

Panoptic-PartFormer++: A Unified and Decoupled View for Panoptic Part Segmentation

Xiangtai Li, Shilin Xu, Yibo Yang[†], Haobo Yuan, Guangliang Cheng, Yunhai Tong[†], Zhouchen Lin, Ming-Hsuan Yang, Dacheng Tao

Abstract—Panoptic Part Segmentation (PPS) unifies panoptic and part segmentation into one task. Previous works utilize separate approaches to handle things, stuff, and part predictions without shared computation and task association. We aim to unify these tasks at the architectural level, designing the first end-to-end unified framework, Panoptic-PartFormer. Moreover, we find the previous metric PartPQ biases to PQ. To handle both issues, we first design a meta-architecture that decouples part features and things/stuff features, respectively. We model things, stuff, and parts as object queries and directly learn to optimize all three forms of prediction as a unified mask prediction and classification problem. We term our model as Panoptic-PartFormer. Second, we propose a new metric Part-Whole Quality (PWQ), better to measure this task from pixel-region and part-whole perspectives. It also decouples the errors for part segmentation and panoptic segmentation. Third, inspired by Mask2Former, based on our meta-architecture, we propose Panoptic-PartFormer++ and design a new part-whole cross-attention scheme to boost part segmentation qualities further. We design a new part-whole interaction method using masked cross attention. Finally, extensive ablation studies and analysis demonstrate the effectiveness of both Panoptic-PartFormer and Panoptic-PartFormer++. Compared with previous Panoptic-PartFormer, our Panoptic-PartFormer++ achieves 2% PartPQ and 3% PWQ improvements on the Cityscapes PPS dataset and 5% PartPQ on the Pascal Context PPS dataset. On both datasets, Panoptic-PartFormer++ achieves new state-of-the-art results. Our models can serve as a strong baseline and aid future research in PPS. The source code and trained models will be available at <https://github.com/xtGH/Panoptic-PartFormer>.

Index Terms—Scene Understanding, Part-Whole Modeling, Panoptic Part Segmentation, Vision Transformer

1 INTRODUCTION

ONE fundamental problem in computer vision is to understand a scene at *multiple levels of concepts*. In particular, when people look at a figure, they can not only catch each visual entity, such as cars, buses, and background contexts like sky or road, but also understand the parts of entities, such as person-head and car-wheel, etc. The former is named as *scene parsing*, while the latter is termed as *part parsing*. One representative direction of unified scene parsing is Panoptic Segmentation (PS) [1]. It predicts a class label and an instance ID for each pixel and unifies the foreground objects (named as things) and background context (named as stuff). Part parsing has many cases, such as human or animal parts or car parts [2], [3]. While these tasks are often addressed in separate contents, both are equally important for numerous vision systems, including autonomous driving and robot navigation [4].

Recently, the Panoptic Part Segmentation (PPS, or Part-aware Panoptic Segmentation) [5] is proposed to unify these multiple




Fig. 1: Illustration of the Panoptic Part Segmentation task. It combines Panoptic Segmentation and Part Segmentation in a unified manner that provides a multi-level concept understanding of the image. Best viewed in color.

levels of abstraction into one single task. As shown in Fig. 1, it requires a model to output per-pixel scene-level classification for background stuff, segment things into individual instances, and segment each instance into specific parts. Several approaches [6], [7], [8] with hybrid formulations have been proposed to tackle this task. As shown in Fig. 2(a), these methods fuse individual predictions for different tasks to obtain the PPS results, *i.e.*, obtaining panoptic and part segmentation results individually based on separate models and then merging them for the PPS task. This makes the entire process complex, with substantial engineering efforts. In addition, the shared computation and task association are ignored, which leads to huge costs and inferior results. Another solution to this task is to make the part segmentation an extra head with a shared backbone [9], as shown in Fig. 2(b). Such design is well explored in the early PS studies [10], [11], [12], [13],

- X. Li, S. Xu, Y. Tong, and Z. Lin are with the State Key Laboratory of General Artificial Intelligence, School of Intelligence Science and Technology, Peking University, Beijing, China. Z. Lin is also with Puzhou Laboratory (Huangpu), Guangzhou, Guangdong, China. This work is supported by the National Key Research and Development Program of China (No.2023YFC3807600). Z. Lin was supported by the NSF China (No. 62276004). The first two authors contributed equally to this extension.
- Y. Yang is with KAUST, Jeddah, KSA.
- G. Cheng is with the Department of Computer Science at the University of Liverpool, UK.
- M.-H. Yang and H. Yuan are with the Department of Computer Science and Engineering at the University of California, Merced, US.
- D. Tao is with Nanyang Technological University, Singapore.
- Corresponding to: Y. Tong and Y. Yang

[14], [15], [16], [17], [18], [19]. However, most of them treat PS as separated tasks [10] or sequential tasks with post-processing components [18], which still cannot explore the mutual effect of the scene-level and part-level understandings.

Since Detection Transformer (DETR) [20], there have been several works [21], [22], [23], [24] unifying both thing and stuff learning via *object queries* in PS, which makes the entire pipeline elegant and achieves strong results. These results show that many complex components, including NMS and box detection heads, can be removed, and mask classification and prediction can be performed directly. In particular, such designs consider the complete scene understanding via performing interactions among object queries and image features. Joint training with things, stuff, and parts likely leads to better part segmentation results since the complete scene information renders more discriminative information, such as global context. Motivated by this, we propose a unified model to tackle the PPS task.

This work presents a simple yet effective framework named Panopic-PartFormer, a unified model for PPS tasks. As shown in Fig. 2(c), we introduce three different queries for modeling thing, stuff, and part predictions, respectively. Then, a decoupled decoder is proposed to generate fine-grained features that decode things, stuff, and part mask predictions. Rather than directly using the pixel-level self-attention in Transformer, we consider the recent work K-Net [23] that performs self-attention on the instance level for panoptic segmentation. Specifically, we focus on refining queries via their corresponding query features, which are defined as the masked decoder features using the mask associated with each query. We generate masks via dot product between queries and decoder features and perform updating object queries with the query features. This operation is implemented with one dynamic convolution [23], [25] and multi-head self-attention layers [26] between query and query features iteratively. The former poses instance-wise information from features to enhance query learning. The latter performs inner reasoning among different queries to model the relationship among thing, stuff, and part. Moreover, the entire procedure avoids the pixel-level computation in other vision Transformer decoders [20], [22]. Extensive experiments (Sec.4.3) show that our approach achieves much better results than the previous designs in Fig. 2 (a) and (b). It achieves the new state-of-the-art results on two challenging PPS benchmarks, including the Pascal Context PPS dataset (PPP) with about 6-7% PartPQ gain on ResNet101, 10% PartPQ gains using Swin Transformer [27], and the Cityscapes PPS dataset (CPP), with about 1-2% PartPQ gain.

Beyond Panopic-PartFormer framework, published in ECCV 2022 [29], we make more significant contributions in this work. **First**, by analyzing the PartPQ metric in PPS, we discover that this metric unfairly overemphasizes the effect of panoptic segmentation. To address this issue, we propose a new metric called Part-Whole Quality (PWQ), inspired by the principles of part-whole modeling and pixel-wise segmentation quality. The PWQ metric evaluates segmentation quality in two ways: first, by considering both pixel-wise and region-wise segmentation quality, and second, by using a hierarchical approach that equally weighs part and scene results. The PWQ metric decouples part segmentation errors from scene segmentation errors and distinguishes between pixel-wise segmentation (mIoU) and region-wise segmentation (PQ). By introducing the PWQ metric, we hope to advance the PPS field by providing a fair evaluation of part and scene segmentation quality at both the pixel-region level and the part-whole level, while also

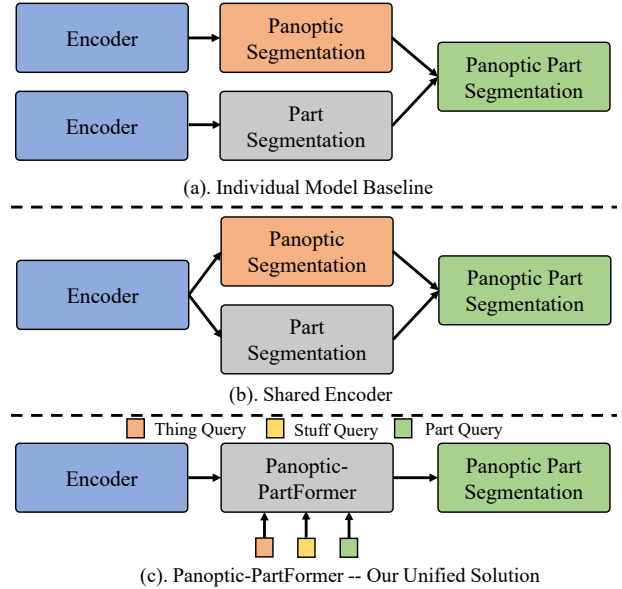


Fig. 2: (a) The baseline method proposed in [5] combines results of panoptic segmentation and part segmentation. (b) Panopic-FPN-like baseline [10], [11], [28] adds part segmentation into the current panoptic segmentation frameworks. (c) Our proposed approach represents things, stuff, and parts via object queries and performs joint learning in a unified manner.

balancing the ratio of PQ in PartPQ.

Second, we present a new enhanced model named Panopic-PartFormer++. In particular, we have *three* improvements over the original Panopic-PartFormer. **1.** Rather than using the coarse masked pooling for query-level reasoning, we adopt the recent stronger baseline Mask2Former design by replacing query learning with masked cross attention of both queries and multiscale features. Since most part objects in both PPP and CPP datasets are tiny, masked pooling ignores the fine details of the part. We present a global-first strategy. In particular, we first perform the global-part masked cross attention, where the part, thing, and stuff queries are jointly learned first. Then, we adopt part masked cross attention using the part query from the previous step to learn a better part feature. **2.** We present an enhanced version of the decoupled decoder. We append extra semantic segmentation loss to enhance part features learning. Moreover, we adopt a deformable encoder to extract multiscale features in a way similar to Mask2Former [24]. **3.** We jointly perform thing, stuff, and part query learning for each scale. We only supervise the part query learning for the highest resolution to keep fine-grained information.

Third, we present a more detailed analysis of our newly proposed architecture and metric. In particular, we verify the effectiveness of different architectural designs in the experiment part. Our model can directly output thing, stuff, and part segmentation predictions in a box-free and NMS-free manner. A sub-task of PPS, such as Panoptic segmentation, can also be evaluated. In the experiment part, we verify our panoptic segmentation predictions on the Cityscapes dataset [4], and achieve better results than the recent work [12]. Moreover, we explore the recent stronger backbones on both CPP and PPP datasets. Our final proposed Panopic-PartFormer++ achieves the new state-of-the-art results on three metrics, including PQ, PartPQ, and our proposed PWQ.

In particular, it achieves 63.1% PartPQ on the CPP dataset. Compared to the previous conference version, we observe about 1%-2% PartPQ and PWQ improvements. When scaling up our model, it also achieves 53.2% PartPQ on the PPP dataset, which is more than 15% improvement over the previous baselines and 5% improvement over the Panoptic-PartFormer in our conference version.

2 RELATED WORK

Part Segmentation. Current research on part segmentation can be divided into two categories: supervised learning and unsupervised learning. Supervised learning approaches for both instance and semantic part segmentation typically rely on human analysis [30]. In contrast, recent works on unsupervised learning [31], [32] focus on single objects with multiple parts and propose various contrastive learning approaches. However, most of these works do not consider multiple objects and background context. Therefore, in this paper, we mainly review supervised learning approaches. Several works [33], [34] have designed specific methods for semantic part segmentation, which only considers per-pixel classification. These methods focus on modeling global-part context more effectively. For human instance part segmentation, there are two paradigms: top-down pipelines [35], [36], [37], and bottom-up pipelines [38], [39], [40]. Top-down methods use two-stage detectors [6] to jointly detect and segment each instance. For each detected instance, they perform semantic part segmentation. In contrast, bottom-up methods first segment each human and then perform part grouping from each detected human. Several works also focus on task-specific part segmentation, such as car part segmentation [8], [41]. However, each of these settings requires a specific design. In comparison, this paper focuses on solving the PPS task, which naturally contains part segmentation as a sub-task.

