

Instances as Queries

Yuxin Fang^{1*}, Shusheng Yang^{1,2*}, Xinggang Wang^{1†}, Yu Li²,
Chen Fang³, Ying Shan², Bin Feng¹, Wenyu Liu¹

¹School of EIC, Huazhong University of Science & Technology

²Applied Research Center (ARC), Tencent PCG ³Tencent

Abstract

We present *QueryInst*, a new perspective for instance segmentation. *QueryInst* is a multi-stage end-to-end system that treats instances of interest as learnable queries, enabling query based object detectors, e.g., Sparse R-CNN, to have strong instance segmentation performance. The attributes of instances such as categories, bounding boxes, instance masks, and instance association embeddings are represented by queries in a unified manner. In *QueryInst*, a query is shared by both detection and segmentation via dynamic convolutions and driven by parallelly-supervised multi-stage learning. We conduct extensive experiments on three challenging benchmarks, i.e., COCO, CityScapes, and YouTube-VIS to evaluate the effectiveness of *QueryInst* in object detection, instance segmentation, and video instance segmentation tasks. For the first time, we demonstrate that a simple end-to-end query based framework can achieve the state-of-the-art performance in various instance-level recognition tasks. Code is available at <https://github.com/hustvl/QueryInst>.

1. Introduction

Instance segmentation is a fundamental yet challenging computer vision task that requires an algorithm to assign a pixel-level mask with a category label for each instance of interest in image. Prevalent state-of-the-art instance segmentation methods are based on high performing object detectors and follow a multi-stage paradigm. Among which, the Mask R-CNN family [22, 25, 32, 5, 10, 44] is the most successful one, where the regions-of-interest (RoI) for instance segmentation is extracted via a region-wise pooling operation (e.g., RoIPool [23, 20] or RoIAlign [22]) based on the box-level localization information from the region proposal network (RPN) [41], or the previous stage bounding-

*Equal contributions. This work was done while Shusheng Yang was interning at Applied Research Center (ARC), Tencent PCG.

†Corresponding author, E-mail: xgwang@hust.edu.cn.

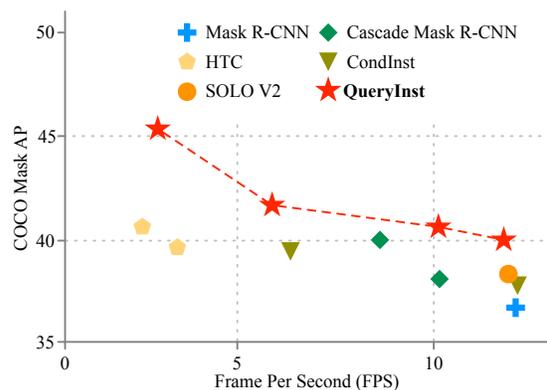


Figure 1: AP vs. FPS on COCO test-dev. *QueryInst* outperforms current state-of-the-art methods in terms of both accuracy and speed. The speed is measured using a single Titan Xp GPU.

box prediction [4, 5]. The final instance mask is obtained via feeding the RoI feature into the mask head, which is a small fully convolutional network (FCN) [35].

Recently, DETR [7] is proposed to reformulate object detection as a query based direct set prediction problem, whose input is mere 100 learned object queries. Follow-up works [63, 44, 45, 19, 60, 16] in object detection improve this query based approach and achieve comparable performance with state-of-the-art detectors such as Cascade R-CNN [4]. The results show that query based instance-level perception is a very promising research direction. Thus, enabling query based detection framework to perform instance segmentation is highly desirable. However, we find that it is inefficient to integrate the previous successful practices in Cascade Mask R-CNN [5] and HTC [10], which are state-of-the-art mask generation solutions in the non-query based paradigm, directly into query based detectors for instance mask generation. Therefore, an instance segmentation method tailored for the query based end-to-end framework is urgently needed.

To bridge this gap, we propose **QueryInst** (Instances as

Queries), a query based end-to-end instance segmentation method driven by parallel supervision on dynamic mask heads [26, 46, 44]. The key insight of QueryInst is to leverage *the intrinsic one-to-one correspondence in object queries* across different stages, and *one-to-one correspondence between mask RoI features and object queries* in the same stage. Specifically, we set up dynamic mask heads in parallel with each other, which transform each mask RoI feature adaptively according to the corresponding query, and are simultaneously trained in all stages. The mask gradient not only flows back to the backbone feature extractor, but also to the object query, which is intrinsically one-to-one interlinked in different stages. The queries implicitly carry the multi-stage mask information, which is read by RoI features in dynamic mask heads for final mask generation. There is no explicit connection between different stage mask heads or mask features. Moreover, the queries are shared between object detection and instance segmentation sub-networks in each stage, enabling cross-task communications that one task can take advantage of the information from the other task. We demonstrate that this shared query design can fully leverage the synergy between object detection and instance segmentation. When the training is completed, we throw away all the dynamic mask heads in the intermediate stages and only use the final stage predictions for inference. Under such a scheme, QueryInst surpasses the state-of-the-art HTC in terms of AP while runs much faster. Concretely, our main contributions are summarized as follows:

- We attempt to solve instance segmentation from a new perspective that uses parallel dynamic mask heads in the query based end-to-end detection framework. This novel solution enables such a new framework to outperform well-established and highly-optimized non-query based multi-stage schemes such as Cascade Mask R-CNN and HTC in terms of both accuracy and speed (see Fig. 1). Our best model achieves 56.1 AP^{box} and 49.1 AP^{mask} on COCO test-dev.
- We set up a task-joint paradigm for query based object detection and instance segmentation by leveraging the shared query and multi-head self-attention design. This paradigm establishes a kind of communication and synergy between detection and segmentation tasks, which encourages these two tasks to benefit from each other. We demonstrate that our architecture design can also significantly improve the object detection performance.
- We extend the QueryInst to video instance segmentation task (VIS) [59] task by simply adding a vanilla track head. Experiments on YouTube-VIS dataset [59] indicate that with same tracking approach, our methods outperforms MaskTrack R-CNN [59] and

SipMask-VIS [6] by a large margin. QueryInst-VIS can even outperform well-designed VIS approaches such as STEM-Seg [1] and VisTR [55].

