ouyang, W., Wang, X.: Not all tokens are equal: Human-centric visual analysis via token clustering transformer. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 11101–11111 (2022)
76. Zhang, G., Luo, Z., Cui, K., Lu, S., Xing, E.P.: Meta-detr: Image-level few-shot detection with inter-class correlation exploitation. IEEE Trans. Pattern Anal. Mach. Intell. (2022)
77. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection. Int. Conf. Learn. Represent. (2023)
78. Zhang, Y., Zeng, W., Jin, S., Qian, C., Luo, P., Liu, W.: When pedestrian detection meets multi-modal learning: Generalist model and benchmark dataset. In: Eur. Conf. Comput. Vis. (2024)
79. Zhu, C., Chen, F., Ahmed, U., Shen, Z., Savvides, M.: Semantic relation reasoning for shot-stable few-shot object detection. In: IEEE Conf. Comput. Vis. Pattern Recog. (2021)
80. Zhu, X., Zhu, J., Li, H., Wu, X., Li, H., Wang, X., Dai, J.: Uni-perceiver: Pre-training unified architecture for generic perception for zero-shot and few-shot tasks. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 16804–16815 (2022)