

Fashionformer: A Simple, Effective and Unified Baseline for Human Fashion Segmentation and Recognition

Shilin Xu^{1,3*} Xiangtai Li^{1,3 *} Jingbo Wang²
Guangliang Cheng³ Yunhai Tong^{1✉} Dacheng Tao⁴

¹ Key Laboratory of Machine Perception, MOE, School of Artificial Intelligence, Peking University

² CUHK-SenseTime Joint Lab, The Chinese University of Hong Kong

³ SenseTime Research

⁴ The University of Sydney

lxtpku@pku.edu.cn, xushilin@stu.pku.edu.cn, chengguangliang@sensetime.com

Abstract. Human fashion understanding is one crucial computer vision task since it has comprehensive information for real-world applications. This focus on joint human fashion segmentation and attribute recognition. Contrary to the previous works that separately model each task as a multi-head prediction problem, our insight is to bridge these two tasks with one unified model via vision transformer modeling to benefit each task. In particular, we introduce the object query for segmentation and the attribute query for attribute prediction. Both queries and their corresponding features can be linked via mask prediction. Then we adopt a two-stream query learning framework to learn the decoupled query representations. We design a novel Multi-Layer Rendering module for attribute stream to explore more fine-grained features. The decoder design shares the same spirit as DETR. Thus we name the proposed method *Fashionformer*. Extensive experiments on three human fashion datasets illustrate the effectiveness of our approach. In particular, our method with the same backbone achieve **relative 10% improvements** than previous works in case of a joint metric ($AP_{IoU+F_1}^{mask}$) for both segmentation and attribute recognition. To the best of our knowledge, we are the first unified end-to-end vision transformer framework for human fashion analysis. We hope this simple yet effective method can serve as a new flexible baseline for fashion analysis. Code will be available <https://github.com/xushilin1/FashionFormer>.

Keywords: Human Fashion, Fine-Grained Attribute Analysis, Segmentation, Vision Transformer

* The first two authors contribute equally. The work was done at Sensetime Research. This research is also supported by the National Key Research and Development Program of China under Grant No.2020YFB2103402. We thank the computation resource provided by SenseTime Research.

1 Introduction

The capability of understanding human fashion is essential for numerous real-world applications, such as digital human modeling, AR/VR, and on-line business. Limited by the representation of human fashion, previous works [39,69,58,57,61] mainly focus on fashion parts localization and ignore the fine-grained human attributes, for example, the description of clothes. To analyze such a richer context for human fashion, researchers [25] begin to categorize human fashion attributes into different aspects (*e.g.* plain, tight, normal waist) shown in Fig. 1(a). They solve this problem by extending the widely used instance segmentation framework [21] for these attributes as a multi-head prediction task, as shown in Fig. 1(b).

Although the Attribute-Mask R-CNN [25] is a strong baseline for such fashion analysis task, there are several remaining critical issues. *Firstly*, the relationship between instance segmentation and attribute recognition is not well explored and thus the fine-grained knowledge provided by these attributes for segmentation is unknown. *Secondly*, adopting single RoI (Region of Interest) representation leads to incomplete and inferior results for fine-grained attributes, causing huge gaps between segmentation results and attribute recognition. *Thirdly*, adopting box based detector limits the mask resolution and results in imprecise mask quality, which leads to uncomfortable experience for auto-segmentation when shopping. Thus, a better and unified baseline is needed for this track.

In recent years, Vision Transformers have been proven effective in feature representation learning, unification, and simplification of various downstream tasks [2,47,55,13,49]. In particular, Detection Transformer (DETR) [2] adopts *object queries* to model the detected objects. Meanwhile, several works [66,8] provide box-free segmentation solutions. Adopting query based design can naturally link segmentation and attribute recognition. More details can be obtained via corresponding masks. Moreover, as proven by ViT [13], Vision Transformers have a huge capacity, which is helpful for the shopping industry since the data is easy to obtain. These findings motivate us to formulate a new solution for fashion analysis using transformers.

In this work, we present a simple, effective and unified baseline named Fashionformer for this task. As shown in Fig. 1(c), the entire framework is an encode-decoder framework in DETR style. Our key insights are using *object queries* and *attribute queries* to encode instance-wised and attribute-wised information, respectively. Then we perform joint learning for both instance segmentation and attribute recognition. Joint learning leads to better instance segmentation results, since attribute labels constrain the scope of mask label. For example, for sleeve object, there will not exist attributes like plain or tight. In particular, we propose to match each object query into each attribute query via residual addition for implementation. This gives a good initialization of attributes and serves as the starting points for attribute recognition. Since each object query owns its corresponding mask, we name the query feature as grouped features via these corresponding masks. Then we perform updating object queries with the query features. This operation is implemented with dynamic convolution [66,47]




Fig. 1: Task introduction and method overview. (a) Illustration for human fashion segmentation and attribute recognition task. It requires the model to predict the masks, their labels (shown in red), and corresponding attributes (shown in yellow rectangles). (b) Overview of previous pipeline used in [25]. It contains several independent task heads and solves each task separately. (c) The proposed Fashionformer pipeline. It is an encoder-decoder framework that takes object queries and attribute queries as inputs and directly outputs masks and attributes. Best view in color and zoom in.

and multi-head self-attention layers [50] between query and query features iteratively. In the experiments, surprisingly, we find *joint* learning with the object queries and attribute queries boosts the original segmentation results.

Moreover, since both attribute recognition and instance segmentation have different goals, the former needs more fine-grained and discriminative representation for multi-label classification while the latter performs mask based multi-class classification. For attribute recognition, we propose a Multi-Layer Rendering (MLR) module to generate the attributes query features. The multi-level features are grouped via corresponding predicted masks. Such a module is lightweight and bridges the gap between instance segmentation and attribute recognition. Then we perform the same procedure as object query path via dynamic convolution and self-attention layers where we use attributes query feature to re-weight the attribute query. Note that both are *independent*, which leads to a two-stream decoder.