2. Related Work

Query Based Methods. Recently, query based methods emerged to tackle the set-prediction problems. Concretely, DETR [7] first introduces the query based methods with transformer architecture to object detection. Deformable DETR [63], UP-DETR [16], ACT [60] and TSP [45] improve the performance on the top of DETR. The recently proposed Sparse R-CNN [44] builds a query based set-prediction framework upon R-CNN [21, 20, 41] based detector. For segmentation, VisTR [55] introduces a query based sequence matching and segmentation method to video instance segmentation, building a fully end-to-end framework for instance segmentation in video. Max-DeepLab[51] presents the first box-free end-to-end panoptic segmentation model with a global memory as external query. Trackformer [37] and Transtrack [43] build a query based multiple object tracker upon DETR and Deformable DETR, respectively, and attain comparable results to the non-query based methods. AS-Net [12] introduces a query based set-prediction pipeline to human object interaction and obtains promising results. Despite query based set-prediction method is being widely used to many computer vision tasks, few efforts are conducted to build a successful query based instance segmentation framework. We aim to achieve this goal in this paper.

Object Detection. Object detection is a fundamental computer vision task which aims to detect visual objects with bounding boxes. With the propose of R-CNN [21], Fast R-CNN [20] and Faster R-CNN [41], anchor based methods [4, 40, 36, 30, 33] dominate object detection for a long period. CenterNet [61] and FCOS [47] establish anchor-free detectors with competitive detection performance. Recently, with the proposed DETR [7], query based set-prediction methods catch lots of attentions. Deformable DETR [63] introduces deformable convolution [62] to the DETR framework, achieving better performance with faster training convergence. UP-DETR [16] extends DETR to unsupervised scenarios. ACT [60] and TSP [45] introduce the adaptive clustering module and a new bipartite matching method to DETR. Sparse R-CNN [44] build a query based detector on top of R-CNN architecture, while OneNet [42] and DeFCN [52] are end-to-end detector built upon the one-stage FCOS [47]. In this work, we present a query based instance segmentation method on the top of the query based Sparse R-CNN detector.

Instance Segmentation. Instance segmentation is a fundamental yet challenging computer vision task that requires an algorithm to assign a pixel-level mask with a category label for each instance of interest in image. Mask R-CNN [22] introduces a fully convolutional mask head to Faster R-CNN [20] detector. Cascade Mask R-CNN [5] simply combine the Cascade R-CNN [4] with Mask R-CNN. HTC [10] presents interleaved execution and mask information flow and achieves state-of-the-art performance. Follow-up works [25, 13, 50, 27] also contribute to the Mask R-CNN family. In addition to R-CNN based methods, YOLACT [3, 2], SipMask [6], CondInst [48], BlendMask [9] and SOLO [53, 54] build one-stage instance segmentation framework on the top of one-stage framework, achieving comparable results with favorable inference speed. Following the R-CNN based methods, we present a query based instance segmentation framework.

3. Instances as Queries

We propose QueryInst (Instances as Queries), a query based end-to-end instance segmentation method. QueryInst consists of a query based object detector and six dynamic mask heads driven by parallel supervision. Our key insight is to leverage the *intrinsic one-to-one correspondence in queries* across different stages. This correspondence exists in all query based framework [49, 17, 39, 8, 7] regardless of the specific instantiations and applications. The overall architecture of QueryInst is illustrated in Fig. 2 (c).

3.1. Query based Object Detector

QueryInst can be built on any multi-stage query based object detector [7, 63, 44]. We choose Sparse R-CNN [44] as our default instantiation, which has six query stages. The object detection pipeline is depicted in Fig. 2 (a) and can be formulated as follows:

$$\begin{aligned}
 \mathbf{x}_t^{\text{box}} &\leftarrow \mathcal{P}^{\text{box}}(\mathbf{x}^{\text{FPN}}, \mathbf{b}_{t-1}), \\
 \mathbf{q}_{t-1}^* &\leftarrow \text{MSA}_t(\mathbf{q}_{t-1}), \\
 \mathbf{x}_t^{\text{box}*}, \mathbf{q}_t &\leftarrow \text{DynConv}_t^{\text{box}}(\mathbf{x}_t^{\text{box}}, \mathbf{q}_{t-1}^*), \\
 \mathbf{b}_t &\leftarrow \mathcal{B}_t(\mathbf{x}_t^{\text{box}*}),
 \end{aligned} \tag{1}$$

where $\mathbf{q} \in \mathbf{R}^{N \times d}$ denotes the object query. N and d denote the length (number) and dimension of query \mathbf{q} . At stage t , a pooling operator \mathcal{P}^{box} extracts the current stage bounding box features $\mathbf{x}_t^{\text{box}}$ from FPN [29] features \mathbf{x}^{FPN} under the guidance of previous stage bounding box predictions \mathbf{b}_{t-1} . Meanwhile, a multi-head self-attention module MSA_t is applied to the input query \mathbf{q}_{t-1} to get the transformed query \mathbf{q}_{t-1}^* . Then, a box dynamic convolution module $\text{DynConv}_t^{\text{box}}$ takes $\mathbf{x}_t^{\text{box}}$ and \mathbf{q}_{t-1}^* as inputs and enhances the $\mathbf{x}_t^{\text{box}}$ by reading \mathbf{q}_{t-1}^* while generating \mathbf{q}_t for the next stage. Finally, the enhanced bounding box features