Panoptic Segmentation. Earlier work [1] performs segmentation of things and stuff via separated networks [6], [42], and then directly combines the predictions of things and stuff for the panoptic segmentation result. To alleviate the computation cost, recent works [11], [16] are proposed to model both stuff segmentation and thing segmentation in one model with shared backbone and different task heads. Detection-based methods [10], [11] usually represent things with the box prediction, while several bottom-up models [18], [19] perform grouping instance via pixel-level affinity or center heat maps from semantic segmentation results. Several recent works [43], [44] focus on designing a stronger backbone and more effective task association method to merge semantic and instance segmentation. However, the former introduces a complex process, while the latter suffers from a performance drop in complex scenarios. Recently, several works [21], [45], [46], [47] propose to obtain panoptic segmentation masks without box supervision directly. However, these works do *not* cover the *knowledge of part-level semantics* of images, which can provide more comprehensive and hierarchical information for scene understanding.

Part-Whole Modeling. Part-Whole modeling has a long history in computer vision. Before the deep learning era, part-based methods [48], [49] are mainly used for object detection and recognition. Several approaches use part discovery in the deep learning era as essential cues for the fine-grained classification [50]. There are also several works [51], [52] using generative adversarial networks for few-shot part segmentation. Meanwhile, several works [53], [54] explore the part motion information in dynamic video inputs.

These works mainly focus on the parts of a single object for recognition, which ignores the background context or multiple complex objects. To better understand the entire scene and unify the instance-wise part segmentation [38], [55], the PPS task [5] is proposed. This work annotates two datasets (Cityscape PPS [4] and Pascal Context PPS [56]) and proposes a new metric named PartPQ [1], [5] for evaluation. This work also presents several baseline methods to obtain the final results. However, all these methods use separated networks to obtain the panoptic segmentation result. Besides, they use existing panoptic segmentation algorithms with part semantic segmentation as an isolated subnetwork. Recently, another work [57] formulated the PPS task as multi-level recognition by request. However, it still uses the two separated models to handle part and thing segmentation. As a comparison, we focus on the much more complex scene understanding task, considering multiple instances with their parts. We aim to design a *simple and unified* network for all the sub-tasks.

Vision Transformer (ViT). ViTs have recently been explored in numerous vision problems, including representation learning as a feature extractor, vision-language modeling, and using object queries for downstream tasks. For the first aspect, ViTs [27], [58], [59], [60] have more advantages in modeling global-range relations among the image patch features. Most recent works [61], [62] combine the local CNN design with ViTs. Moreover, recent work [63] also adopts ViTs for self-supervised learning via masked images, achieving stronger results than supervised learning. For the second aspect, several works [64] explore large-scale text-image pairs to build Vision Language Models (VLMs), which can be used for many zero-shot or open vocabulary settings [65]. Moreover, there are several works [66] transferring or adapting the knowledge of Transformer for downstream tasks. The third application uses the object query proposed by DETR [20] to unify or simplify the task pipeline. It models the object detection task as a set prediction problem with learnable queries. Other works [25], [67] explore the locality of the learning process or location prior to refine the object query. Query-based learning can also be applied to more fields, including instance segmentation [68], [69], video object detection [70], object tracking [71], and video segmentation [72]. Our method is inspired by these works, enjoying the benefit of ViTs backbone to unify and simplify the PPS task based on query learning. To our knowledge, we propose the *first unified Transformer model* for the PPS task.

Multi-Task Dense Prediction. This task [73], [74], [75] aims to predict pixel-level outputs, including semantic segmentation, depth estimation, surface normal, and object boundary of the scene, simultaneously. Earlier works [76], [77], [78], [79], [80], [81] explore cross-task learning by co-propagating intra-task and inter-task structures with different fusion strategies, such as dense fusion and cross-task affinity modeling. Recently, several concurrent works [82], [83], [84] have also explored transformer architecture to leverage the task association. In particular, MQtransformer [83] adopts task-specific queries to learn the correlation of different task features, while TaskPrompter [84] designs a spatial-channel task prompt to interact with patch tokens along spatial and channel dimensions. These methods output pixel-level format of the input image without considering instance discrimination and part-whole correlation between object and parts. Thus, via our extensive experiment in Sec. 4.3.2, their approaches cannot obtain performance gains on PPS. This is because these methods only explore feature-level interaction and ignore entity-level interaction. However, we argue that entity-level interaction is more important,

as proved in both PanopticPartFormer and PanopticPartFormer++ (Sec. 4.3.1 and Sec. 4.3.2).

3 METHOD

We first review the definition of PPS and the recent Mask Transformers for segmentation as preliminary in Sec. 3.1. Then, we describe the meta-architecture of PanopticPartFormer in Sec. 3.2 and detail the previous PanopticPartFormer design [29] in Sec. 3.3 as a unified model. Via our PanopticPartFormer, we discuss the potential issues of the PPS metric and propose a new balanced metric named Part-Whole Quality in Sec. 3.4. Next, in Sec. 3.5, we present a new enhanced version named PanopticPartFormer++ further to enhance our proposed PWQ metric and previous PartPQ metric. Finally, we describe the training and inference procedure of both models in Sec. 3.6.

3.1 Preliminaries

Problem Definition. Given an input image $I \in \mathbb{R}^{H \times W \times 3}$, where H and W are the height and width of the image, the goal of PPS is to output a group of non-overlapping masks $\{y_i\}_{i=1}^G = \{(m_i, c_i)\}_{i=1}^G$ where c_i denotes the ground truth class label of the mask m_i and G is the number of masks which depends on the input scene. In particular, $c_i \in \mathcal{L} = \{C_{st}, C_{th}, C_{pt}\}$ may be the label of *part*, *thing* or *stuff* and the part masks are within the thing masks. C_{st} , C_{th} , and C_{pt} are the categories of thing, stuff, and part. \mathcal{L} is the whole label set. The above is defined from the region level. At the same time, each pixel is assigned a unique triple tuple from the pixel level, including semantic ID, instance ID, and part ID. If we remove the part class of PPS, this task reduces to panoptic segmentation. Thus, panoptic segmentation can be deemed as a specific case of PPS. We term thing, stuff segments as *scene segments*. Previous works [5] merge individual outputs of panoptic segmentation and part segmentation, which ignores the natural part-whole relationship.

PartPQ Metric. PartPQ is proposed to unify both panoptic segmentation and part segmentation [5] and is defined as:

$$\text{PartPQ} = \frac{\sum_{(p,g) \in TP} \text{IoU}(p,g)}{|TP| + \frac{1}{2}|FP| + \frac{1}{2}|FN|}, \quad (1)$$

where TP is a true positive segment, FP is a false positive segment, and FN is a false negative segment. This metric is based on the Intersection over Union (IoU) between a predicted segment p and a ground-truth segment g for one class l (where $l \in \mathcal{L}$). A prediction is a TP if it has an overlap with a ground-truth segment and the $\text{IoU} > 0.5$. An FP is a predicted segment not matched with the ground truth, and an FN is a ground-truth segment not matched with a prediction. IoU contains two cases (part and non-part):

$$\text{IoU}(p,g) = \begin{cases} \text{mean IoU}_{\text{part}}(p,g), & l \in \mathcal{L}^{\text{parts}} \\ \text{IoU}_{\text{scene}}(p,g), & l \in \mathcal{L}^{\text{no-parts}} \end{cases}$$

If a segment contains the part annotations, PartPQ calculates the mean IoU of each part. If all the things have no part annotations, the PartPQ degrades to PQ.

General Mask Transformers. Recent methods for unified segmentation adopt formulations similar to DETR [20]. In Max-Deeplab [21], the object queries (including thing queries and stuff queries) are used to *perform cross-reasoning with high-resolution features and self-attention among the reasoned queries*. Different

methods usually have different designs for query reasoning. In particular, K-Net [23] proposes to use mask pooling with a dynamic convolution. Mask2Former [24] proposes masked cross attention with multiscale features. After this step, the updated queries are used for mask classification via an MLP. Like DETR [20], the corresponding masks are obtained by multiplying the updated queries with features. These steps are performed multiple times in a cascaded manner. In addition, the intermediate outputs are also supervised with mask labels during training.

3.2 Meta Architecture For PPS

Extending General Mask Transformers for PPS. One simple way to extend Mask Transformers for the PPS task is to add part queries, as shown in Fig. 4(a). Although it can serve as the simplest unified baseline, the limitation lies in the source of features for cross-reasoning. The part features are extremely local and fine-grained, while the thing/stuff features are mostly at global scopes. Thus, they naturally conflict with each other. Moreover, the part queries should be jointly modeled with the thing and stuff queries, so simply splitting the cross-reasoning on heads is not a good solution. We demonstrate the effectiveness of our meta-architecture in Sec. 4.3.3.

Panoptic-PartFormer Meta-Architecture. To handle these problems, we present Panoptic-PartFormer. First, we decouple both part features and thing/stuff features from the encoder. Then, the cross-attention is performed separately while the relationship is explored jointly. Finally, the decoder takes three different queries and their corresponding features as inputs and provides thing, stuff, and part mask predictions as outputs. Fig. 4(b) shows the meta-architecture of our Panoptic-PartFormer, which will be described in Sec. 3.3 and Sec. 3.5. Panoptic-PartFormer++ inherits the design of Panoptic-PartFormer as presented in Sec. 3.3 and Sec. 3.5.

Meta-Architecture Overview. Fig. 3 presents the overall illustration of Panoptic-PartFormer as an example. It contains three parts: (1) an encoder backbone to extract features; (2) a decoupled decoder to obtain the scene features and part features individually. Note that the *scene features* are used to generate thing and stuff masks, while the *part features* are used to generate part masks; (3) a Transformer decoder that takes three different types of queries and backbone features as inputs and provides thing, stuff, and part mask predictions. Both Panoptic-PartFormer and Panoptic-PartFormer++ share the **same** meta-architecture. We present the design details in Sec. 3.3 and Sec. 3.5.

Encoder. We first extract image features using an encoder. It contains a backbone network (Convolution Network [85] or ViTs [27]) with Feature Pyramid Network [86] as the neck. This results in a set of multiscale features, which play as the inputs of the decoupled decoder.

Decoupled Decoder. The decoupled decoder has two separate decoder networks to obtain features for scene features and part features, respectively. The former decodes both thing and stuff predictions, while the latter is applied to decode the part prediction. Our motivation is that part segmentation has different properties from panoptic segmentation. Firstly, part features need a more precise location and fine details. Secondly, the scene features focus on mask proposal level prediction while part features pay more attention to the inner parts of mask proposal, which conflict with each other. We show that the decoupled design leads to better results in the experimental section (see Sec. 4.3). For implementation, we adopt semantic FPN design [11] to fuse features in a

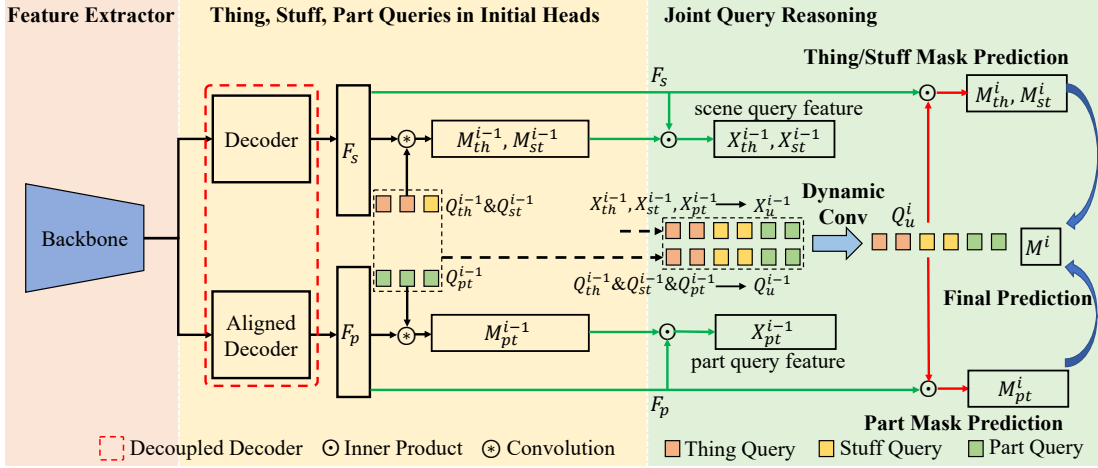


Fig. 3: Our proposed Panoptic-PartFormer. It contains three parts: (a) a backbone to extract features (Red area), (b) a decoupled decoder to generate scene features and part features along with the initial prediction heads to generate initial mask predictions. (Yellow area), (c) A cascaded transformer decoder will conduct the reasoning between the query and query features (Green area). **Green arrows** mean input (come from the previous stage) while **Red arrows** represent current stage output (used for the next stage). The outputs in the Red arrows are the inputs in the Green arrows. We take the last stage outputs as final output [20], [24].

top-down manner. Thus, we obtain the two features named F_s and F_p . In particular, we design a lightweight aligned feature upsampling for part segmentation. Rather than the naive bilinear upsampling, we propose to learn feature flow [40], [87] to warp the low-resolution feature to the high-resolution feature. Then, we sum all the predictions into the highest resolution as semantic FPN. Moreover, to preserve more locational information, we add the positional embedding to each stage of the semantic FPN, as adopted in the previous works [88], [89]. As such, a decoupled decoder outputs two separated features: scene features F_s , and part features F_p . The former is used to generate thing and stuff masks, while the latter is used to generate part masks, and both have the exact resolution. We show the effectiveness of such design for both Panoptic-PartFormer and Panoptic-PartFormer++.