Our framework can support both CNN backbone [22] and Transformer backbone [38] as encoder. Finally, the framework is motivated by the design of box-free architecture, and it can generate high resolution masks without box supervision or RoI-wised operation. Our framework can also adapt to previous fashion benchmarks by replacing the recognition branch with an extra classification head. Extensive experiments show that our approach achieves much better results than previous works [25, 66, 9]. Fig. 2 (d) shows our method achieves better segmentation results when comparing to Attribute Mask-RCNN. As shown in Fig. 2(a)-(c), Fashionformer achieves the new state-of-the-art results on three fashion benchmarks including FashioNpedia [25], ModaNet [69] and Deepfashion [39] with *less computation cost and parameters*. Moreover, *The gap between attribute recognition and instance segmentation can be effectively narrowed through our approach by a significant performance margin to the state-of-the-art methods*. These results demonstrate that Fashionformer can serve as a new baseline for future research in fashion analysis. To the best of our knowledge, this is the first unified end-to-end model for this task.

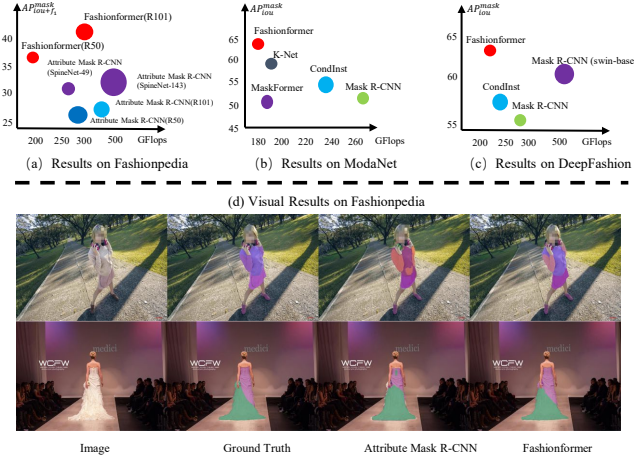


Fig. 2: Comparison Results on (a) Fashionpedia, (b) ModaNet, and (c) DeepFashion. All models in (b) and (c) use ResNet50 as backbone. The GFlop is obtained via $1020 \times 1020 \times 3$ inputs. Our method achieves a significant gain with the fewer GFlops. The number of network parameters are indicated by the radius of circles. (d). Visual comparison results using ResNet-50 backbone. Best view in color.

2 Related Work

Instance Segmentation. This task aims to detect and segment each instance [10, 20]. The two-stage pipeline Mask R-CNN like models [21, 23, 4, 25] first generate object proposals using Region Proposal Network (RPN) [44] and then predict boxes and masks on each RoI where the advanced versions utilize more cues to enhance the mask representation. Several single-stage methods [48, 3, 1, 7] achieve significant progress and comparable results with two-stage pipelines where they use single stage detector. SOLO [53] treats instances into grid representation and then performs instance classification and segmentation in a decoupled manner. Meanwhile, there are several bottom-up approaches [42, 11] where each instance is grouped from the semantic segmentation prediction [6, 71, 29, 31]. Several works [16, 12] use object query to encode instance wise information. However, they still need object detector [74, 47]. Instance Segmentation is one of the sub-task of our framework. Moreover, we show that joint learning of both segmentation and fine-grained attributes lead to better segmentation results.

Related Human Segmentation. Most works in this area focus on human part modeling, including instance part segmentation and semantic part segmentation. Several works [15, 37, 52, 73] design specific methods for semantic part segmentation which are in category-level settings. Recent human segmentation methods focus on human instance part segmentation. There are two paradigms for this direction: *top-down* pipelines [28, 60, 45, 24, 59] and *bottom-up* pipelines [19, 27, 67, 72]. Our framework solves the *instance level* segmentation with *fine-grained attributes recognition*, which is much more challenging.

Fine-grained Image Recognition. Lots of works [17,68] use the localization classification subnetwork to highlight finer feature regions and then achieves better classification results. Various approaches are proposed including attention [64], extra models [65,63] and deep filters [56,54]. Meanwhile, there are several works learning the end-to-end feature encoding via specific loss [46] or high-order feature interactions [36]. Our methods adopt object query and object mask as extra spatial cues to grasp multiscale image features, which lead to better fine-grained recognition results.

Transformer in Computer Vision. There are two research directions for vision transformer. The first is to replace CNN as a feature extractor. Compared with CNN, vision transformers [13,38,49] have more advantages in modeling long-range relation among the patch features. Moreover, they show better performance and higher capacity among the downstream tasks. The second design is to use the object query representation. DETR [2] models the object detection task as a set prediction problem with object queries. The following works [74,66,8] explore the locality of the learning process to improve the performance of DETR. Meanwhile, there are also several works using object queries to solve more complex tasks [70,62,30,55,32]. Our work is inspired by these works, which also uses a transformer architecture to unify and simplify human fashion task. Thus, our main contributions lie in the second part of vision transformer. However, our method can be adapted with vision transformer backbone [38].

3 Method

In this section, we will first review the related works using object query and present our method’s key insights and motivations. Then we briefly describe our Fashionformer and detail the design of our proposed two-stream query learning framework for both object query and attribute query. Finally, we describe the training and inference procedures.

3.1 Overview of Previous Work

Query based models. DETR [2] firstly introduces the concept of object query which is used as the input of transformer decoder to build one by one mapping on objects in the scene. Various approaches in different tasks [55,74,9] prove the effectiveness of such design including segmentation and detection. The MaskFormer [9,8] shows that the pure mask based classification can solve both semantic and instance level segmentation problems. Meanwhile, several works [66,47] show that adopting self-attention on proposal level features with their corresponding object queries can also achieve robust performance. For both aspects, the critical component is the design of object query and interaction with feature maps from backbone.

