

# Inverted Pyramid Multi-task Transformer for Dense Scene Understanding

Hanrong Ye and Dan Xu✉

Department of Computer Science and Engineering, HKUST  
Clear Water Bay, Kowloon, Hong Kong  
{hyeae,danxu}@cse.ust.hk

**Abstract.** Multi-task dense scene understanding is a thriving research domain that requires simultaneous perception and reasoning on a series of correlated tasks with pixel-wise prediction. Most existing works encounter a severe limitation of modeling in the locality due to heavy utilization of convolution operations, while learning interactions and inference in a global spatial-position and multi-task context is critical for this problem. In this paper, we propose a novel end-to-end Inverted Pyramid multi-task Transformer (**InvPT**) to perform simultaneous modeling of spatial positions and multiple tasks in a unified framework. To the best of our knowledge, this is the first work that explores designing a transformer structure for multi-task dense prediction for scene understanding. Besides, it is widely demonstrated that a higher spatial resolution is remarkably beneficial for dense predictions, while it is very challenging for existing transformers to go deeper with higher resolutions due to huge complexity to large spatial size. InvPT presents an efficient UP-Transformer block to learn multi-task feature interaction at gradually increased resolutions, which also incorporates effective self-attention message passing and multi-scale feature aggregation to produce task-specific prediction at a high resolution. Our method achieves superior multi-task performance on NYUD-v2 and PASCAL-Context datasets respectively, and significantly outperforms previous state-of-the-arts. The code is available at <https://github.com/prismformore/InvPT>.

## 1 Introduction

Multi-task visual scene understanding typically requires joint learning and reasoning on a bunch of correlated tasks [41], which is highly important in computer vision and has a wide range of application scenarios such as autonomous driving, robotics, and augmented or virtual reality (AR/VR). Many of visual scene understanding tasks produce pixel-wise predictions for dense understanding of the scene such as semantic segmentation, monocular depth estimation, and human parsing. These dense prediction tasks essentially have abundant explicit and implicit correlation at the pixel level [48], which is very beneficial and can be fully utilized to improve the overall performance of multi-task models. However, how to effectively learn and exploit the cross-task correlation (*e.g.* complementarity and consistency) in a single model remains a challenging open issue.




Fig. 1: Joint learning and inference of global spatial interaction and simultaneous all-task interaction are critically important for multi-task dense prediction.

To advance the multi-task dense scene understanding, existing works mostly rely on the powerful Convolutional Neural Networks (CNN), and significant effort has been made by developing multi-task optimization losses [19] and designing [48] or searching [16] multi-task information sharing strategies and network structures. Although promising performance has been achieved, these works are still limited by the nature of convolution kernels that are heavily used in their deep learning frameworks, which model critical spatial and task related contexts in relatively local perceptive fields (*i.e.* locality discussed by previous works [45,3]). Although the recently proposed attention based methods address this issue [52,55,5], the scope of their cross-task interaction is still highly limited. However, for multi-task dense scene understanding, the capability of capturing long-range dependency and *simultaneously* modeling global relationships of *all* tasks is crucially important for this pixel-wise multi-task problem (see Fig. 1).

On the other hand, recently the transformer models have been introduced to model the long-range spatial relationship for dense prediction problems [34,50] but they only target the setting of *single-task* learning, while the joint modeling of multiple dense prediction tasks with transformer is rarely explored in the literature, and it is not a trivial problem to globally model both the spatial and the cross-task correlations in a unified transformer framework. Besides, the performance of dense prediction tasks is greatly affected by the resolution of the final feature maps produced from the model, while it is very difficult for existing transformers to go deeper with higher resolution because of huge complexity brought by large spatial size, and typically many transformers downsample the spatial resolution dramatically to reduce the computation overhead [14,44].

To tackle the above-mentioned issues, in this work we propose a novel end-to-end Inverted Pyramid Multi-task Transformer (InvPT) framework, which can jointly model the long-range dependency within spatial and all-task contexts, and also efficiently learn fine-grained dense prediction maps at a higher resolution, for multi-task dense scene understanding. Specifically, it consists of three core designs: (i) an InvPT transformer encoder to learn generic image representations for input images, (ii) an InvPT Transformer decoder built by consecutively stacking a proposed efficient UP-Transformer block, to model implicit correlations among all the dense prediction tasks and produce multi-task

features with gradually increased resolutions, and (iii) two effective multi-scale strategies, *i.e.* cross-scale Attention Message Passing (AMP) that aggregates self-attention maps across different transformer blocks, and multi-scale Encoder Feature Aggregation (EFA) that enhances the decoding features with multi-scale information from the InvPT transformer encoder.