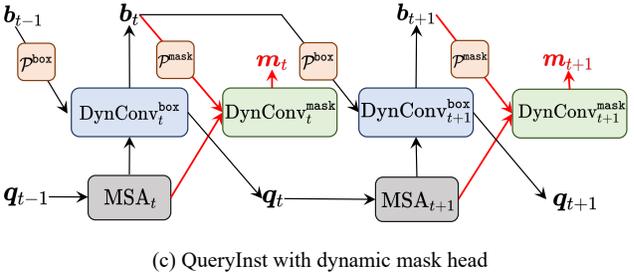
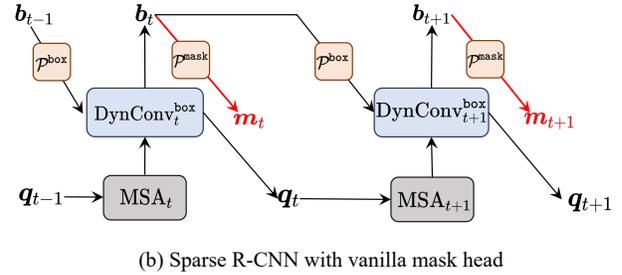
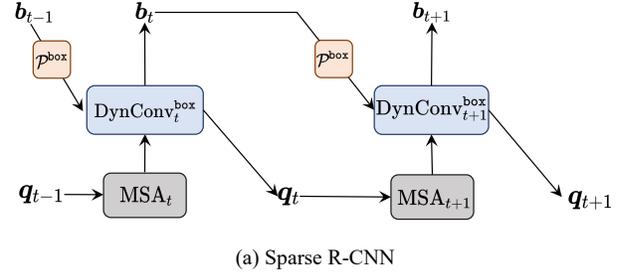


Figure 2: Overview of QueryInst. The red arrows indicate mask branches. Please note that QueryInst consists of 6 stage in parallel, *i.e.*, $t = \{1, 2, 3, 4, 5, 6\}$. The figure only shows 2 stages.

$\mathbf{x}_t^{\text{box}*}$ are fed into the box prediction branch \mathcal{B}_t for current bounding box prediction \mathbf{b}_t .

3.2. Mask Head Architecture

3.2.1 Vanilla Mask Head

For instance mask prediction, we first adopt the widely used vanilla mask head architecture design in Mask R-CNN [22] as our instance segmentation baseline. The model architecture is depicted in Fig. 2 (b). Based on the object detection pipeline described in Sec. 3.1, the mask generation process can be expressed as follows:

$$\begin{aligned}
 \mathbf{x}_t^{\text{mask}} &\leftarrow \mathcal{P}^{\text{mask}}(\mathbf{x}^{\text{FPN}}, \mathbf{b}_t), \\
 \mathbf{m}_t &\leftarrow \mathcal{M}_t(\mathbf{x}_t^{\text{mask}}),
 \end{aligned} \tag{2}$$

where \mathbf{b}_t is the bounding box predictions from the object detector. $\mathcal{P}^{\text{mask}}$ denotes a region-wise pooling operator for mask RoI features extraction. \mathcal{M}_t indicates the mask FCN head consisting of a stack of four consecutive conv-layers,

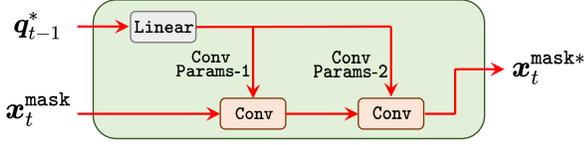


Figure 3: Illustrations of $\text{DynConv}_t^{\text{mask}}$ at stage t . $x_t^{\text{mask}*}$ is enhanced by two consecutive conv-layers, whose kernel parameters are produced by q_{t-1}^* .

one dconv-layer and one 1×1 conv-layer for mask generation [22]. m_t is the current stage mask predictions.

Overall, this vanilla design is an analogy of Cascade Mask R-CNN[5] in a query based framework. However, we find that this design is not as effective as the original Cascade Mask R-CNN. Moreover, establishing explicit mask flow following HTC [10] on top of this design (Fig. 2 (b)) can only bring moderate improvements at a cost of large drops in both training and inference speed. Part of the reasons may be the number of queries in our framework is much smaller than the number of proposals in Cascade Mask R-CNN and HTC, resulting in limited availability of training samples.

3.2.2 Dynamic Mask Head

Our goal is to design a mask prediction head tailor-made for query based instance segmentation frameworks. To this end, we propose to leverage dynamic mask heads driven by parallel supervision to replace the vanilla design in Sec. 3.2.1. The dynamic mask head at stage t consists of a dynamic mask convolution module $\text{DynConv}_t^{\text{mask}}$ (see Fig. 3) [44] following by a vanilla mask head \mathcal{M}_t [22]. The mask generation pipeline is reformulated as follows:

$$\begin{aligned} x_t^{\text{mask}} &\leftarrow \mathcal{P}^{\text{mask}}(x^{\text{FPN}}, b_t), \\ x_t^{\text{mask}*} &\leftarrow \text{DynConv}_t^{\text{mask}}(x_t^{\text{mask}}, q_{t-1}^*), \\ m_t &\leftarrow \mathcal{M}_t(x_t^{\text{mask}*}). \end{aligned} \quad (3)$$

It is noteworthy that the only difference between the proposed dynamic mask head and vanilla mask head is the existence of $\text{DynConv}_t^{\text{mask}}$. We demonstrate that $\text{DynConv}_t^{\text{mask}}$ enables (1) per-mask information flow in queries driven by parallel mask branch supervision, and (2) communication and synergy for joint detection and instance segmentation in the following two sub-sections, respectively. The effectiveness of these two properties is verified in our experiments.

3.3. Per-mask Information Flow with Parallel Supervision

In query based models such as [7, 63, 44], the model learns different specialization for each query slot [7], *i.e.*, $q_t[s]$ is the transformed and refined version of previous

stage $q_{t-1}[s]$ in the same s -th slot. Moreover, $x_t^{\text{mask}}[s]$ corresponds to and is refined by $q_t[s]$ [44]. Therefore there is an one-to-one correspondence across different stage queries inherent in these frameworks, as well as one-to-one correspondence between mask ROI features and object queries in the same stage.

QueryInst is driven by parallel supervision on dynamic mask heads, which fully leverages the intrinsic one-to-one correspondence in object queries across different stages. Specifically, we set up dynamic mask heads in parallel with each other, which transform each mask ROI feature x_t^{mask} adaptively in $\text{DynConv}_t^{\text{mask}}$ according to the corresponding query q_{t-1}^* , and are simultaneously trained in all stages. Inside $\text{DynConv}_t^{\text{mask}}$, the query acts as memory and is read by mask ROI features x_t^{mask} in the forward pass and written by x_t^{mask} in the backward pass.