Motivation of our method. As stated in Sec. 1, recent works [25,39] on fashion analysis have several shortcomings. In this work, we seek a new baseline

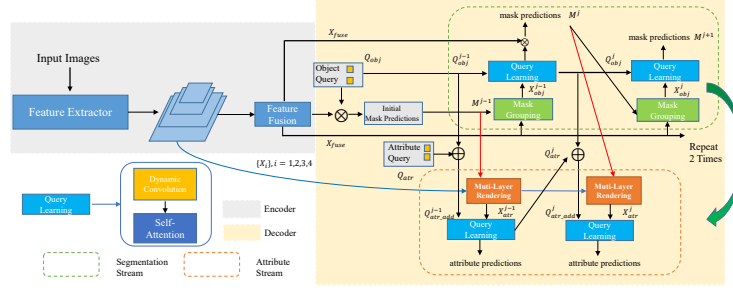


Fig. 3: Fashionformer contains three parts. (1) a feature extractor to extract multiscale features and then fuse these feature into one high resolution. (2) segmentation stream to generate segmentation masks and labels. It takes the object queries and initial mask as inputs and directly outputs refined mask and object queries interactively. (3) attribute stream to obtain the attribute predictions. It takes attribute queries, mask predictions and multiscale features as inputs and outputs the attribute prediction.

for unified fashion analysis to solve both segmentation and attribute recognition jointly. Motivated by the recent process of query based models, a natural question emerges: *can a unified model improves both instance segmentation and attribute recognition, despite the goals of these two tasks being different?* Such significant gaps between two tasks have already existed in the results of previous work [25]. Our insights are: Firstly, we add an extra attribute query with a residual like learning with object query to balance the conflicts of two tasks. Secondly, since both queries share the same instance masks, we decouple the query feature learning via designing specific modules for each task, which results in a two-stream decoder. We detail the design in the following sections.

3.2 Fashionformer as a Unified Baseline

Feature extractor. We first extract image features for each input image using a feature extractor. It contains a backbone network (Convolution Network [22] or Vision Transformer [38]) with Feature Pyramid Network [33] as neck. This results in a set of multiscale features $\{X_i, ..\}$ where $i = 1, 2, 3, 4$ are indexes of different scales. Then we sum up all the feature pyramids into one high resolution feature map X_{fuse} which follows the design of [26]. Moreover, we follow the common design of DETR-like models [2, 48, 66] to add position embeddings on the X_{fuse} . We omit this operation for simple illustration.

Meanwhile, we keep the multiscale features $\{X_i, ..\}$ for further usage. This process is shown in the gray region of Fig. 3. We denote the *object query feature* as the corresponding feature from feature extractor (shown in black line, X_{fuse}) via mask grouping, which is used in segmentation stream. We denote the *attribute query features* as corresponding feature generated by Multi-Layer Rending which is used in attribute stream. We detail the process further.

Initial object query. Following previous works [66,47], the initial weights of object queries are directly obtained from the weights of initial decoder prediction. We adopt one 1×1 convolution layer to obtain the initial instance masks M_0 . The kernel weights are the query weights. Such initialization shows better convergence. The object queries Q_{obj} are in $N \times d$ where N denotes the query number ($N = 100$ by default) and d is the hidden dimension ($d = 256$ by default, which is the same as DETR [2]).

two-stream decoder architecture. As shown in the yellow region of the Fig. 3, we adopt a two-stream design for both segmentation and attribute recognition. In particular, the attribute recognition module is appended to the segmentation module. Since most detection transformers [2,66,9] have cascaded decoder design to converge, we keep the same strategy via repeating the both modules into a cascaded decoder. In each step, the segmentation module directly outputs the binary segmentation masks and mask classification results while the recognition module outputs the attribute probability maps. We use index j to represent step number. All the results are supervised with specific loss functions during the training.

3.3 Linking Object Query and Attributes Query

Adding extra attribute query. Rather than directly learning the attribute query from the object query via a multi-layer fully connected network, we present an extra attribute query design where we initialize a learnable embedding Q_{atr} with the same shape as Q_{obj} . For each step j in the two-stream decoder, we add both queries together as the new attribute query:

$$Q_{atr_add}^{j-1} = Q_{atr}^{j-1} + Q_{obj}^{j-1}. \quad (1)$$

When $j = 1$ for the initial stage ($j \in \{1, 2, 3\}$), we denote $Q_{atr} = Q_{atr}^0$, $Q_{obj} = Q_{obj}^0$ in the following sections.

Multi-Layer Rendering (MLR). This module takes previous multi-level features $\{X_i, \dots\}$, $i = 1, 2, 3, 4$ (blue arrows in Fig. 3) and mask predictions $M^j - 1$ (red arrows in Fig. 3) from the previous stage as the inputs. It outputs a refined attribute query feature X_{atr}^j . We first resize the M^{j-1} into different scales ($i = 1, 2, 3, 4$) and then perform multi-level grouping with multi-level features X_i as following:

$$X_{i,atr}^j = \sum_u \sum_v^{W_i, H_i} M^{j-1}(u, v, i) \cdot X_i(u, v), i \in \{1, 2, 3, 4\}, \quad (2)$$

where j is the interaction number, W_i and H_i are the height and width of the corresponding features X_i , u and v are the spatial index of features. Then we obtain a set of features $X_{i,atr}^j$ with the same shape of attribute queries. Then we adopt one MLP-like to fuse these queries where it contains one fully connected layer, one Layer Norm [50] and one relu activation layer as:

$$X_{atr}^j = MLP(Concat(X_{atr}^{1,j}, \dots, X_{atr}^{i,j})), i \in \{1, 2, 3, 4\}. \quad (3)$$

This operation is shown in **orange nodes** in Fig. 3. The final refined attribute query feature X_{atr}^j will be the input of attribute stream for attribute query learning. In Sec. 4, we find this module significantly improves the attribute recognition results.

Decoupled query learning. For object query feature, we mainly follow the design of previous works [66,47]. The object query feature X_{obj}^{j-1} is obtained from the mask-based grouping from the previous mask prediction and feature map X_{fuse} (Mask Grouping in **green nodes**) where:

$$X_{obj}^{j-1} = \sum_u^W \sum_v^H M^{j-1}(u, v) \cdot X_{fuse}(u, v). \quad (4)$$

Then, following [48,66,47], we perform a Dynamic Convolution (DC) to refine input queries Q_{obj}^{j-1} with the object query features X_{obj}^j . The motivation of DC is to encode instance-wise features into each object query for fast convergence. This is **not** our contribution.

$$Q_{obj}^{j-1} = DC(X_{obj}^j, Q_{obj}^{j-1}), \quad (5)$$

where the dynamic convolution uses the query features X_{obj}^j to re-weight input queries Q_{obj}^{j-1} . To be more specific, *DynamicConv* uses object query features X_{obj}^j to generate gating parameters via MLP and multiply back to the original query input Q_{obj}^{j-1} . For attribute branch, we adopt the same procedure:

$$Q_{atr}^{j-1} = DC(X_{atr}^j, Q_{atr.add}^{j-1}), \quad (6)$$

where X_{atr}^j is the *attribute query feature* from the MLR. Note that both *DC* operations have independent parameters and learned in a separate manner. In this way, both streams are decoupled and lead to better attribute recognition results.