The proposed method yields strong multi-task performance measured by the relative improvement metric [25], which is 2.59% and 1.76% on NYUD-v2 and PASCAL-Context datasets respectively. The proposed method also *largely* outperforms other state-of-the-art methods on all 9 evaluation metrics of these two benchmarks. Notably on NYUD-v2, it surpasses the best competitor by **7.23** points (mIoU) on semantic segmentation, while on PASCAL-Context it outperforms the previous best result by **11.36** and **4.68** points (mIoU) on semantic segmentation and human parsing, respectively.

In summary, our main contribution is three-fold:

- We propose a novel end-to-end Inverted Pyramid Multi-task Transformer (InvPT) framework for jointly learning multiple dense prediction tasks, which can effectively model long-range interaction in both spatial and all-task contexts in a unified architecture. As far as we know, it is the *first* work to present a transformer structure for this problem.
- We design an efficient UP-Transformer block, which allows for multi-task feature interaction and refinement at gradually *increased* resolutions, and can construct a multi-layer InvPT decoder by consecutively stacking multiple blocks to produce final feature maps with a high resolution to largely boost dense predictions. The UP-Transformer block can also flexibly embed multi-scale information through two designed strategies, *i.e.* cross-scale self-attention message passing and InvPT encoder feature aggregation.
- The proposed framework obtains superior performance on multi-task dense prediction and *remarkably* outperforms the previous state-of-the-arts on two challenging benchmarks (*i.e.* Pascal-Context and NYUD-v2).

## 2 Related Work

We review the most related works in the literature from two angles, *i.e.* multi-task deep learning for scene understanding and visual transformers.

**Multi-task Deep Learning for Scene Understanding** As an active research field [20,25,48], multi-task deep learning can greatly help to improve the efficiency of training, as it only needs to optimize once for multiple tasks, and the overall performance of scene understanding when compared with performing several scene understanding tasks separately [42]. Multi-task deep learning methods mainly focus on two directions [41], *i.e.* multi-task optimization and network architecture design. Previous works in the former direction typically investigate loss balancing techniques in the optimization process to address the problem of task competition [10,11,19]. In the latter direction researchers design explicit or implicit mechanisms for modeling cross-task interaction and embed them into the whole deep model [27,35,17] or only the decoder stage [52,55,51].

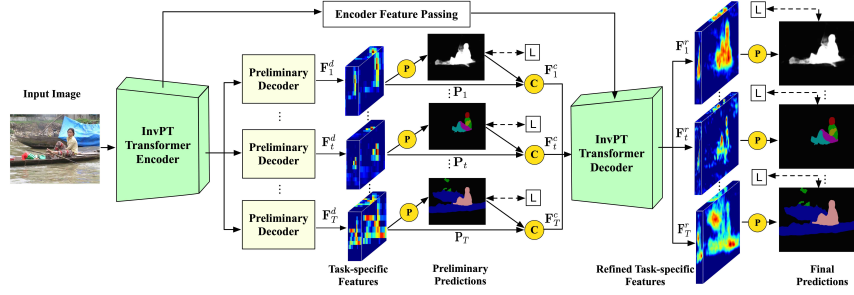


Fig. 2: Framework overview of the proposed Inverted Pyramid Multi-task Transformer (InvPT) for dense scene understanding. The task-shared transformer encoder learns generic visual representations from the input image. Then for each task  $t \in \{1 \dots T\}$ , the preliminary decoder produces task-specific feature  $\mathbf{F}_t^d$  and preliminary prediction  $\mathbf{P}_t$ , which are combined as  $\mathbf{F}_t^c$ , serving as the input of the InvPT decoder to generate *refined and resolution-enlarged* task-specific features via globally modeling spatial and all-task interactions for the final prediction. ©, Ⓟ and L denote the channel-wise concatenation, linear projection layer and loss function, respectively.

Regarding the multi-task dense scene understanding where all the tasks require pixel-wise predictions, many pioneering research works [19,25,27,17,48] have explored this field. Specifically, Xu *et al.* [48] propose PAD-Net with an effective information distillation module with attention guided cross-task message passing on multi-task predictions. MTI-Net [42] designs a sophisticated multi-scale and multi-task CNN architecture to distill information at multiple feature scales. However, most of these methods adopt CNN to learn multi-task representation and model it in a limited local context. To address this issue, there are also some exciting works developing attention-based mechanisms. For instance, [52], [55] and [51] design spatial global or local attention learning strategies within each task-specific branch, and propagate the attention to refine features among all the tasks. Concurrent to our work, [5] proposes to build up decoder via searching for proper cross-task attention structures with neural architecture search (NAS). Despite their innovative designs, these works still fail to jointly model spatial and cross-task contexts in a global manner.