3.3 Panoptic-PartFormer

Thing, Stuff, and Part as Queries with Initial Head Prediction:

Previous works [21], [22], [23] show that single mask classification and prediction can achieve the state-of-the-art results on the COCO dataset [90]. Motivated by this, our model treats thing, stuff, and part as the input queries to directly obtain the final panoptic part segmentation. Similar to prior works [23], [25], the initial weights of these queries are directly obtained from the first stage weights of the initial decoupled decoder prediction. For mask predictions of thing, stuff, and part, we use three 1×1 convolution layers to obtain the initial outputs of thing, stuff, and part masks. These layers are appended at the end of the decoupled decoder.

All these predictions are directly supervised with corresponding ground truth masks. As shown in [23], [25], using such initial heads can avoid heavier Transformer encoder layers for pixel-level computation, since the corresponding query features can be obtained via mask grouping from the initial mask prediction. Similarly, we obtain the three different queries for things, stuff, and parts, along with their initial mask predictions. We term them as Q_{th} , Q_{st} , Q_{pt} and M_{th} , M_{st} , M_{pt} with shapes $N_{th} \times d$, $N_{st} \times d$, $N_{pt} \times d$ and shapes $N_{th} \times H \times W$, $N_{st} \times H \times W$, $N_{pt} \times H \times W$. d , W , H are the channel number, width, height

of feature F_p and F_s . N_{th} , N_{st} , N_{pt} are the numbers of queries for things, stuff, and parts.

Joint Thing, Stuff, and Part Query Reasoning: The cascaded Transformer decoder takes previous mask predictions, previous object queries, and decoupled features as inputs and outputs the current refined mask predictions and object queries. The refined mask predictions, object queries, and decoupled features will be the inputs of the next stage. The relationship between queries and query features is jointly learned and reasoned. The main ideas of our approach are stated as follows. First, joint learning can learn the complete correlation between scene features and part features. For example, car parts must be on the road rather than in the sky. Second, joint reasoning can avoid several scene noisy cases, such as car parts on the human body or human parts on the car. We find that joint learning leads to better results (see Sec. 4.3).

We combine the three queries and the three mask predictions into a unified query Q_u^{i-1} and M_u^{i-1} where $Q_u^{i-1} = \text{concat}(Q_{th}, Q_{st}, Q_{pt}) \in \mathbf{R}^{(N_{th}+N_{st}+N_{pt}) \times d}$ and $M_u^{i-1} = \text{concat}(M_{th}^{i-1}, M_{st}^{i-1}, M_{pt}^{i-1}) \in \mathbf{R}^{(N_{th}+N_{st}+N_{pt}) \times H \times W}$. Here, i is the stage index of our Transformer decoder, and $i = 1$ means the predictions come from the initial heads. Otherwise, it means the predictions come from the outputs of the previous stage. In addition, ‘‘concat’’ is the concatenating operation performed along the first dimension. Similar to [23], we first obtain query features X^i via grouping from previous mask predictions M_u^{i-1} and input features (F_p , F_s) shown in Eq. 2 (dot product in Fig. 3). We present this process in one formulation for simplicity,

$$X^i = \sum_{m=1}^W \sum_{n=1}^H M_u^{i-1}(m, n) \cdot F(m, n), \quad (2)$$

where $X^i \in \mathbf{R}^{(N_{th}+N_{st}+N_{pt}) \times d}$ is the per-instance extracted feature with the same shape as Q_u , M_u^{i-1} is the per-instance mask extracted from the previous stage $i - 1$, and F is the input feature extracted for the decoupled decoder head. In addition, m , n are the indices of spatial location, and i is the layer number and starts from 1. As shown in the center part of Fig. 3, the part mask prediction and scene mask prediction are applied on corresponding features

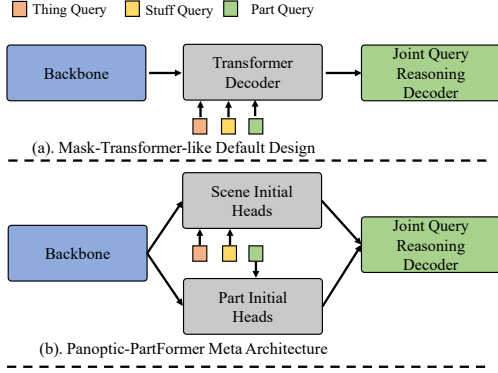


Fig. 4: Meta architecture of Panoptic-PartFormer. (a) is the default general Mask Transformer design for PPS. (b) is our Panoptic-PartFormer that decouples the Scene and Part heads in cases for both learning part features and cross-reasoning stages.

(F_p, F_s) individually where we obtain part query features X_{pt}^i and scene query features X_{th}^i and X_{st}^i . Then we combine these query features through $X_u^i = \text{concat}(X_{th}^i, X_{st}^i, X_{pt}^i)$. These inputs are shown in the green arrows in Fig. 3.

We use dynamic convolution [23], [25] to refine input queries Q_u^{i-1} with features X_u^i which are grouped from their masks,

$$\hat{Q}_u^{i-1} = \text{DynamicConv}(X_u^i, Q_u^{i-1}), \quad (3)$$

where the dynamic convolution uses the query features X_u^i to generate parameters to weight input queries Q_u^{i-1} . Specifically, DynamicConv uses input query features X_u^i to generate gating parameters via MLP and multiply back to the original query input Q_u^{i-1} . The reason is two-fold. First, compared with pixel-wise MHSA [20], [22], dynamic convolution introduces less computation and faster convergence. Second, it dynamically poses the instance-wise information to each query during training and inference, which shows better generalization and has complementary effects with MHSA. More details can be found in Sec. 4.3.

This operation absorbs more fine-grained information to help query and look for more precise locations. We adopt the same design [23] by learning gating functions to update the refined queries. The DynamicConv operation is based on:

$$\hat{Q}_u^{i-1} = \text{Gate}_x(X_u^i)X_u^i + \text{Gate}_q(X_u^i)Q_u^{i-1}, \quad (4)$$

where Gate is implemented with a fully connected (FC) layer followed by Layer Norm (LN) and a sigmoid layer, we adopt two different gate functions, including Gate_x and Gate_q . The former is to weigh the query features, while the latter is to weigh corresponding queries. Next, we adopt one self-attention layer with feed-forward layers [21], [26] to learn the correspondence among each query and update them accordingly. This operation leads to the full correlation among queries, shown as follows:

$$Q_u^i = \text{FFN}(\text{MHSA}(\hat{Q}_u^{i-1}) + \hat{Q}_u^{i-1}), \quad (5)$$

where MHSA means Multi Head Self Attention, FFN is the Feed Forward Network commonly used in the current vision of Transformers [20], [58]. The output refined query has the same shape as the input, i.e., $Q_u^i \in \mathbf{R}^{(N_{th}+N_{st}+N_{pt}) \times d}$.

Finally, the refined masks are obtained via dot product between the refined queries Q_u^i and the input features F_p, F_s . For mask classification, we adopt several feed-forward layers on Q_u^i and

TABLE 1: Upper Bound Analysis using Merge baseline and PanopticPartFormer on the CPP dataset. We report the results using ResNet-50 Backbone. Our new proposed PWQ is not sensitive to the quality of panoptic segmentation, which is more balanced.

Setting	Panoptic-GT	Part-GT	PartPQ	PartPQ _{part}	PWQ
Merged baseline	-	-	56.8	42.5	60.8
	-	✓	60.8	55.9	75.2
Panoptic-PartFormer	✓	-	85.4	54.4	78.3
	-	-	57.4	43.9	62.1
	-	✓	61.6	56.1	76.9
	✓	-	88.4	56.4	79.8

TABLE 2: PPS Metric Analysis. The proposed PWQ has all five important properties.

Metric Properties	mIOU	PQ	PartPQ	HPQ	PWQ
Pixel-Level Evaluation	✓	-	-	-	✓
Region-Level Evaluation	-	✓	✓	✓	✓
Part-Whole Evaluation	-	-	✓	✓	✓
Decouple Errors	-	-	-	-	✓
Balance Part and Scene Segments	-	-	-	-	✓

directly output the class scores (thing, stuff, and part). For mask prediction, we also adopt several feed-forward layers on Q_u^i and then we perform the inner product between the learned queries and features (F_s, F_p) to generate scene masks (thing and stuff) and part masks of stage i . These masks will be used for the following stage input, as shown in the red arrows in Fig. 3. The process of Eq. 2, Eq. 3, and Eq. 5 are repeated several times. We set the iteration number to 3 in this work. The inter-mask predictions are also optimized by mask supervision.

3.4 Part-Whole Quality Metric

Issues of PartPQ. There are several issues with the PartPQ metric, as defined in Eq. 1. First, PartPQ biases against PQ. We use two different models: Panoptic-PartFormer and Merged baseline [23] to analyze the upper bound performance by replacing the panoptic segmentation Ground Truth (GT) or part segmentation GT into our prediction. The Merged baseline contains a panoptic segmentation model [23] and one separate part segmentation model [42]. For both models, we find that replacing panoptic segmentation ground truth leads to a significant gain on PartPQ, while replacing part segmentation only results in a limited gain. This indicates that PartPQ is more *sensitive to panoptic segmentation than part segmentation* on the CPP dataset, which decreases the role of part segmentation. Second, it only considers region-level matching and lacks pixel-level evaluation. Third, it does not decouple the error of part and scene segmentation, which lacks interpretability of model performance. Namely, the PartPQ metric does not balance the two types of relations: part-whole relation and pixel-region relation.

Proposed Part-Whole Quality. Motivated by the previous analysis, we present a new metric named Part-Whole Quality. The main idea is to decouple the error of part and scene segmentation outputs and consider both region-level and pixel-level evaluation in one formulation. Our metric is defined as:

$$\text{PWQ} = \left(\frac{\text{mIoU}_{th,st} \cdot \text{PQ}_{th,st} + \text{mIoU}_{part} \cdot \text{PartPQ}_{part}}{2} \right)^{\frac{1}{2}}, \quad (6)$$

where PartPQ_{part} only considers the PartPQ metric for part segments, and mIoU refers to the mean Intersection over Union for scene segments (including thing (th) and stuff (st) masks) and part segments. In particular, we adopt an average of scene



PartPQ:62.0 PWQ: 63.9 PartPQ:62.3 PWQ: 65.7

Fig. 5: Visual examples for PartPQ and PWQ. Left: Merged baseline. Right: Our PanopticPartFormer++. Both methods use Swin-base as the backbone. Although both methods achieve similar PartPQ results. The part segmentation quality is significantly different. Our proposed PWQ indicates the better difference between both methods for part details.

and part to highlight the ratio of part segments. The PWQ metric can be decoupled into two items: Scene Segment Quality: $SSQ = mIoU_{th,st} \cdot PQ_{th,st}$ and Part Segment Quality: $PSQ = mIoU_{part} \cdot PartPQ_{part}$. The PWQ is defined as the geometric mean of SSQ and PSQ. As shown in the last column of Tab. 2, compared with previous metrics, PartPQ and HPQ, our new metric fully considers all five properties, including pixel-level evaluation, region-level evaluation, part-whole evaluation, decoupling the errors, and balancing the part-scene segments. We note that the recent work [57] presents a hierarchical panoptic quality (HPQ) metric to measure the accuracy of segmentation in different depths. However, it still does not well balance the ratio of part segments. After using the proposed PWQ, as shown in the last column of Tab. 1, by the same upper bound analysis, the proposed PWQ is more friendly to balancing the part segments and scene segments.

We present several cropped examples in Fig. 5. We use a stronger merged baseline [5] for illustration of the visual issues of PartPQ. The left-hand side of this figure is the same merge baseline as adopted in Tab. 1, and the right-hand side is the proposed Panoptic-PartFormer++, which will be detailed further. Both methods use the same Swin-base backbone. Although both methods achieve similar PartPQ results, it is shown from Fig. 5 that our proposed PWQ shows a clearer difference between the two methods, which is in line with the contrast of visual qualities.