We adopt the same design [66] by learning gating functions to update the refined queries including both object queries and attributes queries. The *DC* operation is shown as follows:

$$\hat{Q}_u^{j-1} = gate_x(X_u^j)X_u^j + gate_q(X_u^j)Q_u^{j-1}, u \in \{obj, atr\}, \quad (7)$$

where *gate* is implemented with a fully connected (FC) layers 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 weight the query features, while the latter is to weight corresponding queries. After that, one self-attention layer with feed forward layers [50,51] is used to learn the correspondence among each query. This operation leads to the full correlation among queries shown as follows:

$$Q_u^j = FFN(MHSA(\hat{Q}_u^{j-1}) + \hat{Q}_u^{j-1}), u \in \{obj, atr\}, \quad (8)$$

where *MHSA* means Multi Head Self Attention, *FFN* is the Feed Forward Network that is commonly used in current vision transformers [2,13]. The output

refined query sets contain Q_{obj}^j and Q_{atr}^j . Both query learning works independently, shown in **blue nodes** in Fig. 3. The main reason for using independent query learning is the optimization goal. The attribute query for attribute prediction is a *multi-label classification problem*, while the object query for segmentation is a *single label classification problem*.

Generating mask and attributes prediction. Finally, the refined masks are obtained via dot product between the refined queries Q_{obj}^j and the input features X_{fuse} . For mask classification, we adopt two feed forward layers on Q_{obj}^j and directly output the class scores. For mask segmentation, we also adopt two feed forward layers on Q_{obj}^j and then we perform the inner product between learned queries and features X_{fuse} to generate stage i object masks. These masks will be used for the next step $j + 1$.

For attribute prediction, we also adopt several feed forward layers on Q_{atr}^j and use the outputs with a sigmoid activation function as final output. The process of these equations will be repeated several times. The iteration number is set to 3 by default. All the inter mask predictions are refined and optimized during the training.

3.4 Training and Inference

Loss functions. We mainly follow the design of previous works [9,51,2] to 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 object queries. In particular, we apply focal loss [34] on both classification and mask prediction. We also adopt dice loss [41] on mask predictions. For attribute queries, we follow the default design of the Attribute Mask R-CNN and the prediction is supervised in one hot format, which is the default setting of multi-label classification. It is a binary multi-label classification loss. The total loss can be written as follows:

$$\mathcal{L}_j = \lambda_{cls} \cdot \mathcal{L}_{cls} + \lambda_{mask} \cdot \mathcal{L}_{mask} + \lambda_{atr} \cdot \mathcal{L}_{atr}. \quad (9)$$

Note that the losses are applied to each stage $\mathcal{L}_{final} = \sum_j^N \mathcal{L}_j$, where N is the total stages applied to the framework. We adopt $N = 3$ and all λ s are set to 1 by default.

Inference. We directly get the output masks from the corresponding queries according to their sorted scores, which are obtained from the mask classification branch. Since object queries and attribute queries are in the one-by-one mapping manner. We obtain the attribute predictions with 0.5 threshold via a sigmoid function from attribute prediction.

4 Experiment

4.1 Settings

Dataset. We carry out our experiments on Fashionpedia [25], ModaNet [69] and DeepFashion [39] dataset. Fashionpedia is a new challenging fashion dataset

| RD | DC | MLR | N=1 | I=3 | AP_{IoU}^{mask} | $AP_{IoU+F_1}^{mask}$ |
|----|----|-----|-----|-----|-------------------|-----------------------|
| ✓ | ✓ | ✓ | - | ✓ | 33.2 | 29.2 |
| - | ✓ | ✓ | - | ✓ | 33.0 | 27.0 |
| ✓ | - | ✓ | - | ✓ | 32.1 | 27.5 |
| ✓ | ✓ | - | - | ✓ | 33.0 | 26.5 |
| ✓ | ✓ | ✓ | ✓ | - | 29.1 | 25.8 |
| - | ✓ | - | - | ✓ | 30.1 | - |

(a) Effect of each component.

| Setting | AP_{IoU}^{mask} | $AP_{IoU+F_1}^{mask}$ |
|---------|-------------------|-----------------------|
| NL=1 | 32.2 | 27.5 |
| NL=2 | 33.1 | 28.6 |
| NL=4 | 33.2 | 29.2 |

(b) Ablation on design of MLR. NL: number of layers in MLR.

| Setting | AP_{IoU}^{mask} | $AP_{IoU+F_1}^{mask}$ | Param(M) |
|------------------|-------------------|-----------------------|----------|
| Shared Query | 33.4 | 27.6 | 36.2 |
| Individual Query | 33.3 | 29.1 | 37.7 |

(c) Decoupled query learning. The Shared Query use the same object query via a MLP to generate attribute query.

Table 1: Ablation studies and analysis on Fashionpedia dataset set with ResNet50 as backbone. DC: Dynamic Convolution. RD: Residual Addition. MLR: Multi-Layer Rendering. N: Number of decoder layers.

which contains 45,632 images for training and 1,158 images for validation. ModaNet contains 52,254 images for training. Since the online server of ModaNet is closed, we randomly sample 4,000 images from the training dataset for testing. DeepFashion datasets have 6,817 images for training and 3,112 images for testing. For both ModaNet and DeepFashion datasets, we retrain the several representative baselines [21, 48, 66] for fair comparison. We use the challenging Fashionpedia dataset for ablation and analysis, since it contains both segmentation and attribute labels. We only report the instance segmentation results on the two remaining datasets.