Different from these works, we propose a novel multi-task transformer framework to learn long-range interactions of spatial and cross-task relationships in global contexts, which is critical for this complex multi-task problem.

**Visual Transformer** Research interest in visual transformers grows rapidly nowadays inspired by the recent success of transformers [43] in multi-modal learning [31,33], 2D [29,14,6,7,18,8] and 3D computer vision tasks [28,26,53]. The transformer models are originally designed for natural language processing tasks, and then show strong performance and generalization ability in solving vision problems [14,4,2]. Exciting results have been achieved from different aspects including: (i) Enhancing the self-attention mechanism to incorporate useful

inductive bias [12,32,22,49]. *e.g.* Swin Transformer [22] replaces the global attention with shifted window attention to improve efficiency; Focal Transformer [49] combines coarse-granularity global attention with fine-grained local attention to balance model efficiency and effectiveness; (ii) Combining transformer with CNNs [38,44,46]. *e.g.* BOTNet [38] uses a specially designed multi-head self-attention head to replace final three bottleneck blocks of ResNet; PVT [44] and CVT [46] embed convolutional layers in a hierarchical transformer framework and demonstrate that they can help improve the performance; (iii) Designing special training techniques [39,21,40]. *e.g.* DEIT [39] proposes a special transformer token for knowledge distillation from CNNs; DRLOC [21] designs a self-supervised auxiliary task to make transformer learn spatial relations. Regarding dense scene understanding tasks [44,47], Ranftl *et al.* [34] recently propose a visual transformer framework with transpose convolution for dealing with dense prediction, while HRFormer [50] adopts local attention to keep multi-scale features efficiently in the network.

To the best of our knowledge, this is the first exploration of simultaneously modeling multiple dense prediction tasks in a unified transformer framework for scene understanding. The proposed framework jointly learns spatial and all-task interactions in a global context at gradually *increased* resolutions with a novel and efficient UP-Transformer block, which can produce spatially higher-resolution feature maps to significantly boost dense predictions, and effectively incorporate multi-scale information via the proposed strategies of self-attention message passing and encoder feature aggregation.

### 3 InvPT for Multi-Task Dense Prediction

#### 3.1 Framework Overview

The overall framework of the proposed Inverted Pyramid Multi-task Transformer (InvPT) is depicted in Fig. 2. It mainly consists of three core parts, *i.e.* a task-shared InvPT transformer encoder, the task-specific preliminary decoders, and the InvPT transformer decoder. Specifically, the transformer encoder learns *generic* visual representations from the input images for all tasks. Then, the preliminary decoders produce *task-specific* features and preliminary predictions, which are supervised by the ground-truth labels. The task-specific feature and preliminary prediction of each task are combined and concatenated as a sequence serving as input of the InvPT Transformer decoder, to learn to produce refined task-specific representations within global spatial and task contexts, which are further used to produce the final predictions with task-specific linear projection layers. The details of these parts are introduced as follows.

#### 3.2 InvPT Transformer Encoder

The transformer encoder is shared for different tasks to learn generic visual representations from the input image. The self-attention mechanism of the transformer can help to learn a global feature representation of the input via long-range modeling of the spatial dependency of image pixels or patches. In our