During training, the per-mask information (*i.e.*, the mask gradient) not only flows back to mask ROI features x_t^{mask} , but also to the object query q_{t-1}^* , which is intrinsically one-to-one interlinked in different stages. Therefore the per-mask information flow is naturally established by leveraging the inherent properties of query based frameworks, with no additional connection needed. After the training is completed, the information for mask prediction is stored in queries.

During inference, we throw away all dynamic mask heads in 5 intermediate stages and only use the final stage predictions for inference. The queries implicitly carry the multi-stage information for mask prediction, which is read by mask ROI features x_t^{mask} in dynamic mask convolution $\text{DynConv}_t^{\text{mask}}$ at the last stage for final mask generation.

Without $\text{DynConv}_t^{\text{mask}}$, the link between mask ROI features and the query is lost, and mask heads in different stages are isolated. Even though parallel supervision is applied to all mask heads, the information related to mask generation cannot flow into queries. In this condition, QueryInst degenerates to Cascade Mask R-CNN with a fixed number (*i.e.*, N) of proposals across all stages.

3.4. Shared Query and MSA for Joint Detection and Segmentation

At stage t , a multi-head self-attention MSA_t is applied to the query q_{t-1} . MSA_t projects the query q_{t-1} to a high dimensional embedding space and its output q_{t-1}^* is read by the dynamic box convolution $\text{DynConv}_t^{\text{box}}$ and dynamic mask convolution $\text{DynConv}_t^{\text{mask}}$ respectively, to enhance the task-specific features x_t^{box} and x_t^{mask} .

Throughout this process, the query and MSA are shared between detection and instance segmentation tasks. Both detection and segmentation information flow back into the query through MSA. This task-joint paradigm establishes a kind of communication and synergy between detection and segmentation tasks, which encourages these two tasks to

benefits from each other. The query learns a better instance-level representation with the guidance of two highly correlated tasks. We observe a performance decrease using separate queries or MSA in our experiments.

3.5. QueryInst-VIS for Video Instance Segmentation

Video instance segmentation (VIS) [59] is a highly relevant task to still-image instance segmentation that aims at detecting, classifying, segmenting and tracking visual instances over video frames. We demonstrate that QueryInst can be easily extended to VIS with minimal modifications by simply adding the vanilla track head in MaskTrack R-CNN baseline [59]. The proposed model coined as QueryInst-VIS can perform video instance segmentation in an online manner while operating at real-time. The total training and inference pipeline keep the same as MaskTrack R-CNN. We evaluate QueryInst-VIS on the challenging YouTube-VIS [59] benchmark to demonstrate its effectiveness.

4. Experiments

4.1. Dataset

COCO. Most of our experiments are conducted on the challenging COCO dataset [31]. Following the common practice, we use the COCO train2017 split (115k images) for training and the val2017 split (5k images) as validation for our ablation study. We report our main results on the test-dev split (20k images).

Cityscapes. Cityscapes [14] is an ego-centric street-scene dataset with 8 categories, 2975 train images, and 500 validation images for instance segmentation. The images are with higher resolution (1024×2048 pixels) compared with COCO, and have more pixel-accurate ground-truth.

YouTube-VIS. In addition to static-image instance segmentation, we demonstrate the effectiveness of our QueryInst on video instance segmentation. YouTube-VIS [59] is a challenging dataset for video instance segmentation task, which has a 40-category label set, 4,883 unique video instances and 131k high-quality manual annotations. There are 2,238 training videos, 302 validation videos, and 343 test videos.

4.2. Implementation Details

Training Setup. Our implementation is based on MMDetection [11] and Detectron2 [56]. Following [44], the default training schedule is 36 epochs and the initial learning rate is set to 2.5×10^{-5} , divided by 10 at 27-th epoch and 33-th epoch, respectively. We adopt AdamW optimizer with 1×10^{-4} weight decay. Hyper-parameters, configurations as well as the label assignment procedures follow the setting in [7, 63, 44]. In total, the R-CNN head of QueryInst contains 6 stages in parallel as [44]. The mask head is trained

by minimizing dice loss [38]. Without special mentioning, we adopt QueryInst model trained with 100 queries and ResNet-50-FPN [24, 29] as backbone in our experiments in the ablation studies.

Inference. Given an input image, QueryInst directly outputs top 100 bounding box predictions with their scores and corresponding instance masks without further post-processing. For inference, we use the final stage masks as the predictions and ignore all the parallel $\text{DynConv}^{\text{mask}}$ at the intermediate stages. The inference speed reported is measured using a single Titan Xp GPU with inputs resized to have their shorter side being 800 and their longer side less or equal to 1333.

4.3. Main Results

Comparisons on COCO Instance Segmentation.

The comparison of QueryInst with the state-of-the-art instance segmentation methods on COCO test-dev are listed in Tab. 1. We have tested different backbones and data augmentations. CondInst [48] (with auxiliary semantic branch) and SOLOv2 [54] are the latest state-of-the-art instance segmentation approach based on dynamic convolutions. A 5-stage QueryInst trained with 100 queries outperforms them with over 1.1 mask AP gain under similar inference speed. QueryInst trained with 100 queries can also surpass Cascade Mask R-CNN [22] by 1.5 mask AP while runs with the same FPS. For fair comparisons with HTC [10], we train HTC w/o semantic segmentation branch using the 36 epochs training schedule and multi-scale data augmentations following the standard setting in [22, 56], yielding ~ 1 higher mask AP than original results reported in [10]. Under same experimental conditions, QueryInst outperforms the state-of-the-art HTC in terms of both accuracy and speed. Moreover, QueryInst outperforms HTC in terms of AP at different IoU thresholds (AP_{50} and AP_{75}) as well as AP at different scales (AP_S , AP_M and AP_L), regardless of the experimental configuration. We also find that compared with Cascade Mask R-CNN and HTC, the query based QueryInst can benefit more from stronger data argumentation used in [7, 63, 44]. Specifically, using ResNet-101-FPN [29] backbone and stronger multi-scale data argumentation with random crop, QueryInst surpasses HTC by 2.0 mask AP and 1.8 box AP while runs $2.4\times$ faster. Further, QueryInst with deformable ResNeXt-101-FPN backbone [57, 15, 62] achieves 44.6 mask AP and 50.4 box AP without bells and whistles.