Implementation details. We implement our models in PyTorch [43] with MMDetection toolbox [5]. We use the distributed training framework with 8 GPUs. *For Fashionpedia dataset*, we adopt similar settings as the original work. We use large scale jittering that resizes an image to a random ratio between [0.5, 2.0] of the target input image size. ResNet-50, ResNet-101 [22] and Swin-Transformer [38] are used as the backbone network and the remaining layers use Xavier [18] initialization. The optimizer is AdamW [40] with weight decay 0.0001. The training batch size is set to 16 and all models are trained with 8 GPUs. We adopt 12 epochs training (one-third of $1 \times$ schedule by default, 12 epochs in original size.) for ablation studies. We empirically found more training epochs leads to better results. All the models adopt the single scale inference. *For both ModaNet and DeepFashion dataset*, we follow the same setting from COCO [35] for fair comparison. Note that since both datasets have no attribute labels, we treat the mask classification labels as one attribute without the modification of network architecture. We present more implementation details in appendix.

Metrics. Following previous works, we adopt mask based mean average precision (AP_{IoU}^{mask}) which is the default COCO setting for instance segmentation evaluation. Moreover, we also adopt mAP that also consider F1-score ($AP_{IoU+F_1}^{mask}$) for joint evaluating segmentation and attribute recognition. This is the **main metric** for comparison. The GFlops are obtained with $3 \times 1020 \times 1020$ inputs following the original work [25]. Moreover, we also report the gap G between AP_{IoU}^{mask} and $AP_{IoU+F_1}^{mask}$ since the latter is more challenging.

| Method | backbone | AP _{IoU} ^{mask} |
|----------------|----------|-----------------------------------|
| Mask-RCNN [21] | ResNet50 | 30.3 |
| CondInst [48] | ResNet50 | 26.6 |
| QueryInst [16] | ResNet50 | 31.7 |
| Ours | ResNet50 | 33.2 |

(a) Comparison on recent box based approaches.

| Method | backbone | AP _{IoU} ^{mask} | schedule |
|----------------|----------|-----------------------------------|----------|
| K-Net [66] | ResNet50 | 30.3 | 12e |
| MaskFormer [9] | ResNet50 | 31.4 | 12e |
| Our | ResNet50 | 33.2 | 12e |
| K-Net [66] | Swin-b | 46.4 | 3× |
| Our | Swin-b | 49.5 | 3× |

(b) Comparison on recent Transformer based models. 12e means 12 epochs training.

Table 2: Comparison with recent representative methods on Fashionpedia.

4.2 Ablation Studies.

Effectiveness analysis of each component. In Tab. 1a, we explore the effectiveness of each component in our Fashionformer. Compared with instance segmentation baseline (the last row), our method (the first row) achieves significant improvements on instance segmentation (3.1%). This indicates the effectiveness of our framework that the joint learning leads to better segmentation results. It means fine-grained attribute recognition can improve the fashion classification. Removing RD leads to inferior results on attribute recognition (about 2.0% drop) which means the mask classification conflicts with attribute recognition. Removing DC also leads to bad results since the model does not converge which is also observed in previous works [47,66]. Moreover, from the first row and the fourth row, we find that adding MLR leads to a significant improvement (2.7%) on attribute recognition *without* hurting the performance of instance segmentation. MLR explores the multi-level features to make the attribute query features. This leads to a more distinctive query representation. Decreasing the number of decoder layers leads to a significant drop (from 33.2% to 29.1 %) which means more steps are needed for the decoder. However, adding more steps in decoder will not increase performance ($I = 4$, 33.3%). Details can be found in the appendix.

Ablation on design of MLR module. In Tab. 1b, we explore the layer number that affects the final attribute recognition performance. We increase the number in a top-down manner (low resolution to high resolution). As shown in that table, we find increasing the number of layers can lead to better results. This proves that better attribute recognition needs multi-level feature representation. Moreover, it only takes 3% *GFlops* increase when appending the MLR.

Ablation on design of decoupled query learning. We also explore the decoupled query learning in Equ. 5, Equ. 6 and Equ. 8. We find using independent dynamic convolution and self-attention modules can lead to better results than shared query learning where both Equ. 5 and Equ. 6 share *the same parameters*. This verifies our design of two-stream architecture and shows that both tasks need decoupled query feature learning to avoid the conflicts.

| method | backbone | schedule | GFlops | params(M) | $AP_{IoU}^{mask} \uparrow$ | $AP_{IoU+F_1}^{mask} \uparrow$ | $G \downarrow$ |
|----------------------|--------------|----------|--------|-----------|----------------------------|--------------------------------|----------------|
| Attribute-Mask R-CNN | R50-FPN | 1× | 296.7 | 46.4 | 34.3 | 25.5 | 8.8 |
| | | 2× | | | 38.1 | 28.5 | 9.6 |
| | | 3× | | | 39.2 | 29.5 | 9.7 |
| Attribute-Mask R-CNN | R101-FPN | 1× | 374.3 | 65.4 | 36.7 | 27.6 | 9.1 |
| | | 2× | | | 39.2 | 29.8 | 9.4 |
| | | 3× | | | 40.7 | 31.4 | 9.3 |
| Attribute-Mask R-CNN | SpineNet-49 | 6× | 267.2 | 40.8 | 39.6 | 31.4 | 8.2 |
| | SpineNet-96 | | 314.0 | 55.2 | 41.2 | 31.8 | 9.4 |
| | SpineNet-143 | | 498.0 | 79.2 | 43.1 | 33.3 | 9.8 |
| Fashionformer | R50-FPN | 1× | 198.0 | 37.7 | 40.3 | 36.6 | 3.7 |
| | | 3× | | | 42.5 | 39.4 | 3.1 |
| Fashionformer | R101-FPN | 1× | 275.7 | 56.6 | 43.2 | 40.5 | 2.7 |
| | | 3× | | | 45.6 | 42.8 | 2.8 |
| Attribute-Mask R-CNN | Swin-b | 3× | 508.3 | 107.3 | 47.5 | 40.6 | 6.9 |
| Fashionformer | Swin-b | 3× | 442.5 | 100.6 | 49.5 | 46.5 | 3.0 |

Table 3: Benchmark results on Fashionpedia. We report AP_{IoU}^{mask} , $AP_{IoU+F_1}^{mask}$ (main metric) and their gaps G . Swin-b means Swin-base [38]. Our methods achieve significant improvement over these three metrics.