3.5 PanopticFormer++

Motivation. As shown in the analysis of Sec. 3.4, to obtain a better PPS model, both scene-level and part-level segment outputs should be improved. In Panoptic-Partformer, we argue that three different issues still exist. **First**, the masked grouping operation (Eq. 2) misses the fine-grained details for both scene and part features. **Secondly**, although local part queries should be guided by high-level scene queries, the local details are still important, and they should not be dominated by scene features. **Thirdly**, a stronger feature extractor can lead to both scene-level and part-level segmentation, which is still not explored. Solving these issues can improve the PSS performance for various metrics.

Enhanced Feature Extractor. For the last issues, we integrate our Meta-Architecture in Fig. 4 with the Mask2Former design. We show that a better PPS model can improve both scene and

part masks by using Mask2Former [24] as the baseline model for feature extraction and fine-grained cross-reasoning. We use the deformable Feature Pyramid Networks [67] (deformable FPN) in Mask2Former to replace the feature extractor of Panoptic-PartFormer, and retain the decoupled decoder. The aligned decoder is appended after the output of deformable FPN. Similar to [24], we also adopt the positional embedding on both scene and part features.

Revisiting Joint Query Reasoning Design. The PanopticPartFormer uses a joint query reasoning design to simplify the PPS pipeline. However, as shown in Sec. 3.4, due to the unbalanced issues of PartPQ, previous design [29] overlooks the key role of part segmentation design. To enhance the part segmentation quality, we present a compelling scene and part query reasoning method by leveraging the benefits of joint global reasoning and local part features, *i.e.*, global-part masked cross attention.

Global-Part Masked Cross Attention. This design mainly solves the first two issues. This process contains two steps: First, we perform joint query reasoning by extending masked cross attention on the three object queries and scene features. (**Global Masked Cross attention**) Second, we use refined part queries to perform masked cross-attention on part features. (**Part Masked Cross attention**)

As shown in the experiment results of Panoptic-PartFormer in Sec. 4.3, joint learning with scene features and part features leads to better part segmentation. However, adding part supervision leads to a minor effect on scene segmentation. Thus, a better joint query reasoning design should be well-designed. We analyze the masked cross attention:

$$Q_l = \text{softmax}(M_{l-1} + \text{MLP}(Q_l)K_l^T)V_l + Q_{l-1}, \quad (7)$$

where M_{l-1} is a 2D attention mask from the previous stage l , $M_{l-1} \in \mathbb{R}^{n \times H \times W}$, and MLP is a Multilayer Perceptron network that contains two fully connected layers.

As Global Masked Cross (GMC) attention mainly models the global relation of scene and part queries, we combine all mask predictions into M_u^{i-1} and corresponding queries Q_u^{i-1} . We perform joint attention and reasoning as in Eq. 3 and Eq. 5. We first replace the Q_l with Q_u^{i-1} , M_{l-1} with M_u^{i-1} in Eq. 7. K_l and V_l are obtained via different MLPs on shared scene feature F_s . The operation works similarly to the previous dynamic convolution step in Eq. 3. However, previous pooling-based approaches [23], [29] miss the detailed information, which is more important for part queries. Using pixel-level cross attention injects more fine-grained information into all three queries, which leads to better results (Sec. 4.2). Then, we perform the self-attention and FFN on refined Q_u to learn the global correction of thing, stuff, and part query. This operation works the same as Eq. 5. For brevity, we ignore this operation for notation. Although Mask2Former contains multiscale features, for notation brevity, we use the feature with the highest mask resolution F_s for formulation purposes. The scene masks are classified with output queries. The thing/stuff mask predictions are obtained via inner production with F_s . For scene features, we adopt the multiscale features to produce the thing/stuff masks. Via pixel-level cross attention, we can achieve stronger results on both scene segmentation and part segmentation.

Then, for the second issue, we further present a refined step to inject more local features into part queries. In particular, as for Part Masked Cross (PMC) attention, we take part queries Q_{pt}^{i-1} from GMC attention outputs and corresponding part masks M_{pt}^{i-1} to replace Q_l and M_{l-1} in Eq. 7. In this module, K_l and V_l

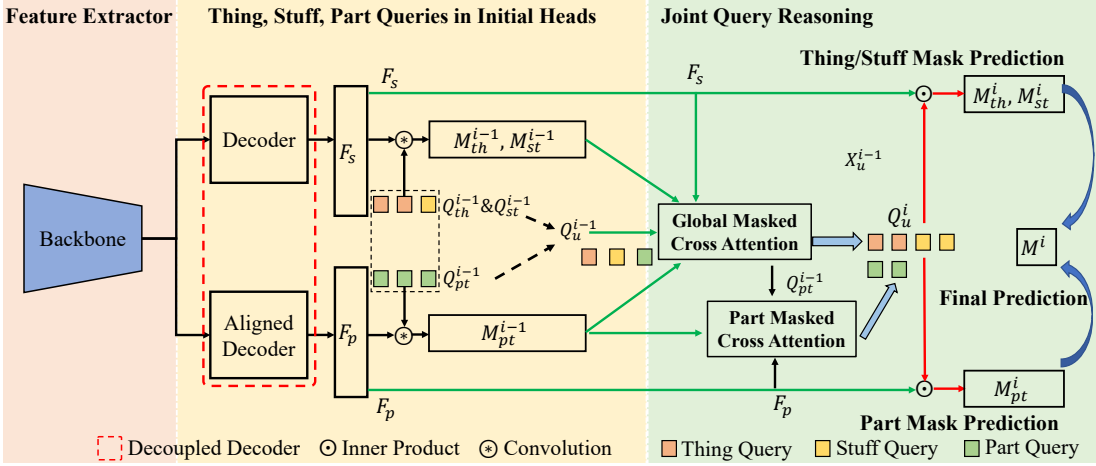


Fig. 6: Proposed Panoptic-PartFormer++. The Panoptic-PartFormer++ follows the meta-architecture design of Panoptic-PartFormer. It proposes a new Global-Part Masked Cross Attention design by performing joint query reasoning and then doing cross-attention on local part features. Global Masked Cross (GMC) attention takes object queries, object masks, and scene features F_s as inputs and outputs the refined object queries and scene masks. Part Masked Cross (PMC) attention takes the refined part queries from GMC, and part features F_p as inputs and outputs the refine part queries and part masks. Best viewed in color and zoom in.

are obtained via MLP on part feature F_p , and Q_{pt}^{i-1} has already absorbed the global context information and can be further refined by the local fine detailed feature F_p . The part masks are classified with final part queries. The part mask predictions are obtained via inner production with F_p . For both global and part masked cross attention, we perform MHSA operation in Eq. 5 after the cross attention. Unlike the Global Masked Cross Attention performed multiple times on different resolution features, the Part Masked Cross Attention is only performed on the *highest resolution* feature. In our implementation, we add an extra dense prediction head to force the part features sensitive to the part structure. We use the same Mask2Former design, *i.e.*, both Masked Cross Attentions are repeated multiple times. Moreover, after each cross attention in Eq. 7, we perform self-attention and feed-forward networks on output queries, which is similar to Eq. 5.

Discussion on Cross Attention Design. While our model uses dynamic convolution, masked cross-attention, and self-attention among queries, the main contribution of this work is the system-level unified model for this challenging task (PPS). We show that joint learning of the thing, stuff, and part learning benefits PPS tasks more than many other designs. In addition, we demonstrate that learning joint global relations first and then performing part segmentation is a good design for part segmentation. Furthermore, we verify these designs in the experiment section (Sec. 4.3.2).

3.6 Training and inference

Training. To train Panoptic-PartFormer and Panoptic-PartFormer++, we need to assign ground truth according to the pre-defined cost since all the outputs are encoded via queries. Similar to [22], we use bipartite matching as a cost by considering both mask and classification results. After the bipartite matching, we apply a loss jointly considering mask prediction and classification for each thing, stuff, and part. We apply cross-entropy loss on both classification and mask prediction. In addition, we adopt dice loss [91] on mask predictions (L_{part} , L_{thing} , L_{stuff}). Such settings are *the same*

as previous works [20], [22]. The loss for each stage i is:

$$\mathcal{L}_i = \lambda_{part} \cdot \mathcal{L}_{part} + \lambda_{thing} \cdot \mathcal{L}_{thing} + \lambda_{stuff} \cdot \mathcal{L}_{stuff} + \lambda_{cls} \cdot \mathcal{L}_{cls} \quad (8)$$

Note that the losses are applied to each stage, $\mathcal{L}_{final} = \sum_i^N \mathcal{L}_i$, where N is the total stages applied to the framework. Moreover, for Panoptic-PartFormer++, we add an extra semantic segmentation loss on part segmentation with part feature F_p .

Inference: We directly get the output masks from the corresponding queries according to their sorted scores. To obtain the final panoptic part segmentation, we first obtain the panoptic segmentation results and then merge part masks into panoptic segmentation results. For panoptic segmentation, we adopt the method in Panoptic-FPN [11] to merge the panoptic mask. We copy the predictions from panoptic segmentation for scene-level semantic classes that do not have part classes. For predicted instances with parts, we extract the part predictions for the pixels corresponding to this segment. Otherwise, if a part prediction contains a part class that does not correspond to the scene-level class, we set it to *void* label.

4 EXPERIMENTS

4.1 Experiment Setup

Datasets. We carry out experiments on the Cityscapes Panoptic Parts (CPP) and PASCAL Panoptic Parts (PPP) datasets. The CPP dataset is constructed based on the Cityscapes set [4] with part-level semantics and is annotated with 23 part-level semantic classes. Five scene-level semantic classes from the human and vehicle high-level categories are annotated with parts. The CPP dataset contains 2,975 training and 500 validation images. The PPP dataset extends the PASCAL VOC 2010 benchmark [56] with part-level and scene-level semantics. PPP has 4,998 training and 5,105 validation images. For fair comparisons, we use the same settings as [5], [56], for experiments with 59 scene-level classes (20 things, 39 stuff) and 57 part classes. We further report the Cityscapes Panoptic Segmentation validation set [4] results for performance evaluation.

Implementation Details. In the proposed method, ResNet [85] and Vision Transformer [27], [92] are adopted as the backbone networks, and other layers use Xavier initialization [93]. The optimizer is AdamW [94] with a weight decay of 0.0001. The training batch size is set to 16, and all models are trained with 8 GPUs. For the PPP dataset, we first pre-train our model on the COCO dataset [90] since most previous baselines [5] are also pre-trained on the COCO dataset. For the PPP dataset, we adopt the multiscale training [20] by resizing the input images from a scale of 0.5 to a scale of 2.0. We also apply random crop augmentations during training, where the training images are cropped with a probability of 0.5. Each random rectangular patch is then resized again to 800 (height), 1,333 (width) for training. For the CPP dataset, we use the same setting in Panoptic-Deeplab [18] where we resize the images with a scale range from 0.5 to 2.0 and randomly crop the whole image during training with a batch size of 8. We adopt V100 GPU servers with 16 32GB GPUs for training and testing. All the results are obtained via single-scale inference. We refer to the COCO-pretraining and Mapillary-pretraining for large models in the appendix.

Evaluation Metric. Aside from the PartPQ and PQ measures, we use our proposed PWQ metric to analyze the evaluated methods. For PanopticPart-Former, we adopt the PartPQ and PQ metrics for analysis. For PanopticPart-Former++, we use the PartPQ and PWQ measures (including SSQ and PSQ) for analysis.

4.2 Main Results

Results on Cityscape Panoptic Part Dataset. Tab. 3 shows the results of our Panoptic-PartFormer and other baselines. All the models use single-scale inference without test time augmentation. Our method with ResNe50 backbone achieves 57.4% PartPQ, which outperforms the previous work using complex pipelines [6], [42] with even stronger backbone [102]. For the same backbone, our method achieves **2.3% PartPQ gain** over the previous baseline. For the large model comparison, our method with Swin-Transformer achieves 61.9 % PartPQ. It outperforms the previous works that use state-of-the-art individual models [8], [98], [99], [100] by 0.5%. Note that the best model from HRNet [98] is pre-trained using the Mapillary dataset [103]. We use the same pipeline for fair comparisons. For our new proposed Panoptic-PartFormer++, we *only* use the COCO dataset for pertaining. As shown in Tab. 3, the proposed Panoptic-PartFormer++ achieves about **1.3%-1.6% PartPQ and 1.6% PWQ gain** over Panoptic-PartFormer. With the recently proposed backbone [104], our method obtains the new state-of-the-art results with 63.1% PartPQ and 66.5% PWQ. Compared to the recent method using separated vision transformers [57], our methods perform favorably with both ResNet50 and a larger backbone with fewer GFlops and a more direct pipeline. In conclusion, compared to Panoptic-PartFormer, Panoptic-PartFormer++ achieves better results on all three metrics, including PQ, PartPQ, and PWQ, which can be a new baseline for the PPS task.