We demonstrate that the instance segmentation performance of QueryInst not simply come from the accurate bounding box provided by Sparse R-CNN [44] object detector. On the contrary, QueryInst can largely improve the detection performance. The best result of Sparse R-CNN (ResNet-101-FPN, 300 queries, $480 \sim 800$ w/ crop, 36 epochs) reported in [11] is 46.3 box AP. Under the same

| Method | Backbone | Aug. | Epochs | AP ^{box} | AP | AP ₅₀ | AP ₇₅ | AP _S | AP _M | AP _L | FPS |
|---|----------------|-----------------------|--------|---------------------------|----------------------|------------------|------------------|-----------------|-----------------|-----------------|------------------|
| Mask R-CNN [22] | ResNet-50-FPN | 640 ~ 800 | 36 | 41.3 | 37.5 | 59.3 | 40.2 | 21.1 | 39.6 | 48.3 | 14.0 |
| CondInst w/ sem. [48] | | | | – | 38.6 | 60.2 | 41.4 | 20.6 | 41.0 | 51.1 | 14.1 |
| SOLOv2 [53] | | | | 40.4 | 38.8 | 59.9 | 41.7 | 16.5 | 41.7 | 56.2 | 13.8 |
| QueryInst (5 Stage, 100 Queries) | | | | 44.5 | 39.9 | 62.2 | 43.0 | 22.9 | 41.7 | 51.9 | 13.5 |
| Cascade Mask R-CNN [5] | ResNet-50-FPN | 640 ~ 800 | 36 | 44.5 | 38.6 | 60.0 | 41.7 | 21.7 | 40.8 | 49.6 | 10.4 |
| HTC [10] | | | | 44.9 | 39.7 | 61.4 | 43.1 | 22.6 | 42.2 | 50.6 | 3.1 |
| QueryInst (100 Queries) | | | | 44.8 | 40.1 | 62.3 | 43.4 | 23.3 | 42.1 | 52.0 | 10.5 |
| QueryInst (300 Queries) | | | | 45.6 | 40.6 | 63.0 | 44.0 | 23.4 | 42.5 | 52.8 | 7.0 |
| Cascade Mask R-CNN | ResNet-101-FPN | 640 ~ 800 | 36 | 45.7 | 39.8 | 61.6 | 43.0 | 22.4 | 42.2 | 50.8 | 8.7 |
| HTC | | | | 46.2 | 40.7 | 62.7 | 44.2 | 23.1 | 43.4 | 52.7 | 2.5 |
| QueryInst (300 Queries) | | | | 47.0 | 41.7 | 64.4 | 45.3 | 24.2 | 43.9 | 53.9 | 6.1 |
| Cascade Mask R-CNN | | | | 46.2 | 40.0 | 61.7 | 43.5 | 22.5 | 42.5 | 51.2 | 8.7 |
| HTC | ResNet-101-FPN | 480 ~ 800 w/ crop | 36 | 46.3 | 40.8 | 62.6 | 44.3 | 23.0 | 43.5 | 52.6 | 2.5 |
| Sparse R-CNN (300 Queries) | | | | 46.3 | – | – | – | – | – | – | 6.9 |
| QueryInst (300 Queries) | | | | 48.1 | 42.8 | 65.6 | 46.7 | 24.6 | 45.0 | 55.5 | 6.1 |
| QueryInst (300 Queries) | | | | ResNeXt-101-FPN w/ DCN | 480 ~ 800 w/ crop | 36 | 50.4 | 44.6 | 68.1 | 48.7 | 26.6 |
| QueryInst (300 Queries) @ val | Swin-L | 400 ~ 1200 w/ crop | 50 | 56.1 | 48.9 | 74.0 | 53.9 | 30.8 | 52.6 | 68.3 | 3.3 ^T |
| QueryInst (300 Queries) | Swin-L | 400 ~ 1200 w/ crop | 50 | 56.1 | 49.1 | 74.2 | 53.8 | 31.5 | 51.8 | 63.2 | 3.3 ^T |

Table 1: Main results on COCO test-dev. The numbers under “Aug.” indicate the scale range of the shorter size of inputs with a stride of 32. AP^{box} denotes box AP. AP without superscript denotes mask AP. The best results are in **bold** for each configuration. Superscript “T” indicates the FPS data are measured on a single RTX 2080Ti GPU with batch size 1.

experimental setting, QueryInst can achieve 48.1 box AP, which outperform Sparse R-CNN by 1.8 box AP. We also show in the ablation study that QueryInst can outperform Cascade Mask R-CNN and HTC based on a weaker query based detector.

We also apply QueryInst to the recent state-of-the-art Swin Transformer [34] backbone without further modifications, and we find the proposed model is quite capable of adapting with Swin-L. Without bells and whistles, QueryInst can achieve the state-of-the-art performance in instance segmentation as well as object detection. For the first time, we demonstrate that an end-to-end query based framework driven by parallel supervision can achieve the state-of-the-art performance in various instance-level recognition tasks.

Comparisons on Cityscapes Instance Segmentation.

We also conduct experiments on Cityscapes dataset to demonstrate the generalization of QueryInst. Following the standard setting in [48, 22], all models are first pre-trained on COCO train2017 split then finetuned on Cityscapes using fine annotations for 24k iterations with batch size 8. The initial learning rate is linearly scaled to 1.25×10^{-5} and is reduced by a factor of 10 at step 18k.

The results are shown in Tab. 2. QueryInst achieves 39.4 AP on val split and 34.4 AP on test split, sur-

passing several strong baselines. Notably, compared to the dynamic convolution based method CondInst [48], QueryInst with ResNet-50 backbone outperforms CondInst with both ResNet-101-DCN-BiFPN backbone and semantic branch. Overall, our QueryInst achieves leading results on Cityscapes dataset without bells and whistles.