| Method | backbone | AP_{IoU}^{mask} | GFlops |
|-----------------|----------|-------------------|--------|
| Mask R-CNN [21] | ResNet50 | 51.8 | 264.8 |
| Mask R-CNN [21] | Swin-b | 54.8 | 508.8 |
| CondInst [48] | ResNet50 | 54.9 | 234.2 |
| K-Net [66] | ResNet50 | 57.9 | 198.5 |
| QueryInst [16] | ResNet50 | 53.6 | 469.2 |
| MaskFormer [9] | ResNet50 | 51.5 | 191.5 |
| Fashionformer | ResNet50 | 62.5 | 198.0 |

(a) Results on ModaNet dataset.

| Method | backbone | AP_{IoU}^{mask} | GFlops |
|-----------------|----------|-------------------|--------|
| Mask R-CNN [21] | ResNet50 | 56.9 | 264.8 |
| Mask R-CNN [21] | Swin-b | 61.8 | 508.8 |
| CondInst [48] | ResNet50 | 58.6 | 234.2 |
| K-Net [66] | ResNet50 | 62.0 | 198.5 |
| QueryInst [16] | ResNet50 | 61.4 | 469.2 |
| Fashionformer | ResNet50 | 64.4 | 198.0 |

(b) Results on DeepFashion dataset.

Table 4: Benchmark results on ModaNet [69] (a) and DeepFashion [39] (b).

4.3 Comparison with State-of-the-art Methods

Comparison with box-based approaches on Fashionpedia. In Tab. 2a, we compare several representative results on Fashionpedia dataset using 12 epochs training. Our method achieves better result than recent instance segmentation methods including dynamic convolution based [48], query based approaches [16]. This proves our query design is a better choice for fashion analysis.

Comparison with transformer based models on Fashionpedia. We also compare several recent pure mask based approaches in Tab. 2b. All the methods use the same codebase with the same setting. Our method outperforms previous works [66,9] by a large margin under various settings (both ResNet and Swin Transformer).

Results on Fashionpedia benchmark. In Tab. 3, we compare our model results with previous state-of-the-art approaches. Models that use the Swin Transformer are re-trained in our settings. We compare our methods with Attribute-Mask R-CNN in the different settings. In particular, using ResNet50 backbone, our method outperforms it by 6% mAP_{IoU}^{mask} and nearly 10% $mAP_{IoU+F_1}^{mask}$. The gap G is decreased nearly 50%. When adopting the ResNet101 backbone, our

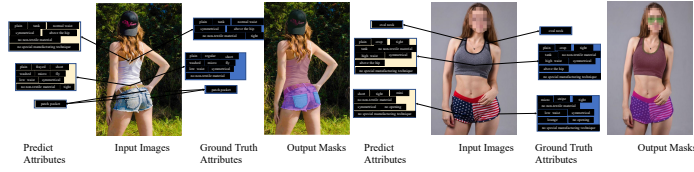


Fig. 4: Visualization results on Fashionpedia dataset. Our method obtains good segmentation results and almost right attribute predictions. Best view it on screen.




Fig. 5: Visual comparison results on ModaNet dataset. Our method has better consistent segmentation results. Best view it on screen and zoom in.

method still achieves nearly 4% $\text{mAP}_{\text{IoU}}^{\text{mask}}$ gain and 10% $\text{mAP}_{\text{IoU}+\text{F}_1}^{\text{mask}}$ gain. It even outperforms the more complex backbone SpineNet [14] with more training iterations. Finally, we equip our method with the Swin-base backbone, it also outperforms the Attribute-Mask RCNN by 5.7% $\text{mAP}_{\text{IoU}}^{\text{mask}}$ and achieves the best performance. In the above three metrics, the improvements on $\text{mAP}_{\text{IoU}+\text{F}_1}^{\text{mask}}$ are more significant which means our model can achieve much better attribute recognition. This indicates the effectiveness of our framework including two stream decoder design and the proposed Multi-Layer Rendering module.

Results on ModaNet benchmark. We further report results on the ModaNet benchmark in Tab. 4a. Our Fashionformer achieves the best performance among these works with less GFlops. Note that our work with ResNet50 backbone even outperforms Mask-RCNN with Swin-base backbone [38]. This indicates the effectiveness of our approach.

Results on DeepFashion benchmark. In Tab. 4b, we compare our methods with recent representative works [9,66]. Our works also achieve state-of-the-art results using ResNet50 backbone which further shows the generalization ability of our approach.

4.4 Analysis and Visualization

Results visualization. In Fig. 4, we present several visualization examples on Fashionpedia using ResNet101 backbone. Our model can predict perfect instance segmentation masks and almost right attributes. Moreover, we compare

| Category | AP | AP50 | AP75 | APl | APm | APs |
|------------------------|-------------|-------------|-------------|-------------|-------------|------------|
| overall(AM-RCNN) | 43.1 / 33.3 | 60.3 / 42.3 | 47.6 / 37.6 | 50.0 / 35.4 | 40.2 / 27.0 | 17.3 / 9.4 |
| outerwear | 64.1 / 40.7 | 77.4 / 49.0 | 72.9 / 46.2 | 67.1 / 43.0 | 44.4 / 29.3 | 19.0 / 4.4 |
| parts | 19.3 / 13.4 | 35.5 / 20.8 | 18.4 / 14.4 | 28.3 / 14.5 | 23.9 / 16.4 | 12.5 / 9.8 |
| accessory | 56.1 / - | 77.9 / - | 63.9 / - | 57.5 / - | 60.5 / - | 25.0 / - |
| overall(Fashionformer) | 45.7 / 42.8 | 58.2 / 49.7 | 45.7 / 43.3 | 63.3 / 48.9 | 43.9 / 31.8 | 11.4 / 6.5 |
| outerwear | 72.7 / 52.4 | 78.7 / 56.5 | 73.5 / 53.2 | 76.9 / 55.9 | 48.3 / 31.5 | 8.4 / 2.2 |
| parts | 19.2 / 20.3 | 31.8 / 30.1 | 17.2 / 19.7 | 42.6 / 29.0 | 29.8 / 26.2 | 8.2 / 7.1 |
| accessory | 56.4 / - | 75.0 / - | 58.6 / - | 73.4 / - | 60.9 / - | 18.1 / - |

Table 5: Per super-category results. Result format follows $[\text{mAP}_{\text{IoU}}^{\text{mask}} / \text{mAP}_{\text{IoU}+\text{F}_1}^{\text{mask}}]$. AM-RCNN means Attribute-Mask R-CNN with SpineNet-143 backbone. Our results are obtained with ResNet101 backbone. We follow the same COCO sub-metrics for overall and three super-categories for apparel objects.

our model on the ModaNet dataset with other methods. Our model achieves better segmentation results, including more consistent masks and correct labels. More visualization examples can be found in the appendix file.