GFlops and Parameter Comparison on the CPP dataset. The proposed unified model, Panoptic-PartFormer, is compact and efficient in terms of model parameters and run-time. Since the source code of [100] is not publicly available, we estimate the lower bound by its baseline model [6]. As shown in Tab. 5, our model obtains a decrease of 55%-60% in GFlops and 65%-70% in parameters. Compared with Panoptic-PartFormer, the proposed Panoptic-PartFormer++ only introduces an increase of 5% GFlops

and parameter cost. As shown in Tab. 5, our methods achieve significant improvements over previous models.

Results on the Pascal Panoptic Part Dataset. We evaluate our method with the works on the Pascal Panoptic Part dataset. Tab. 4 shows that, on different settings, our methods achieve state-of-the-art results on both PQ and PartPQ with significant gains. For the ResNet backbone, our methods achieve **6%-7% gains** on PartPQ. Moreover, our ResNet101 model performs favorably against the prior work using larger backbones [43]. For our new proposed Panoptic-PartFormer++, under the same backbone, we can achieve **1.9%-3.3% gains** on PQ, **1.8% gains** on PartPQ, and **3.0% gains** on PWQ. For all metrics, this indicates the newly proposed models are good at handling all level segmentation predictions.

Scaling Up Panoptic-PartFormer++. We analyze the performance of Panoptic-PartFormer++ for both CPP and PPP datasets using the recent state-of-the-art backbones. All the models are pre-trained on the COCO dataset and then fine-tuned on the CPP and PPP datasets. Different from previous works [5], [29], we do not use Mapillary data for pertaining. As shown in Tab. 6, we first explore various backbones on the CPP datasets, where we find that Convnext models [104] achieve the best results. We further explore the recent self-supervised pre-trained backbone BEiT-V2 [92]. However, we do not find extra gains in three metrics. The results show that a stronger pre-trained backbone is not critical for a better PPS model. Then, we analyze the performance of our model on the PPP dataset using various backbones. The model with Swin transformer performs better in terms of PartPQ and lower performance for PWQ. However, this is not the case with the Convnext models. These results show that the Convnext balances well for learning both part and scene segmentation tasks. Using the Convnext-base backbone, Panoptic-PartFormer++ model achieves state-of-the-art results.

Results on Cityscapes Panoptic Segmentation. We evaluate our method against the recent works on the Cityscapes Panoptic validation set. As shown in Tab. 8, our Panoptic-PartFormer also achieves favorable results against the previous works [12], [18], [23]. Panoptic-PartFormer++ achieves better results than the recent Mask2Former using the same backbone under the same settings. These results demonstrate the generalization ability of the proposed models on both panoptic segmentation and PPS tasks.

4.3 Ablation Study and Analysis

We present ablation studies and design analysis of the Panoptic-PartFormer and Panoptic-PartFormer++ models. We first present the core design and analysis of Panoptic-PartFormer (Sec. 4.3.1). Note that since both models share the same meta-architecture design, several designs are the same. Thus, we only explore the core design of Panoptic-PartFormer++ in Sec. 4.3.2. We then present the comparisons with Mask-Transformer-like baselines.

4.3.1 Ablation on Panoptic-PartFormer

effectiveness of each component. As shown in Tab. 7a, we show the effectiveness of each component of our framework by removing it from the original design. Removing the Decoupled Decoder (DD) results in a 1.4% drop in PartPQ. Removing Dynamic Convolution (DC) or Self Attention (SA) leads to a large drop in PQ, which means both are important for the interaction between queries and corresponding query features. Decreasing the stage number to 1 also leads to a significant drop. Performing more interaction results in a more accurate feature location for each query, which is the same as previous works [23], [25].

TABLE 3: **Experiment Results on the CPP dataset.** Prior works [5], [57] combine results from commonly used (top), and state-of-the-art methods (bottom) for semantic segmentation, instance segmentation, panoptic segmentation, and part segmentation. Metrics split into P and NP are evaluated on scene-level classes with and without parts, respectively. Our proposed Panoptic-PartFormer and Panoptic-PartFormer++ provide a simple yet unified solution with better performance and much fewer GFlops. We list the improvements of PanopticPartFormer++ by black \uparrow .

Panoptic seg. method	Part seg. method	PQ			PartPQ			PWQ
		All	P	NP	All	P	NP	
<i>Settings: Cityscapes Panoptic Parts validation set</i>								
UPSNNet [10](ResNet50)	DeepLabv3+ [42](ResNet50)	59.1	57.3	59.7	55.1	42.3	59.7	-
DeepLabv3+(ResNet50) & Mask R-CNN(ResNet50) [6]	DeepLabv3+ [42] (Xception- 65)	61.0	58.7	61.9	56.9	43.0	61.9	-
SegFormer-B1 & CondinInst [57]	SegFormer-B1 [95]	-	-	-	57.9	47.0	61.9	-
Panoptic-PartFormer (ResNet50) [29]	None	61.6	60.0	62.2	57.4	43.9	62.2	60.5
Panoptic-PartFormer++ (ResNet50)	None	63.6 (2.0 \uparrow)	61.0	63.1	59.2(1.8 \uparrow)	42.5	65.1	62.1 (1.6 \uparrow)
EfficientPS [96](EfficientNet) [97]	BSANet [8](ResNet101)	65.0	64.2	65.2	60.2	46.1	65.2	-
HRNet-OCR (HRNetv2-W48) [98], [99] & PolyTransform [100]	BSANet [8](ResNet101)	66.2	64.2	67.0	61.4	45.8	67.0	-
SegFormer-B5 & CondinInst [57]	SegFormer-B5 [95]	-	-	-	61.2	48.1	65.8	-
Part Based SegFormer-B5 & CondinInst [57]	SegFormer-B5 [95]	-	-	-	62.4	50.5	66.6	-
Panoptic-PartFormer (Swin-base) [29]	None	66.6	65.1	67.2	61.9	45.6	68.0	64.4
Panoptic-PartFormer++ (Swin-base)	None	68.0 (1.4 \uparrow)	67.0	68.3	62.3 (0.4 \uparrow)	46.0	68.2	65.7 (1.7 \uparrow)
Panoptic-PartFormer++ (Convnext-base)	None	68.2	67.2	69.0	63.1	46.4	69.1	66.5

TABLE 4: **Experiment Results on the PPP dataset.** Our methods using a *smaller* backbone ResNet101 achieve new state-of-the-art results with much lower GFlops. Both models show the effectiveness of the end-to-end design. We list the improvements of PanopticPartFormer++ by \uparrow .

Panoptic seg. method	Part seg. method	PQ	PartPQ	PWQ
<i>Settings: Pascal Panoptic Parts validation set</i>				
DeepLabv3+ & Mask R-CNN [6](ResNet50)	DeepLabv3+ [42](ResNet50)	35.0	31.4	-
DLv3-ResNet50 [101] & DetectoRS [43]	BSANet [8]	42.0	38.3	-
Our Unified Approach				
Panoptic-PartFormer (ResNet50)	None	47.6	43.2	42.3
Panoptic-PartFormer (ResNet101)	None	49.2	44.0	43.6
Panoptic-PartFormer++ (ResNet50)	None	51.6 (1.9 \uparrow)	45.1 (1.8 \uparrow)	45.2 (2.0 \uparrow)
Panoptic-PartFormer++ (ResNet101)	None	52.5 (3.3 \uparrow)	45.6 (1.9 \uparrow)	46.0 (2.4 \uparrow)

TABLE 5: **GFlops and parameter comparison on the Cityscapes PPS dataset.** The GFlops are measured with 1200×800 input. Both proposed Panoptic-PartFormer and Panoptic-PartFormer++ achieve better results with much fewer GFlops and parameters.

Method	PQ	PartPQ	Param(M)	GFlops
UPSNNet + DeepLabv3+ (ResNet50)	59.1	55.1	>87	>890
Panoptic-PartFormer (ResNet50)	61.6	57.4	37.4	185.8
Panoptic-PartFormer++ (ResNet50)	62.8	59.2	45.6	215.4
HRNet(OCR) +PolyTransform + BSANet	66.2	61.4	>181	>1154
Panoptic-PartFormer (Swin-base)	66.6	61.9	100.3	408.5
Panoptic-PartFormer++ (Convnext-base)	68.2	63.1	120.2	519.5

TABLE 6: **Scaling Up the Panoptic-PartFormer++.** We adopt the recent state-of-the-art backbones to explore the effectiveness of representation learning on PPS tasks. All the methods are pre-trained on COCO.

Dataset	Backbone	PQ	PartPQ	PWQ
CPP	Swin-base	68.0	62.3	65.7
CPP	Convnext-base	68.2	63.1	66.5
CPP	Convnext-large	67.8	62.8	65.4
CPP	BEIT-V2	67.5	62.7	66.0
PPP	Swin-base	59.8	51.3	52.7
PPP	Convnext-base	58.7	52.6	54.2
PPP	Convnext-large	60.2	52.8	54.7

Part query depends more on the thing query? With our framework, we can easily analyze the relationship between the stuff query, thing query, and part query. We present several ways of reasoning and fusing different queries. From intuitive thought, part information is more related to thing query. We design two different query interaction methods, shown in Tab. 7b. For separate

reasoning, we use DD and SA on two query pairs: the stuff-thing query and the thing-part query. For sequential reasoning, we perform DD and SA with the thing-part query first and the stuff-thing query second. However, we find the best model is the joint reasoning, which is the default setting described in Sec. 3.3. These results show that better part segmentation needs the whole scene context rather than thing features only.

Joint query modeling or separate modeling on part queries?

Using the same settings as PanopticFPN [11], we use the semantic-FPN-like model for part segmentation. Dense Prediction (DP) is the baseline method shown in Fig. 1(b). We adopt the same merging process for panoptic segmentation and part segmentation. As shown in Tab. 7c, our joint query-based method achieves better results and outperforms the previous dense prediction-based approach and an improved model [105]. These results indicate that joint learning greatly benefits part segmentation and demonstrate the effectiveness of the proposed framework.

Aligned decoder is more important for part segmentation. As shown in Tab. 7d, using the aligned part decoder results in better PartPQ, especially for the things with parts. Adding both paths with the aligned decoder does not bring extra gain. These results show that part segmentation needs more detailed information, while thing and stuff predictions do not rely on it.

Necessity of Positional Encoding on X_p and X_u . The results in Tab. 7e show that removing positional encoding leads to inferior results on both PQ and PartPQ, indicating the importance of position information [20], [23], [89].

Will boundary supervision help for part segmentation? In Tab. 7g, we also add boundary supervision for part segmentation, where we use the dice loss [109] and binary cross entropy loss. Although boundary supervision may seem to be useful for part segmentation [110], [111], there is no performance gain observed. As our mask is generated from the aligned decoder, it already contains detailed information. These results motivate us to focus on the relationship between the scene, and part features rather than modeling part segmentation solely.

Will joint training help for panoptic segmentation? As shown in Tab. 7f, joint learning benefits the panoptic segmentation baseline. However, the performance gains are limited since both thing, and stuff predictions do not need much detailed information.

4.3.2 Ablation on Panoptic-PartFormer++

Effectiveness of Each Component. In Tab. 9a, we demonstrate the effectiveness of each component by removing it from the

TABLE 7: **Ablation studies and analysis on Panoptic-PartFormer using the CPP validation set.** All the models use ResNet50 as the backbone, with 64 training epochs. DD: Decoupled Decoder. DC: Dynamic Convolution. SA: Self Attention. I: Interaction number. DP: Dense Prediction. w: with. ASPP: Atrous Spatial Pyramid Pooling [105].

(a) Effect of each component.							(b) Ablation on Query Reasoning Design			(c) Dense Prediction or Query Prediction on Part.		
DD	DC	SA	I=1	I=3	PQ	PartPQ	Setting	PQ	PartPQ	Method	PQ	PartPQ
✓	✓	✓	-	✓	61.6	57.4	Joint Reasoning	61.6	57.4	Joint Query	61.6	57.4
-	✓	✓	-	✓	61.2	55.9	Separate Reasoning	61.1	56.8	DP-Based	59.8	55.9
✓	-	✓	-	✓	57.0	52.2	Sequential Reasoning	60.8	56.3	DP-Based w ASPP [105]	59.9	56.1
✓	✓	-	-	✓	57.3	53.4						
✓	✓	✓	✓	-	58.3	54.2						

(d) Effect of Aligned Decoder Design.					(e) Effect of Position Encoding (PE) on F_s and F_p .			(f) Effect on adding part annotations (anno).		(g) Effect of adding boundary supervision, boundary loss (b-loss)		
Settings	PQ	PartPQ	P	NP	Method	PQ	PartPQ	Method	PQ	Method	PQ	PartPQ
w Aligned	61.6	57.4	43.9	62.2	w PE	61.6	57.4	w/o part anno	61.2	baseline	61.6	57.4
w/o Aligned	61.4	56.3	41.2	62.1	w/o PE	59.0	55.1	w part anno	61.6	+ b-loss	61.5	57.2
on Both Features	61.4	57.2	43.7	62.0								

TABLE 8: **Experiment results on the Cityscapes Panoptic validation set.** * indicates using DCN [106]. All the methods use single-scale inference. Both K-Net and Mask2Former are pre-trained on COCO datasets. Thus, the performance is better than the original papers.