Video Instance Segmentation Results on YouTube-VIS.

Tab. 3 shows the video instance segmentation results on YouTube-VIS val set. Following the standard setting in [59], we first pre-train the instance segmentation model on COCO train2017, then we finetune the corresponding VIS model on YouTube-VIS train set for 12 epochs. The maximum number of instances in one frame in YouTube-VIS dataset is 10, so we set the number of queries to 10 in QueryInst-VIS. The setting enables the model to operate at real-time (> 30 FPS).

As mentioned in Sec. 3.5, QueryInst-VIS adopts the vanilla track method of MaskTrack R-CNN [59] and SipMask-VIS [6], while it obtains 4.3 AP improvement compared to MaskTrack R-CNN and 2.1 AP improvement compared to SipMask-VIS. Moreover, QueryInst can outperform many well-established and highly-optimized VIS approaches, such as SEm-Seg, CompFeat and VisTR in terms of both accuracy and speed.

| Method | Backbone | AP _{val} | AP | AP ₅₀ | person | rider | car | trunk | bus | train | mcycle | bicycle |
|-----------------------|---------------|-------------------|-------------|------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| Mask R-CNN [22] | ResNet-50 | 36.4 | 32.0 | 58.1 | 34.8 | 27.0 | 49.1 | 30.1 | 40.9 | 30.9 | 24.1 | 18.7 |
| BShapeNet+ [28] | ResNet-50 | – | 32.9 | 58.8 | 36.6 | 24.8 | 50.4 | 33.7 | 41.0 | 33.7 | 25.4 | 17.8 |
| UPSNet [58] | ResNet-50 | 37.8 | 33.0 | 59.7 | 35.9 | 27.4 | 51.9 | 31.8 | 43.1 | 31.4 | 23.8 | 19.1 |
| CondInst [48] | ResNet-50 | 37.5 | 33.2 | 57.2 | 35.1 | 27.7 | 54.5 | 29.5 | 42.3 | 33.8 | 23.9 | 18.9 |
| CondInst [48] w/ sem. | DCN-101-BiFPN | 39.3 | 33.9 | 58.2 | 35.6 | 28.1 | 55.0 | 32.1 | 44.2 | 33.6 | 24.5 | 18.6 |
| QueryInst | ResNet-50 | 39.4 | 34.4 | 59.6 | 40.4 | 30.7 | 56.8 | 29.1 | 40.5 | 30.8 | 26.0 | 21.1 |

Table 2: Instance segmentation results on Cityscapes val (AP_{val} column) and test (remain columns) split. The best results are in **bold**.

| Method | Backbone | AP | AP ₅₀ | AP ₇₅ | AR ₁ | AR ₁₀ | FPS |
|-----------------------|------------|-------------|------------------|------------------|-----------------|------------------|-------------|
| MaskTrack R-CNN [59] | ResNet-50 | 30.3 | 51.1 | 32.6 | 31.0 | 35.5 | 22.1 |
| SipMask-VIS [6] | ResNet-50 | 32.5 | 53.0 | 33.3 | 33.5 | 38.9 | 30.9 |
| SipMask-VIS* | ResNet-50 | 33.7 | 54.1 | 35.8 | 35.4 | 40.1 | 30.9 |
| STEm-Seg [1] | ResNet-50 | 30.6 | 50.7 | 33.5 | 31.6 | 37.1 | 4.4 |
| STEm-Seg | ResNet-101 | 34.6 | 55.8 | 37.9 | 34.4 | 41.6 | 2.1 |
| CompFeat [18] | ResNet-50 | 35.3 | 56.0 | 38.6 | 33.1 | 40.3 | – |
| VisTR [55] | ResNet-50 | 34.4 | 55.7 | 36.5 | 33.5 | 38.9 | 30.0 |
| VisTR | ResNet-101 | 35.3 | 57.0 | 36.2 | 34.3 | 40.4 | 27.7 |
| QueryInst-VIS | ResNet-50 | 34.6 | 55.8 | 36.5 | 35.4 | 42.4 | 32.3 |
| QueryInst-VIS* | ResNet-50 | 36.2 | 56.7 | 39.7 | 36.1 | 42.9 | 32.3 |

Table 3: Comparisons with state-of-the-art video instance segmentation methods on YouTube-VIS val set. Methods with superscript “*” indicates using multi-scale data argumentation during training. The best results are in **bold**.

4.4. Ablation Study

Study of Parallel Supervision and DynConv^{mask}.

We show that applying parallel mask head supervision and DynConv^{mask} are both indispensable for good performance. As shown in Tab. 5, using parallel supervision on vanilla mask head cannot bring large improvement, because mask heads in different stages are isolated and there is no cross-stage per-mask information flow established (Sec. 3.2.1). Using DynConv^{mask} without parallel supervision on each stage can only bring moderate improvement, for mask gradients injected from the final stage cannot fully driven the per-mask information flow across queries in all stages. When DynConv^{mask} in all stages are driven by parallel supervision simultaneously, QueryInst achieves significant improvement in accuracy with only little drops in inference speed. The reason is that during inference, we throw away all the parallel DynConv^{mask} in the intermediate stage and only use the final stage mask predictions. The per-mask information is written and preserved in queries during training, which only need to be read out at the final stage during inference.

Study of Query and MSA.

Tab. 6 studies the impact of using shared query and MSA. As expected in Sec. 3.4, using shared query and MSA simultaneously establishes a kind of communication and syn-

ergy between detection and segmentation tasks, which encourages this two tasks to benefits from each other and achieves the highest box AP and mask AP. Moreover, this configuration consumes minimal parameters and computation budgets. Therefore we choose using shared query and MSA as the default instantiation of our QueryInst.

Study of Different Mask Heads.

Tab. 4 studies the impact of different mask head architectures on query and non-query based frameworks. All stages are simultaneously trained. For non-query based frameworks, the 1-st row is the results of Cascade Mask R-CNN [5] and the 2-nd row is HTC [10]. We have the following 3 major observations.