Detailed comparison with Attribute-Mask R-CNN. In Tab. 5, we show the detailed comparison with Attribute-Mask R-CNN. Although our backbone is weak (ResNet101 vs. SpineNet-143), our method can achieve better results in most detailed metrics, including AP75, AP50, APl, and APm. However, our method leads to bad results on small objects. The main reason is that our work only considers the single scale feature representation with a simple fusion strategy. In the case of detailed categories, we find that our model achieves better results on “outerwear” categories because our method directly outputs masks on a high resolution feature. Moreover, for $\text{AP}_{\text{IoU}+\text{F}_1}$, our method achieves better results. For “part” categories, our method obtains similar results since most part objects are small. However, our method has *better* attribute recognition results for AP50 and AP75.

5 Conclusion

This paper presents a new baseline named Fashionformer for joint fashion segmentation and attribute recognition. Using both object and attribute queries, we present a unified solution in the case of task association for fashion segmentation and attribute recognition. We design a two-stream query update framework in the decoder part with a novel Muti-Layer Rendering module for better attribute recognition. Extensive experiments on the Fashionpedia dataset show the reciprocal benefits on both tasks, where their gap is minimized significantly. We achieve state-of-the-art results on three datasets (Fashionpedia, ModaNet, and Deepfashion) with fewer GFlops and parameters. Our models also show better results than the recent transformer models in various settings. We hope our method can serve as a new baseline for human fashion understanding.

References

1. Bolya, D., Zhou, C., Xiao, F., Lee, Y.J.: Yolact: Real-time instance segmentation. In: ICCV (2019)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020)
3. Chen, H., Sun, K., Tian, Z., Shen, C., Huang, Y., Yan, Y.: BlendMask: Top-down meets bottom-up for instance segmentation. In: CVPR (2020)
4. Chen, K., Pang, J., Wang, J., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: Hybrid task cascade for instance segmentation. In: CVPR (2019)
5. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 (2019)
6. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV (2018)
7. Chen, X., Girshick, R., He, K., Dollár, P.: Tensormask: A foundation for dense object segmentation. In: ICCV (2019)
8. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. arXiv (2021)
9. Cheng, B., Schwing, A.G., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. arXiv (2021)
10. Dai, J., He, K., Li, Y., Ren, S., Sun, J.: Instance-sensitive fully convolutional networks. In: ECCV. Springer (2016)
11. De Brabandere, B., Neven, D., Van Gool, L.: Semantic instance segmentation with a discriminative loss function. arXiv preprint arXiv:1708.02551 (2017)
12. Dong, B., Zeng, F., Wang, T., Zhang, X., Wei, Y.: Solq: Segmenting objects by learning queries. arXiv preprint arXiv:2106.02351 (2021)
13. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
14. Du, X., Lin, T.Y., Jin, P., Ghiasi, G., Tan, M., Cui, Y., Le, Q.V., Song, X.: Spinenet: Learning scale-permuted backbone for recognition and localization. In: CVPR. pp. 11592–11601 (2020)
15. Fang, H.S., Lu, G., Fang, X., Xie, J., Tai, Y.W., Lu, C.: Weakly and semi supervised human body part parsing via pose-guided knowledge transfer. In: CVPR (2018)
16. Fang, Y., Yang, S., Wang, X., Li, Y., Fang, C., Shan, Y., Feng, B., Liu, W.: Instances as queries. arXiv preprint arXiv:2105.01928 (2021)
17. Fu, J., Zheng, H., Mei, T.: Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In: CVPR (2017)
18. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the thirteenth international conference on artificial intelligence and statistics. pp. 249–256. JMLR Workshop and Conference Proceedings (2010)
19. Gong, K., Liang, X., Li, Y., Chen, Y., Yang, M., Lin, L.: Instance-level human parsing via part grouping network. In: ECCV (2018)
20. Hariharan, B., Arbeláez, P., Girshick, R., Malik, J.: Simultaneous detection and segmentation. In: ECCV. Springer (2014)
21. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV (2017)

22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
23. Huang, Z., Huang, L., Gong, Y., Huang, C., Wang, X.: Mask scoring r-cnn. In: CVPR (2019)
24. Ji, R., Du, D., Zhang, L., Wen, L., Wu, Y., Zhao, C., Huang, F., Lyu, S.: Learning semantic neural tree for human parsing. In: ECCV (2020)
25. Jia, M., Shi, M., Sirotenko, M., Cui, Y., Cardie, C., Hariharan, B., Adam, H., Belongie, S.: Fashionpedia: Ontology, segmentation, and an attribute localization dataset. In: ECCV. pp. 316–332. Springer (2020)
26. Kirillov, A., Girshick, R., He, K., Dollár, P.: Panoptic feature pyramid networks. In: CVPR (2019)
27. Li, J., Zhao, J., Wei, Y., Lang, C., Li, Y., Sim, T., Yan, S., Feng, J.: Multiple-human parsing in the wild. arXiv preprint arXiv:1705.07206 (2017)
28. Li, Q., Arnab, A., Torr, P.H.: Holistic, instance-level human parsing. arXiv preprint arXiv:1709.03612 (2017)
29. Li, X., Li, X., Zhang, L., Cheng, G., Shi, J., Lin, Z., Tan, S., Tong, Y.: Improving semantic segmentation via decoupled body and edge supervision. In: ECCV (2020)
30. Li, X., Xu, S., Yang, Y., Cheng, G., Tong, Y., Tao, D.: Panoptic-partformer: Learning a unified model for panoptic part segmentation. In: arxiv (2022)
31. Li, X., Zhang, L., You, A., Yang, M., Yang, K., Tong, Y.: Global aggregation then local distribution in fully convolutional networks. In: BMVC (2019)
32. Li, X., Zhang, W., Pang, J., Chen, K., Cheng, G., Tong, Y., Loy, C.C.: Video k-net: A simple, strong, and unified baseline for video segmentation. In: CVPR (2022)
33. Lin, T.Y., Dollár, P., Girshick, R.B., He, K., Hariharan, B., Belongie, S.J.: Feature pyramid networks for object detection. In: CVPR (2017)
34. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV (2017)
35. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
36. Lin, T.Y., RoyChowdhury, A., Maji, S.: Bilinear cnn models for fine-grained visual recognition. In: ICCV. pp. 1449–1457 (2015)
37. Liu, S., Sun, Y., Zhu, D., Ren, G., Chen, Y., Feng, J., Han, J.: Cross-domain human parsing via adversarial feature and label adaptation. In: AAAI (2018)
38. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. ICCV (2021)
39. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: CVPR (June 2016)
40. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization (2017)
41. Milletari, F., Navab, N., Ahmadi, S.: V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In: 3DV (2016)
42. Neven, D., Brabandere, B.D., Proesmans, M., Gool, L.V.: Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth. In: CVPR (2019)
43. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. arXiv preprint arXiv:1912.01703 (2019)
44. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NeurIPS (2015)
45. Ruan, T., Liu, T., Huang, Z., Wei, Y., Wei, S., Zhao, Y.: Devil in the details: Towards accurate single and multiple human parsing. In: AAAI (2019)