Method	Backbone	PQ	PQ_{th}	PQ_{st}
UPNet [10]	ResNet50	59.3	54.6	62.7
SOGNet [16]	ResNet50	60.0	56.7	62.5
Seamless [15]	ResNet50	60.2	55.6	63.6
Unifying [107]	ResNet50	61.4	54.7	66.3
Panoptic-DeepLab [18]	ResNet50	59.7	-	-
Panoptic FCN* [12]	ResNet50	61.4	54.8	66.6
Panoptic FCN++ [108]	Swin-large	64.1	55.6	70.2
K-Net [23]	ResNet50	61.2	52.4	66.8
Mask2Former [24]	ResNet50	63.0	54.3	67.2
Mask2Former [24]	Convnext-base	67.5	61.1	71.8
Panoptic-PartFormer	ResNet50	61.6	54.9	66.8
Panoptic-PartFormer++	ResNet50	63.6	57.5	68.1
Panoptic-PartFormer	Swin-base	66.6	61.7	70.3
Panoptic-PartFormer++	Swin-base	68.0	61.9	72.0
Panoptic-PartFormer++	Convnext-base	68.2	62.3	72.5

proposed model. Removing the decoupled decoder leads to a 1.2% PartPQ drop and a 1.8% PSQ drop from the Panoptic-PartFormer++ model. These results indicate the decoupled design is a key component for both Panoptic-PartFormer and Panoptic-PartFormer++. Removing the part masked cross attention module leads to a 1.5% PSQ drop. This shows the effectiveness of the extra part query attention with part features. In both cases, the scene segmentation results (SSQ) are the same. The results indicate that part attention and decoupled decoder design do not affect panoptic segmentation. Moreover, these results also indicate the interpretability of our proposed PWQ metric, where we achieve improvements mainly from the better part segmentation quality.

Ablation on Global-Part Masked Cross Attention Design. In Tab. 9b, we analyze the query interaction design in Panoptic-PartFormer++. We use three different attention methods, including joint reasoning, part-first reasoning, and global-first reasoning. For joint reasoning, we perform masked cross-attention on three queries with the corresponding scene features. For part-first attention module, we first perform part cross attention for part query and part feature, and then we perform the joint reasoning on three




Fig. 7: Ablation Study on Global-Part Masked Cross Attention Design. (a), Our Proposed Method. (b), Local Part Query Cross Attention First. (c), Scene Feature for Part Query Cross Attention. (d), Solo Local Part Query Cross Attention.

queries. In the global-first attention module, we first perform joint reasoning, and then we perform the part cross attention. Different from the results in Tab. 7b, after adopting the masked cross attention, we find adopting global-first leads to the best results. This is because masked attention contains more fine-grained information, and it can be improved with the guidance of joint global scene query learning. In the appendix, we present more detailed ablation studies in our global part-masked cross-attention design. Specifically, we present three more different designs on the location of the refined part query, where our proposed global reasoning and local refined design lead to better results.

Ablation on Extra Part Dense Prediction. Our Panoptic-PartFormer++ also adopts an extra semantic part segmentation head during training for part features. Removing such supervision leads to about 0.3% PartPQ drop and 0.8% PWQ drop. This result means the decoupled design can be further improved by adding task-specific loss.

Ablation on Final PPS merging. Since our Panoptic-PartFormer++ has an extra dense semantic segmentation prediction head for part feature learning, we also explore the results of its part segmentation quality. As shown in Tab. 9c, replacing the dense prediction results leads to a significant drop. This indicates the effectiveness of our unified modeling for PPS.

Ablation on Global-Part Masked Cross Attention Design. We show detailed ablation designs of our Panoptic-PartFormer++ head in Fig. 7. Fig. 7(a) is our default design. Fig. 7(b) is Part Query

TABLE 9: **Ablation studies and analysis on Panoptic-PartFormer++ using the CPP validation set.** GMC: Global Masked Cross attention. PMC: Part Masked Cross attention. DD: Decoupled Decoder. DPH: Dense Prediction Head from the auxiliary head. ResNet50 is adopted as the backbone.

(a) Effect of Each Component							(b) Query Reasoning Design			(c) Choice Of PPS Fusion.		
DD	GMC	PMC	PartPQ	PWQ	SSQ	PSQ	Setting	PartPQ	PWQ	Merge Method	PartPQ	PWQ
✓	✓	✓	59.2	62.1	49.6	27.5	Joint Reasoning	58.2	59.9	with Part Query	59.2	62.1
-	✓	✓	58.0	60.2	49.5	25.3	Part First	58.3	60.5	with DPH	57.4	59.2
✓	✓	-	58.5	60.5	49.3	26.0	Global First	59.2	62.1	Merged DPH/Part Query	58.4	61.5

TABLE 10: **Comparison of Global-Part Masked Cross Attention Design in Fig. 7** All models use ResNet50 as backbone and meta-architecture is PanopticPartFormer++.

Method	PQ	PartPQ	PWQ
Global First (Fig. 7 (a))	63.6	59.2	62.1
Part First (Fig. 7 (b))	63.3	58.3	60.5
Using Scene Features (Fig. 7 (c))	63.6	58.0	61.5
Solo Local Part Query (Fig. 7 (d))	62.8	58.3	60.0

TABLE 11: **Exploration Multi-Task Dense Prediction Approaches for PPS.** All models use ResNet50 as backbone and meta-architecture is PanopticPartFormer++.

Method	PQ	PartPQ	PWQ	GFlops
Our Default Baseline	63.6	59.2	62.1	210.5
+ MQFormer [83]	62.3	56.4	60.2	223.4
+ TaskPrompter [84]	63.0	58.2	61.2	215.6

TABLE 12: **Compared with Single Task Baseline.** All models use ResNet50 as backbone. mIoU evaluates part segmentation.

Method	PQ	PartPQ	PWQ	mIoU
Our Default Baseline	63.6	59.2	62.1	52.0
Single Part Segmentation	-	-	-	50.2
Single Panoptic Segmentation	63.2	-	-	-

TABLE 13: **Comparison of Meta-Architecture Design in Fig. 4 (a) and (b) on the CPP validation set.** Both models achieve better results on all three metrics.

Method	PQ	PartPQ	PWQ	GFlops
K-Net + Part Query	60.8	56.0	58.1	183.2
Panoptic-PartFormer	61.6	57.5	60.1	185.8
Mask2Former + Part Query	62.7	58.2	60.2	210.5
Panoptic-PartFormer++	63.4	59.2	62.1	215.4

TABLE 14: **Effect Of COCO pretraining.** We use ResNet50 as the backbone on the CPP validation set. Different from Tab. 7 and Tab. 9, all the methods use the same training setting.

Method	COCO pre-train	PQ	PartPQ	PWQ
Panoptic-PartFormer	-	54.5	57.8	54.6
Panoptic-PartFormer	✓	57.4	61.6	60.5
Panoptic-PartFormer++	-	61.4	57.5	58.8
Panoptic-PartFormer++	✓	63.6	59.2	62.1

Cross attention First design. Fig. 7(c) uses scene features F_s for both global and local cross-attention. Fig. 7(d) is solo Part Query Cross Attention without Global Refinement. We report the main results in Tab. 10. The proposed method achieves the best result

on all three metrics. Using part query first leads to inferior results on Part Segmentation. This is because, after global cross-attention with the scene feature, the local part information is lost. Using scene features also leads to bad results, which demonstrates the effectiveness of our decoupled decoder head design. Part segmentation needs more fine-grained feature representation. Moreover, using solo local part queries without global reasoning has worse results. This indicates the part segmentation still needs the global prior, which shares a similar finding in Tab. 7 (c).

Explore Methods in Multi-Task Dense Prediction For PPS. In Tab. 11, we also explore two representative approaches [83], [84] for the cross-task feature learning into our PanopticPartFormer++ framework. In particular, we use their open-source implementation¹. We merge their operation into the decoupled feature decoder but keep the remaining unchanged. From the results of Tab. 11, though increasing the GFlops, we do not find any performance gains over our baselines. This is because feature-level interaction is not suitable for PPS tasks, while the entity-level interaction between the scene query feature and part query feature is more important.

Comparison with Single Head Baseline. We also verify the effectiveness of multi-task co-training in our framework in Tab. 12. In particular, we remove the scene head or part head and the remaining components are the same. From that table, we find significant improvement in part segmentation and nearly no improvements in panoptic segmentation, which shares the same conclusion in Tab. 7 (f) of PanopticPartFormer. Joint co-training leads to better part segmentation results.

4.3.3 Comparison with Baseline Methods

Compared with MaskFormer Meta-Architecture Tab. 13 shows the evaluation results of the proposed and MaskFormer models (Fig. 4(a)). All the methods use ResNet-50 as the backbone. Although both works [23], [24] achieve better results than the previous separate baselines [5]. Our methods achieve better results with fewer extra GFlops costs. These results demonstrate the effectiveness of our Meta-Architecture (Fig. 4(b)). In addition, our methods provide a unified and decoupled view for the PPS task, which is significantly different from the MaskFormer study.

Effect of COCO-pretraining. In Tab. 14, we demonstrate the effect of COCO pertaining on the CPP dataset. The performance of both Panoptic-PartFormer and Panoptic-PartFormer++ drops without COCO-pretraining. However, the new proposed Panoptic-PartFormer++ is less sensitive to COCO-pretraining. These results indicate that Panoptic-PartFormer++ is also a data-efficient model. For the effect of COCO-pretraining on the PPP dataset, we refer the reader to the appendix.

1. <https://github.com/prismformore/Multi-Task-Transformer>, <https://github.com/yangyangxu0/MQTransformer>

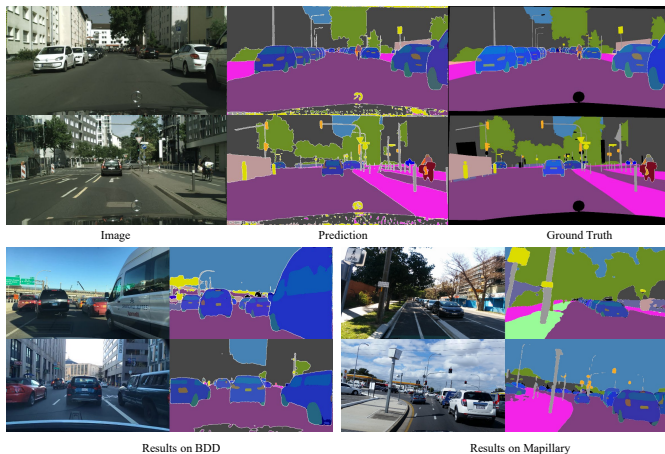


Fig. 8: Visualization of our Panoptic-PartFormer. Top: results on Cityscapes PPS validation set. Bottom left: prediction on BDD dataset [112]. Bottom right: prediction on Mapillary dataset [103]. Best view it on screen.



Fig. 9: Visualization results on the Pascal Context Panoptic Part validation set using Panoptic-PartFormer. Note that stuff classes have the same color, while thing and part classes are not. Best view it on screen.

4.4 Visualization Analysis

Visualization and Generalization With Panoptic-PartFormer.

We present several visualization examples using our model on Cityscapes PPS validation set. Moreover, we also visualize several examples on the Mapillary dataset [103] and BDD dataset [112] to show the generalization ability of our method. As shown in the first row of Fig. 8, our method achieves considerable results. Moreover, on the Mapillary [103] and BDD datasets [112], which do not have part annotations, our method can still work well as shown in the last row of Fig. 8. In addition, we also visualize the results on the PPP datasets in Fig. 9. The first two rows show the crowded human scene and outdoor scene. Both cases show that our model can obtain convincing results. The last row shows the small object cases. The failure cases are due to tiny objects, including their inner parts.

Visual Improvements by Panoptic-PartFormer++. We show some sample results of PanopticPartFormer++ in Fig. 10. Overall, Panoptic-PartFormer++ has better part segmentation quality (better boundaries, aligned semantic consistency within the part) on both the CPP and PPP datasets, as shown in the red boxes of Fig. 10. We present more visualization examples in the appendix.