First, we find that directly integrating cascade mask head [5] and HTC mask flow [10] into the query based model is not as effective as in its original framework. When cascade mask head is applied (3-th row), the query based model is 0.5 AP^{box} and 0.6 AP^{mask} lower than the original Cascade Mask R-CNN (1-th row). When HTC mask flow is applied (4-th row), the query based model is 0.6 AP^{box} and 0.4 AP^{mask} lower than the original HTC (2-th row). These results demonstrate that previous successful empirical practice from non-query based multi-stage models is possibly inadequate for query based models (Sec. 3.2.1).

Second, when the proposed parallel DynConv^{mask} is ap-

| Type | Cascade Mask Head [5] | HTC Mask Flow [10] | DynConv ^{mask} | Fig. | AP ^{box} | Δ^{box} | AP ^{mask} | Δ^{mask} | FPS | Δ^{FPS} |
|-----------------|-----------------------|--------------------|-------------------------|------------|-------------------|-----------------------|--------------------|------------------------|------|-----------------------|
| Non-query Based | ✓ | | | | 44.3 | | 38.5 | | 10.4 | |
| | | ✓ | | | 44.4 | + 0.1 | 39.3 | + 0.8 | 3.1 | - 7.3 |
| Query Based | ✓ | | | Fig. 2 (b) | 43.8 | | 37.9 | | 11.1 | |
| | | ✓ | | | 43.8 | + 0.0 | 38.9 | + 1.0 | 6.0 | - 5.1 |
| | | | ✓ | Fig. 2 (c) | 44.5 | + 0.7 | 39.8 | + 1.9 | 10.5 | - 0.6 |
| | | ✓ | ✓ | | 44.4 | + 0.6 | 40.0 | + 2.1 | 5.4 | - 5.7 |

Table 4: Impacts of different mask head architectures on different frameworks. The setting in blue is our default instantiation.

| Parallel | DynConv ^{mask} | Fig. | AP ^{box} | AP ^{mask} | FPS |
|----------|-------------------------|------------|-------------------|--------------------|------|
| | | Fig. 2 (b) | 43.5 | 37.4 | 11.1 |
| ✓ | | Fig. 2 (b) | 43.8 | 37.9 | 11.1 |
| | ✓ | Fig. 2 (c) | 43.8 | 38.8 | 10.5 |
| ✓ | ✓ | Fig. 2 (c) | 44.5 | 39.8 | 10.5 |

Table 5: Impacts of parallel supervision and DynConv^{mask}.

| Shared MSA | Shared Query | AP ^{box} | Δ^{box} | AP ^{mask} | Δ^{mask} |
|------------|--------------|-------------------|-----------------------|--------------------|------------------------|
| | | 43.4 | | 38.1 | |
| ✓ | | 43.9 | + 0.5 | 38.3 | + 0.2 |
| | ✓ | 44.1 | + 0.7 | 39.5 | + 1.4 |
| ✓ | ✓ | 44.5 | + 1.1 | 39.8 | + 1.7 |

Table 6: Impacts of using shared query and MSA.

plied to the query based model, QueryInst (5-th row) outperforms the baseline (3-th row) by 0.7 AP^{box} and 1.9 AP^{mask}, while maintaining a high FPS. Moreover, QueryInst also beat original HTC (2-th row) in terms of both AP^{box} and AP^{mask} while runs about 3× faster. Fig. 4 demonstrates the effects of DynConv^{mask} qualitatively.

Last, we also find that for query based approaches, HTC mask flow cannot bring further improvement on top of parallel DynConv^{mask} architecture (6-th row). This indicates that the proposed parallel DynConv^{mask} enables adequate mask information flow propagating across queries in different stage for high quality mask generation. Therefore establishing explicit mask feature flow as HTC is redundant and is harmful for model efficiency. In consideration of the speed-accuracy trade-off, we choose Fig. 2 (c) as the default instantiation of our QueryInst.

5. Conclusion

In this paper, we propose an efficient query based end-to-end instance segmentation framework, QueryInst, driven by parallel supervision on dynamic mask heads. To our knowledge, QueryInst is the first query based instance segmentation method that outperforms previous state-of-the-art

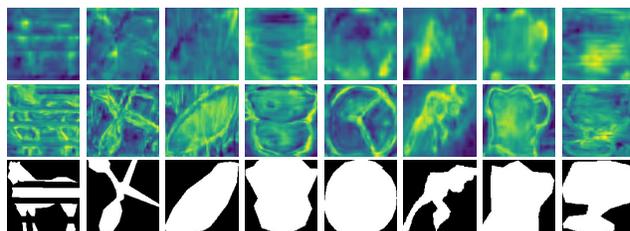


Figure 4: Effects of DynConv^{mask}. The first row shows mask features x^{mask} directly extracted from FPN. The second row shows mask features $x^{\text{mask}*}$ enhanced by queries in DynConv^{mask}. Last row is ground-truth instance masks. The results show that mask features enhanced by queries yield more genuine and accurate details and carry more information of instances.

non-query based instance segmentation approaches. Extensive study proves that parallel mask supervision can bring great performance improvement without any decrease for inference speed, while dynamic mask head with both shared query and MSA joints two sub-tasks of detection and segmentation naturally. We hope this work can strength the understanding of query based frameworks and facilitate future research.

Acknowledgement

This work was in part supported by NSFC (No. 61876212, No. 61733007 and No. 61773176) and the Zhejiang Laboratory under Grant 2019NB0AB02.

References

- [1] Ali Athar, S. Mahadevan, Aljosa Osep, L. Leal-Taixé, and B. Leibe. Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In *ECCV*, 2020.
- [2] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT++: better real-time instance segmentation. *ArXiv*, 2019.
- [3] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT: real-time instance segmentation. In *ICCV*, 2019.