46. Sun, M., Yuan, Y., Zhou, F., Ding, E.: Multi-attention multi-class constraint for fine-grained image recognition. In: ECCV. pp. 805–821 (2018)
47. Sun, P., Zhang, R., Jiang, Y., Kong, T., Xu, C., Zhan, W., Tomizuka, M., Li, L., Yuan, Z., Wang, C., Luo, P.: SparseR-CNN: End-to-end object detection with learnable proposals. CVPR (2021)
48. Tian, Z., Shen, C., Chen, H.: Conditional convolutions for instance segmentation. arXiv preprint arXiv:2003.05664 (2020)
49. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: ICML. PMLR (2021)
50. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. arXiv preprint arXiv:1706.03762 (2017)
51. Wang, H., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: Max-deeplab: End-to-end panoptic segmentation with mask transformers. CVPR (2021)
52. Wang, W., Zhang, Z., Qi, S., Shen, J., Pang, Y., Shao, L.: Learning compositional neural information fusion for human parsing. In: ICCV (2019)
53. Wang, X., Kong, T., Shen, C., Jiang, Y., Li, L.: Solo: Segmenting objects by locations. In: ECCV (2020)
54. Wang, Y., Morariu, V.I., Davis, L.S.: Learning a discriminative filter bank within a cnn for fine-grained recognition. In: CVPR. pp. 4148–4157 (2018)
55. Wang, Y., Xu, Z., Wang, X., Shen, C., Cheng, B., Shen, H., Xia, H.: End-to-end video instance segmentation with transformers. In: CVPR (2021)
56. Xiao, T., Xu, Y., Yang, K., Zhang, J., Peng, Y., Zhang, Z.: The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In: CVPR. pp. 842–850 (2015)
57. Yamaguchi, K., Kiapour, M.H., Ortiz, L.E., Berg, T.L.: Parsing clothing in fashion photographs. In: CVPR (2012)
58. Yamaguchi, K., Kiapour, M.H., Ortiz, L.E., Berg, T.L.: Retrieving similar styles to parse clothing. PAMI **37**(5), 1028–1040 (2014)
59. Yang, L., Song, Q., Wang, Z., Hu, M., Liu, C., Xin, X., Jia, W., Xu, S.: Renovating parsing R-CNN for accurate multiple human parsing. In: ECCV (2020)
60. Yang, L., Song, Q., Wang, Z., Jiang, M.: Parsing R-CNN for instance-level human analysis. In: CVPR (2019)
61. Yang, W., Luo, P., Lin, L.: Clothing co-parsing by joint image segmentation and labeling. In: CVPR. pp. 3182–3189 (2014)
62. Yuan, H., Li, X., Yang, Y., Cheng, G., Zhang, J., Tong, Y., Zhang, L., Tao, D.: Polyphonicformer: Unified query learning for depth-aware video panoptic segmentation (2021)
63. Zhang, H., Xu, T., Elhoseiny, M., Huang, X., Zhang, S., Elgammal, A., Metaxas, D.: Spda-cnn: Unifying semantic part detection and abstraction for fine-grained recognition. In: CVPR. pp. 1143–1152 (2016)
64. Zhang, L., Huang, S., Liu, W., Tao, D.: Learning a mixture of granularity-specific experts for fine-grained categorization. In: ICCV. pp. 8331–8340 (2019)
65. Zhang, N., Donahue, J., Girshick, R., Darrell, T.: Part-based r-cnns for fine-grained category detection. In: ECCV. pp. 834–849. Springer (2014)
66. Zhang, W., Pang, J., Chen, K., Loy, C.C.: K-net: Towards unified image segmentation. NeurIPS (2021)
67. Zhao, J., Li, J., Cheng, Y., Sim, T., Yan, S., Feng, J.: Understanding humans in crowded scenes: Deep nested adversarial learning and a new benchmark for multi-human parsing. In: MM (2018)

- 68. Zheng, H., Fu, J., Mei, T., Luo, J.: Learning multi-attention convolutional neural network for fine-grained image recognition. In: ICCV. pp. 5209–5217 (2017)
- 69. Zheng, S., Yang, F., Kiapour, M.H., Piramuthu, R.: Modanet: A large-scale street fashion dataset with polygon annotations. In: ACM Multimedia (2018)
- 70. Zhou, Q., Li, X., He, L., Yang, Y., Cheng, G., Tong, Y., Ma, L., Tao, D.: Transvod: End-to-end video object detection with spatial-temporal transformers (2022)
- 71. Zhou, T., Wang, W., Konukoglu, E., Van Gool, L.: Rethinking semantic segmentation: A prototype view. In: CVPR (2022)
- 72. Zhou, T., Wang, W., Liu, S., Yang, Y., Van Gool, L.: Differentiable multi-granularity human representation learning for instance-aware human semantic parsing. In: CVPR (2021)
- 73. Zhou, T., Wang, W., Qi, S., Ling, H., Shen, J.: Cascaded human-object interaction recognition. In: CVPR. pp. 4263–4272 (2020)
- 74. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. ICLR (2020)