Failure Cases Analysis. In Fig. 11, we show some failure cases of the Panoptic-PartFormer++ models. As shown in the red boxes, there are two issues: wrong background prediction during PPS




Fig. 10: Visualization improvements on the CPP and PPP datasets. Left: Panoptic-PartFormer. Right: Panoptic-PartFormer++. Both models use the ResNet-50 backbone. Shown in the red boxes, Panoptic-PartFormer++ achieves better part segmentation results on both CPP and PPP datasets.




Fig. 11: Failure cases on the CPP (top) and PPP (bottom) datasets.

merging and inconsistency prediction within the part. These results suggest that a better fusion procedure or a new part segmentation branch may be the future research direction of PPS. We present more failure cases in the appendix.

5 CONCLUSION

This paper addresses the panoptic part segmentation task by designing a unified and decoupled model named Panoptic-PartFormer. We present a decoupled meta-architecture and perform joint reasoning along with thing, stuff, and part queries and their corresponding features. Using this model, we find the unbalanced issues of the PartPQ metric. To handle this problem, we present a balanced and error-decoupled metric from two perspectives: pixel-region and part-scene. Then, we propose a new enhanced model based on our meta-architecture and Mask2Former, named Panoptic-PartFormer++. We design global-part masked cross attention to perform interactions between scene features and part features effectively. We further conduct extensive experiments on two widely used benchmark datasets. Furthermore, our models achieve state-of-the-art results on both datasets with much fewer GFlops and parameters than existing approaches.

REFERENCES

- [1] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár, “Panoptic segmentation,” in *CVPR*, 2019. 1, 3
- [2] X. Liang, C. Xu, X. Shen, J. Yang, S. Liu, J. Tang, L. Lin, and S. Yan, “Human parsing with contextualized convolutional neural network,” in *ICCV*, 2015. 1
- [3] Q. Geng, H. Zhang, F. Lu, X. Huang, S. Wang, Z. Zhou, and R. Yang, “Part-level car parsing and reconstruction in single street view images,” *TPAMI*, 2021. 1
- [4] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The cityscapes dataset for semantic urban scene understanding,” in *CVPR*, 2016. 1, 2, 3, 8
- [5] D. de Geus, P. Meletis, C. Lu, X. Wen, and G. Dubbelman, “Part-aware panoptic segmentation,” in *CVPR*, 2021. 1, 2, 3, 4, 7, 8, 9, 10, 12
- [6] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask r-cnn,” in *ICCV*, 2017. 1, 3, 9, 10
- [7] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network,” in *CVPR*, 2017. 1
- [8] Y. Zhao, J. Li, Y. Zhang, and Y. Tian, “Multi-Class Part Parsing With Joint Boundary-Semantic Awareness,” in *ICCV*, 2019. 1, 3, 9, 10
- [9] S. K. Jagadeesh, R. Schuster, and D. Stricker, “Multi-task fusion for efficient panoptic-part segmentation,” in *ICPRAM*, 2023. 1
- [10] Y. Xiong, R. Liao, H. Zhao, R. Hu, M. Bai, E. Yumer, and R. Urtasun, “Upsnet: A unified panoptic segmentation network,” in *CVPR*, 2019. 1, 2, 3, 10, 11
- [11] A. Kirillov, R. Girshick, K. He, and P. Dollár, “Panoptic feature pyramid networks,” in *CVPR*, 2019. 1, 2, 3, 4, 8, 10
- [12] Y. Li, H. Zhao, X. Qi, L. Wang, Z. Li, J. Sun, and J. Jia, “Fully convolutional networks for panoptic segmentation,” in *CVPR*, 2021. 1, 2, 9, 11
- [13] Y. Chen, G. Lin, S. Li, O. Bourahla, Y. Wu, F. Wang, J. Feng, M. Xu, and X. Li, “Banet: Bidirectional aggregation network with occlusion handling for panoptic segmentation,” in *CVPR*, 2020. 1
- [14] J. Li, A. Raventos, A. Bhargava, T. Tagawa, and A. Gaidon, “Learning to fuse things and stuff,” *arXiv:1812.01192*, 2018. 1
- [15] L. Porzi, S. R. Bulo, A. Colovic, and P. Kotschieder, “Seamless scene segmentation,” in *CVPR*, 2019. 1, 11
- [16] Y. Yang, H. Li, X. Li, Q. Zhao, J. Wu, and Z. Lin, “Sognet: Scene overlap graph network for panoptic segmentation,” in *AAAI*, 2020. 1, 3, 11
- [17] Y. Wu, G. Zhang, H. Xu, X. Liang, and L. Lin, “Auto-panoptic: Cooperative multi-component architecture search for panoptic segmentation,” in *NeurIPS*, 2020. 1
- [18] B. Cheng, M. D. Collins, Y. Zhu, T. Liu, T. S. Huang, H. Adam, and L.-C. Chen, “Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation,” in *CVPR*, 2020. 1, 2, 3, 9, 11
- [19] H. Wang, Y. Zhu, B. Green, H. Adam, A. Yuille, and L.-C. Chen, “Axial-deeplab: Stand-alone axial-attention for panoptic segmentation,” in *ECCV*, 2020. 1, 3
- [20] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, “End-to-end object detection with transformers,” in *ECCV*, 2020. 2, 3, 4, 5, 6, 8, 9, 10
- [21] H. Wang, Y. Zhu, H. Adam, A. Yuille, and L.-C. Chen, “Max-deeplab: End-to-end panoptic segmentation with mask transformers,” in *CVPR*, 2021. 2, 3, 4, 5, 6
- [22] B. Cheng, A. G. Schwing, and A. Kirillov, “Per-pixel classification is not all you need for semantic segmentation,” in *NeurIPS*, 2021. 2, 5, 6, 8
- [23] W. Zhang, J. Pang, K. Chen, and C. C. Loy, “K-net: Towards unified image segmentation,” in *NeurIPS*, 2021. 2, 4, 5, 6, 7, 9, 10, 11, 12
- [24] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, “Masked-attention mask transformer for universal image segmentation,” in *CVPR*, 2022. 2, 4, 5, 7, 11, 12
- [25] P. Sun, R. Zhang, Y. Jiang, T. Kong, C. Xu, W. Zhan, M. Tomizuka, L. Li, Z. Yuan, C. Wang, and P. Luo, “SparseR-CNN: End-to-end object detection with learnable proposals,” in *CVPR*, 2021. 2, 3, 5, 6, 9
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *NeurIPS*, 2017. 2, 6
- [27] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, “Swin transformer: Hierarchical vision transformer using shifted windows,” in *ICCV*, 2021. 2, 3, 4, 9
- [28] Y. Li, X. Chen, Z. Zhu, L. Xie, G. Huang, D. Du, and X. Wang, “Attention-guided unified network for panoptic segmentation,” in *CVPR*, 2019. 2
- [29] X. Li, S. Xu, Y. Yang, G. Cheng, Y. Tong, and D. Tao, “Panoptic-partformer: Learning a unified model for panoptic part segmentation,” in *ECCV*, 2022. 2, 4, 7, 9, 10
- [30] S. Qi, W. Wang, B. Jia, J. Shen, and S.-C. Zhu, “Learning human-object interactions by graph parsing neural networks,” in *ECCV*, 2018. 3
- [31] S. Liu, L. Zhang, X. Yang, H. Su, and J. Zhu, “Unsupervised part segmentation through disentangling appearance and shape,” in *CVPR*, 2021. 3
- [32] S. Choudhury, I. Laina, C. Rupprecht, and A. Vedaldi, “Unsupervised part discovery from contrastive reconstruction,” in *NeurIPS*, 2021. 3
- [33] H.-S. Fang, G. Lu, X. Fang, J. Xie, Y.-W. Tai, and C. Lu, “Weakly and semi supervised human body part parsing via pose-guided knowledge transfer,” in *CVPR*, 2018. 3
- [34] W. Wang, Z. Zhang, S. Qi, J. Shen, Y. Pang, and L. Shao, “Learning compositional neural information fusion for human parsing,” in *ICCV*, 2019. 3
- [35] Q. Li, A. Arnab, and P. H. Torr, “Holistic, instance-level human parsing,” *arXiv:1709.03612*, 2017. 3
- [36] L. Yang, Q. Song, Z. Wang, and M. Jiang, “Parsing R-CNN for instance-level human analysis,” in *CVPR*, 2019. 3
- [37] R. Ji, D. Du, L. Zhang, L. Wen, Y. Wu, C. Zhao, F. Huang, and S. Lyu, “Learning semantic neural tree for human parsing,” in *ECCV*, 2020. 3
- [38] K. Gong, X. Liang, Y. Li, Y. Chen, M. Yang, and L. Lin, “Instance-level human parsing via part grouping network,” in *ECCV*, 2018. 3
- [39] J. Zhao, J. Li, Y. Cheng, T. Sim, S. Yan, and J. Feng, “Understanding humans in crowded scenes: Deep nested adversarial learning and a new benchmark for multi-human parsing,” in *ACM-MM*, 2018. 3
- [40] T. Zhou, W. Wang, S. Liu, Y. Yang, and L. Van Gool, “Differentiable multi-granularity human representation learning for instance-aware human semantic parsing,” in *CVPR*, 2021. 3, 5
- [41] U. Michieli, E. Borsato, L. Rossi, and P. Zanuttigh, “GMNet: Graph Matching Network for Large Scale Part Semantic Segmentation in the Wild,” in *ECCV*, 2020. 3
- [42] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *ECCV*, 2018. 3, 6, 9, 10
- [43] S. Qiao, L.-C. Chen, and A. Yuille, “Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution,” in *CVPR*, 2021. 3, 9, 10
- [44] R. Hou, J. Li, A. Bhargava, A. Raventos, V. Guizilini, C. Fang, J. Lynch, and A. Gaidon, “Real-time panoptic segmentation from dense detections,” in *CVPR*, 2020. 3
- [45] Z. Li, W. Wang, E. Xie, Z. Yu, A. Anandkumar, J. M. Alvarez, P. Luo, and T. Lu, “Panoptic segformer: Delving deeper into panoptic segmentation with transformers,” in *CVPR*, 2022. 3
- [46] H. Yuan, X. Li, Y. Yang, G. Cheng, J. Zhang, Y. Tong, L. Zhang, and D. Tao, “Polyphonicformer: Unified query learning for depth-aware video panoptic segmentation,” in *ECCV*, 2022. 3
- [47] Q. Yu, H. Wang, S. Qiao, M. Collins, Y. Zhu, H. Adam, A. Yuille, and L.-C. Chen, “k-means mask transformer,” in *ECCV*, 2022. 3
- [48] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, “Object detection with discriminatively trained part-based models,” *TPAMI*, 2010. 3
- [49] R. Fergus, P. Perona, and A. Zisserman, “Object class recognition by unsupervised scale-invariant learning,” in *CVPR*, 2003. 3
- [50] N. Zhang, R. Farrell, F. Iandola, and T. Darrell, “Deformable part descriptors for fine-grained recognition and attribute prediction,” in *CVPR*, 2013. 3
- [51] N. Tritrong, P. Rewatbowornwong, and S. Suwajanakorn, “Repurposing gans for one-shot semantic part segmentation,” in *CVPR*, 2021. 3
- [52] Y. Zhang, H. Ling, J. Gao, K. Yin, J.-F. Lafleche, A. Barriuso, A. Torralba, and S. Fidler, “Datasetgan: Efficient labeled data factory with minimal human effort,” in *CVPR*, 2021. 3
- [53] S. Sabour, A. Tagliasacchi, S. Yazdani, G. Hinton, and D. J. Fleet, “Unsupervised part representation by flow capsules,” in *ICML*, 2021. 3
- [54] Z. Xu, Z. Liu, C. Sun, K. Murphy, W. T. Freeman, J. B. Tenenbaum, and J. Wu, “Unsupervised discovery of parts, structure, and dynamics,” *arXiv:1903.05136*, 2019. 3
- [55] L. Yang, Q. Song, Z. Wang, M. Hu, C. Liu, X. Xin, W. Jia, and S. Xu, “Renovating parsing R-CNN for accurate multiple human parsing,” in *ECCV*, 2020. 3
- [56] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, “The Pascal Visual Object Classes (VOC) Challenge,” *IJCV*, 2010. 3, 8
- [57] C. Tang, L. Xie, X. Zhang, X. Hu, and Q. Tian, “Visual recognition by request,” *arXiv:2207.14227*, 2022. 3, 7, 9, 10