- [4] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: delving into high quality object detection. In *CVPR*, 2018.
- [5] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: High quality object detection and instance segmentation. *TPAMI*, 2019.
- [6] Jiale Cao, Rao Muhammad Anwer, Hisham Cholakkal, Fahad Shahbaz Khan, Yanwei Pang, and Ling Shao. Sipmask: Spatial information preservation for fast image and video instance segmentation. In *ECCV*, 2020.
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [8] William Chan, Chitwan Saharia, Geoffrey Hinton, Mohammad Norouzi, and Navdeep Jaitly. Imputer: Sequence modelling via imputation and dynamic programming. In *ICML*, 2020.
- [9] Hao Chen, Kunyang Sun, Zhi Tian, Chunhua Shen, Yongming Huang, and Youliang Yan. BlendMask: Top-down meets bottom-up for instance segmentation. In *CVPR*, 2020.
- [10] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *CVPR*, 2019.
- [11] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.
- [12] Mingfei Chen, Yue Liao, Si Liu, Zhiyuan Chen, Fei Wang, and Chen Qian. Reformulating hoi detection as adaptive set prediction. *arXiv preprint arXiv:2103.05983*, 2021.
- [13] Tianheng Cheng, Xinggang Wang, Lichao Huang, and Wenyu Liu. Boundary-preserving mask r-cnn. In *ECCV*, 2020.
- [14] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, M. Enzweiler, Rodrigo Benenson, Uwe Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [15] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, 2017.
- [16] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. Up-detr: Unsupervised pre-training for object detection with transformers. *arXiv preprint arXiv:2011.09094*, 2020.
- [17] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [18] Yang Fu, Linjie Yang, Ding Liu, Thomas S Huang, and Humphrey Shi. Compfeat: Comprehensive feature aggregation for video instance segmentation. *arXiv preprint arXiv:2012.03400*, 2020.
- [19] Peng Gao, Minghang Zheng, Xiaogang Wang, Jifeng Dai, and Hongsheng Li. Fast convergence of detr with spatially modulated co-attention. *arXiv preprint arXiv:2101.07448*, 2021.
- [20] Ross Girshick. Fast r-cnn. In *ICCV*, 2015.
- [21] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [22] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask r-cnn. In *ICCV*, 2017.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *TPAMI*, 2015.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [25] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring r-cnn. In *CVPR*, 2019.
- [26] Xu Jia, Bert De Brabandere, Tinne Tuytelaars, and Luc Van Gool. Dynamic filter networks. In *NeurIPS*, 2016.
- [27] Lei Ke, Yu-Wing Tai, and Chi-Keung Tang. Deep occlusion-aware instance segmentation with overlapping bilayers. In *CVPR*, 2021.
- [28] Ha Young Kim and Ba Rom Kang. Instance segmentation and object detection with bounding shape masks. *CoRR abs/1810.10327*, 2018.
- [29] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [30] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [31] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, 2014.
- [32] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *CVPR*, 2018.
- [33] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. In *ECCV*, 2016.
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021.
- [35] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [36] Xin Lu, Buyu Li, Yuxin Yue, Quanquan Li, and J. Yan. Grid r-cnn. In *CVPR*, 2019.
- [37] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. *arXiv preprint arXiv:2101.02702*, 2021.

- [38] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, 2016.
- [39] Aaron Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, et al. Parallel wavenet: Fast high-fidelity speech synthesis. In *ICML*, 2018.
- [40] Jiangmiao Pang, Kai Chen, Jianping Shi, Huajun Feng, Wanli Ouyang, and Dahua Lin. Libra r-cnn: Towards balanced learning for object detection. In *CVPR*, 2019.
- [41] Shaoqing Ren, Kaiming He, Ross B. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *TPAMI*, 2015.
- [42] Peize Sun, Yi Jiang, Enze Xie, Zehuan Yuan, Changhu Wang, and Ping Luo. Onenet: Towards end-to-end one-stage object detection. *arXiv preprint arXiv:2012.05780*, 2020.
- [43] Peize Sun, Yi Jiang, Rufeng Zhang, Enze Xie, Jinkun Cao, Xinting Hu, Tao Kong, Zehuan Yuan, Changhu Wang, and Ping Luo. Transtrack: Multiple-object tracking with transformer. *arXiv preprint arXiv:2012.15460*, 2020.
- [44] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. *arXiv preprint arXiv:2011.12450*, 2020.
- [45] Zhiqing Sun, Shengcao Cao, Yiming Yang, and Kris Kitani. Rethinking transformer-based set prediction for object detection. *arXiv preprint arXiv:2011.10881*, 2020.
- [46] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *ECCV*, 2020.
- [47] Zhi Tian, Chunhua Shen, H. Chen, and Tong He. Fcos: A simple and strong anchor-free object detector. *TPAMI*, 2020.
- [48] Zhi Tian, Bowen Zhang, Hao Chen, and Chunhua Shen. Instance and panoptic segmentation using conditional convolutions. *arXiv preprint arXiv:2102.03026*, 2021.
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- [50] Thang Vu, Kang Haeyong, and Chang D Yoo. Snet: Training inference sample consistency for instance segmentation. In *AAAI*, 2021.
- [51] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. *arXiv preprint arXiv:2012.00759*, 2020.
- [52] Jianfeng Wang, Lin Song, Zeming Li, Hongbin Sun, Jian Sun, and Nanning Zheng. End-to-end object detection with fully convolutional network. *arXiv preprint arXiv:2012.03544*, 2020.
- [53] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. SOLO: Segmenting objects by locations. In *ECCV*, 2020.
- [54] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic, faster and stronger. *arXiv preprint arXiv:2003.10152*, 2020.
- [55] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. *arXiv preprint arXiv:2011.14503*, 2020.
- [56] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [57] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017.
- [58] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. Upsnet: A unified panoptic segmentation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8818–8826, 2019.
- [59] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *ICCV*, 2019.
- [60] Minghang Zheng, Peng Gao, Xiaogang Wang, Hongsheng Li, and Hao Dong. End-to-end object detection with adaptive clustering transformer. *arXiv preprint arXiv:2011.09315*, 2020.
- [61] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019.
- [62] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *CVPR*, 2019.
- [63] X. Zhu, Weijie Su, Lewei Lu, Bin Li, X. Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.