- [58] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv:2010.11929*, 2020. 3, 6
- [59] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, “Training data-efficient image transformers & distillation through attention,” in *ICML*, 2021. 3
- [60] J. Zhang, X. Li, Y. Wang, C. Wang, Y. Yang, Y. Liu, and D. Tao, “Eatformer: improving vision transformer inspired by evolutionary algorithm,” *arXiv:2206.09325*, 2022. 3
- [61] K. Li, Y. Wang, P. Gao, G. Song, Y. Liu, H. Li, and Y. Qiao, “Uni-former: Unified transformer for efficient spatiotemporal representation learning,” *TPAMI*, 2022. 3
- [62] J. Guo, K. Han, H. Wu, C. Xu, Y. Tang, C. Xu, and Y. Wang, “Cmt: Convolutional neural networks meet vision transformers,” in *CVPR*, 2022. 3
- [63] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, “Masked autoencoders are scalable vision learners,” *arXiv:2111.06377*, 2021. 3
- [64] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning transferable visual models from natural language supervision,” in *ICML*, 2021. 3
- [65] G. Ghiasi, X. Gu, Y. Cui, and T.-Y. Lin, “Scaling open-vocabulary image segmentation with image-level labels,” in *ECCV*, 2022. 3
- [66] Z. Chen, Y. Duan, W. Wang, J. He, T. Lu, J. Dai, and Y. Qiao, “Vision transformer adapter for dense predictions,” in *ICLR*, 2023. 3
- [67] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, “Deformable detr: Deformable transformers for end-to-end object detection,” in *ICLR*, 2020. 3, 7
- [68] B. Dong, F. Zeng, T. Wang, X. Zhang, and Y. Wei, “Solq: Segmenting objects by learning queries,” in *NeurIPS*, 2021. 3
- [69] 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 *ECCV*, 2022. 3
- [70] Q. Zhou, X. Li, L. He, Y. Yang, G. Cheng, Y. Tong, L. Ma, and D. Tao, “Transvod: End-to-end video object detection with spatial-temporal transformers,” *TPAMI*, 2022. 3
- [71] T. Meinhardt, A. Kirillov, L. Leal-Taixe, and C. Feichtenhofer, “Trackerformer: Multi-object tracking with transformers,” in *CVPR*, 2022. 3
- [72] X. Li, W. Zhang, J. Pang, K. Chen, G. Cheng, Y. Tong, and C. C. Loy, “Video k-net: A simple, strong, and unified baseline for video segmentation,” in *CVPR*, 2022. 3
- [73] K. Tateno, N. Navab, and F. Tombari, “Distortion-aware convolutional filters for dense prediction in panoramic images,” *ECCV*, 2018. 3
- [74] D. Xu, W. Ouyang, X. Wang, and N. Sebe, “Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing,” in *CVPR*, 2018. 3
- [75] Z. Zhang, Z. Cui, C. Xu, Y. Yan, N. Sebe, and J. Yang, “Pattern-affinitive propagation across depth, surface normal and semantic segmentation,” in *CVPR*, 2019. 3
- [76] C. Shu, Y. Liu, J. Gao, Z. Yan, and C. Shen, “Channel-wise knowledge distillation for dense prediction,” *ICCV*, 2021. 3
- [77] N. Takahashi and Y. Mitsufuji, “Densely connected multi-dilated convolutional networks for dense prediction tasks,” in *CVPR*, 2021. 3
- [78] Z. Ling, C. Zhen, X. Chunyan, Z. Zhenyu, W. Chaoqun, Z. Tong, and Y. Jian, “Pattern-structure diffusion for multi-task learning,” *CVPR*, 2020. 3
- [79] G. Ghiasi, B. Zoph, E. D. Cubuk, Q. V. Le, and T.-Y. Lin, “Multi-task self-training for learning general representations,” in *ICCV*, 2021. 3
- [80] J. N. Kundu, N. Lakkakula, and R. V. Babu, “Um-adapt: Unsupervised multi-task adaptation using adversarial cross-task distillation,” in *ICCV*, 2019. 3
- [81] S. Vandenhenne, S. Georgoulis, and L. Van Gool, “Mti-net: Multi-scale task interaction networks for multi-task learning,” in *ECCV*, 2020. 3
- [82] H. Ye and D. Xu, “Inverted pyramid multi-task transformer for dense scene understanding,” in *ECCV*, 2022. 3
- [83] Y. Xu, X. Li, H. Yuan, Y. Yang, and L. Zhang, “Multi-task learning with multi-query transformer for dense prediction,” *IEEE-TCSVT*, 2023. 3, 12
- [84] H. Ye and D. Xu, “Taskprompter: Spatial-channel multi-task prompting for dense scene understanding,” in *ICLR*, 2023. 3, 12
- [85] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016. 4, 9
- [86] T.-Y. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie, “Feature pyramid networks for object detection,” in *CVPR*, 2017. 4
- [87] X. Li, A. You, Z. Zhu, H. Zhao, M. Yang, K. Yang, and Y. Tong, “Semantic flow for fast and accurate scene parsing,” in *ECCV*, 2020. 5
- [88] X. Wang, T. Kong, C. Shen, Y. Jiang, and L. Li, “SOLO: Segmenting objects by locations,” in *ECCV*, 2020. 5
- [89] X. Wang, R. Zhang, T. Kong, L. Li, and C. Shen, “SOLOv2: Dynamic and fast instance segmentation,” in *NeurIPS*, 2020. 5, 10
- [90] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft coco: Common objects in context,” in *ECCV*, 2014. 5, 9
- [91] F. Milletari, N. Navab, and S. Ahmadi, “V-Net: Fully convolutional neural networks for volumetric medical image segmentation,” in *3DV*, 2016. 8
- [92] H. Bao, L. Dong, and F. Wei, “Beit: Bert pre-training of image transformers,” *arXiv:2106.08254*, 2021. 9
- [93] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *AISTATS*, 2010. 9
- [94] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *ICLR*, 2019. 9
- [95] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo, “Segformer: Simple and efficient design for semantic segmentation with transformers,” in *NeurIPS*, 2021. 10
- [96] R. Mohan and A. Valada, “Efficientps: Efficient panoptic segmentation,” *IJCV*, 2021. 10
- [97] M. Tan and Q. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” in *ICML*, 2019. 10
- [98] Y. Yuan, X. Chen, and J. Wang, “Object-contextual representations for semantic segmentation,” in *ECCV*, 2020. 9, 10
- [99] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang *et al.*, “Deep high-resolution representation learning for visual recognition,” *TPAMI*, 2020. 9, 10
- [100] J. Liang, N. Homayounfar, W.-C. Ma, Y. Xiong, R. Hu, and R. Urtasun, “Polytransformer: Deep polygon transformer for instance segmentation,” in *CVPR*, 2020. 9, 10
- [101] H. Zhang, C. Wu, Z. Zhang, Y. Zhu, H. Lin, Z. Zhang, Y. Sun, T. He, J. Mueller, R. Manmatha *et al.*, “Resnest: Split-attention networks,” in *CVPR*, 2022. 10
- [102] F. Chollet, “Xception: Deep learning with depthwise separable convolutions,” in *CVPR*, 2017. 9
- [103] G. Neuhold, T. Ollmann, S. Rota Bulò, and P. Kontschieder, “The mapillary vistas dataset for semantic understanding of street scenes,” in *ICCV*, 2017. 9, 13
- [104] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, “A convnet for the 2020s,” in *CVPR*, 2022. 9
- [105] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, “Re-thinking atrous convolution for semantic image segmentation,” *arXiv:1706.05587*, 2017. 10, 11
- [106] X. Zhu, H. Hu, S. Lin, and J. Dai, “Deformable convnets v2: More deformable, better results,” in *CVPR*, 2019. 11
- [107] Q. Li, X. Qi, and P. H. Torr, “Unifying training and inference for panoptic segmentation,” in *CVPR*, 2020. 11
- [108] Y. Li, H. Zhao, X. Qi, Y. Chen, L. Qi, L. Wang, Z. Li, J. Sun, and J. Jia, “Fully convolutional networks for panoptic segmentation with point-based supervision,” *TPAMI*, 2022. 11
- [109] F. Milletari, N. Navab, and S.-A. Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation,” in *3DV*, 2016. 10
- [110] X. Li, X. Li, L. Zhang, G. Cheng, J. Shi, Z. Lin, S. Tan, and Y. Tong, “Improving semantic segmentation via decoupled body and edge supervision,” in *ECCV*, 2020. 10
- [111] A. Kirillov, Y. Wu, K. He, and R. Girshick, “Pointrend: image segmentation as rendering. arxiv,” in *CVPR*, 2020. 10
- [112] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell, “Bdd100k: A diverse driving dataset for heterogeneous multitask learning,” in *CVPR*, 2020. 13



Xiangtai Li received the Ph.D. degree from, Peking University in 2022. He currently works as a Research Scientist at Tiktok. He worked as a Research Fellow in S-Lab, Nanyang Technological University. His research interests include computer vision and machine learning with a focus on scene understanding, segmentation, and video understanding. Several of his works have been published in top-tier conferences and journals. He regularly reviews top-tier conferences and journals, including CVPR, ECCV, ICCV, ICLR, T-PAMI, and IJCV.



Shilin Xu is a Ph.D student at Peking University. His research interests include computer vision, machine learning, and scene understanding. He has published several works at related conferences, including ECCV, T-PAMI, and ICIP.



Yibo Yang is a research scientist in King Abdullah University of Science and Technology. He obtained his Ph.D. degree from Peking University in 2021. His research interests lie in machine learning and computer vision. He has published more than 20 publications in top-tier journals and conferences. He was awarded the 2021 Wuwenjun-Outstanding Ph.D. Dissertation Award in CAAI. He was also awarded the first prize in BMTT 2021, the Innovative Award in MS COCO 2019, and the Runner-up Award in MS COCO 2018.



Haobo Yuan is a Ph.D student at UC Merced. He worked as a research associate at S-Lab of Nanyang Technological University. He received his Master degree at Wuhan University in 2023. His research interests include computer vision and machine learning. He has published several works in top-tier conferences and journals, including ICCV, ECCV, CVPR, and T-PAMI.



driving, scene understanding, domain adaptation, and remote sensing image processing.

Guangliang Cheng is currently a Reader in the Department of Computer Science at the University of Liverpool. Previously, He was a Senior Research Manager at SenseTime. Before that, he was a Postdoc researcher at the Institute of Remote Sensing and Digital Earth, Chinese Academy of Sciences, China, and he received his Ph.D. degree from the National Laboratory of Pattern Recognition (NLPR) from the Institute of Automation, Chinese Academy of Sciences, Beijing. His research interests include autonomous



national research and development projects in China since 2018.

Yunhai Tong is currently a Professor with the School of Artificial Intelligence, Peking University. Before that, he received his Ph.D. degree in computer science from Peking University in 2002. His main research interests include media intelligent computing, deep learning, and multimodal learning. Most of his works have been published in top-tier conferences and journals in computer vision, natural language processing, and machine learning, including PAMI, IJCV, EMNLP, ECCV, and CVPR. He has led several



Zhouchen Lin received the PhD degree from Peking University, Beijing, China, in 2000. He is currently a professor with the Key Laboratory of Machine Perception, School of Intelligence Science and Technology, Peking University. His research interests include machine learning and numerical optimization. He is also a fellow of the International Association of Pattern Recognition (IAPR) and China Society of Image and Graphics (CSIG). He has been an area chair of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), International Conference on Computer Vision (ICCV), Annual Conference on Neural Information Processing Systems (NIPS/NeurIPS), AAAI Conference on Artificial Intelligence (AAAI), International Joint Conference on Artificial Intelligence (IJCAI), International Conference on Learning Representations (ICLR), and International Conference on Machine Learning (ICML) many times. He was the program co-chair of International Conference on Pattern Recognition (ICPR) 2022 and the senior area chair of ICML 2022, NeurIPS 2022 and CVPR 2023. He is an associate editor of the International Journal of Computer Vision and Optimization Methods and Software.



Ming-Hsuan Yang received the PhD degree from the University of Illinois at Urbana-Champaign. He is a professor of electrical engineering and computer science at the University of California, Merced, and an adjunct professor at Yonsei University. He served as a program co-chair for the 2019 IEEE International Conference on Computer Vision.

He received the Longuet-Higgins Prize in 2023, NSF CAREER award in 2012 and Google faculty award in 2009. He is a fellow of the ACM.



Dacheng Tao (Fellow, IEEE) is with the Nanyang Technological University. He is also an advisor and chief scientist of the digital science institute at the University of Sydney. He mainly applies statistics and mathematics to artificial intelligence and data science, and his research is detailed in one monograph and over 200 publications in prestigious journals and proceedings at leading conferences. He received the 2015 Australian Scopus-Eureka Prize, the 2018 IEEE ICDM Research Contributions Award, and the 2021 IEEE Computer Society McCluskey Technical Achievement Award. He is a fellow of the Australian Academy of Science, AAAS, ACM, and IEEE.