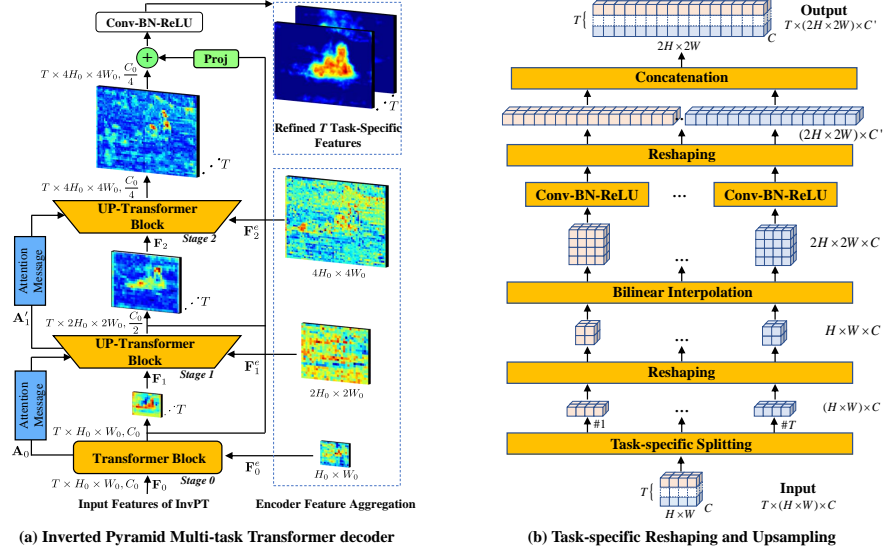


Fig. 3: **(a)** Illustration of InvPT Transformer decoder. Stage 0 uses a transformer block which keeps the feature resolution unchanged, while stage 1 and 2 employ the proposed UP-Transformer block which enlarges the spatial resolution by  $2\times$ . The Attention Message Passing (AMP) enables cross-scale self-attention interaction, and the Encoder Feature aggregation incorporates multi-scale information from the InvPT encoder. “ $\oplus$ ” and “Proj” denote the accumulation operation and the linear projection layer, respectively. **(b)** Pipeline of task-specific reshaping and upsampling block.

implementation, we consider different alternatives for the encoder including ViT [14] and Swin Transformer [22]. We obtain from the encoder a feature sequence and reshape it as a spatial feature map with resolution  $H_0 \times W_0$  where  $H_0$  and  $W_0$  denote its height and width respectively. The feature map is then input into  $T$  preliminary decoders to learn the  $T$  task-specific feature maps.

### 3.3 Task-Specific Preliminary Decoders

To learn task-specific representations for different tasks, we construct a decoding block consisting of a  $3 \times 3$  convolutional layer, a batch normalization layer, and a ReLU activation function (*i.e.* “Conv-BN-ReLU”). The preliminary decoder uses two such blocks to produce task-specific feature for each task. Then suppose that we have  $T$  tasks in total, for the  $t$ -th task, the output task-specific feature  $\mathbf{F}_t^d$  from the preliminary decoder is projected by a linear projection layer (*i.e.*  $1 \times 1$  convolution) to produce a preliminary task prediction  $\mathbf{P}_t$ , which is supervised by the ground-truth labels. Then, we concatenate  $\mathbf{F}_t^d$  and  $\mathbf{P}_t$  along the channel dimension, and adjust the channel number to  $C_0$  with a linear projection layer, to make the different task-specific features with the same channel

dimension to facilitate the processing in the transformer decoder. The combined output is denoted as  $\mathbf{F}_t^c$ . We flatten  $\mathbf{F}_t^c$  spatially to a sequence, and concatenate all  $T$  sequences as  $\mathbf{F}^c$  with  $\mathbf{F}^c \in \mathbb{R}^{TH_0W_0 \times C_0}$ .  $\mathbf{F}^c$  serves as the input of the proposed InvPT transformer decoder, which enables global task interaction and progressively generates refined task-specific feature maps with a high-resolution.

### 3.4 InvPT Transformer Decoder via UP-Transformer Block

As global self-attention is prohibitively expensive when the spatial resolution is high, many visual transformer models typically downsample the feature maps dramatically [14, 46, 44] and output features with low spatial resolution. However, resolution of the feature map is a critical factor for dense prediction problems, as we need to predict task labels for each pixel, and the sizes of semantic objects in images vary tremendously [54]. Another point is that different scales of the feature maps can model different levels of visual information [42, 50], and thus it is particularly beneficial to make different tasks learn from each other at multiple scales. With these motivations, we design a progressively resolution-enlarged transformer decoder termed as ‘‘Inverted Pyramid Transformer Decoder’’ (*i.e.* InvPT decoder), which consists of an efficient UP-Transformer block, cross-scale self-attention message passing, and multi-scale encoder feature aggregation, in a unified network module.

**Main Structure** As shown in Fig. 3 (a), there are three stages in the InvPT decoder, and each stage is the designed UP-Transformer block computing self-attention and updating features at *different* spatial resolutions. The first stage (*i.e.* stage 0) of InvPT decoder learns cross-task self-attention at the output resolution of the InvPT encoder (*i.e.*  $H_0 \times W_0$ ), while the following two stages enlarge the spatial resolution of the feature maps, and calculate cross-task self-attention at higher resolutions. The latter two stages (*i.e.* stage 1 and 2) use the proposed UP-Transformer blocks which refine features at higher resolution and enable cross-scale self-attention propagation as well as multi-scale transformer feature aggregation from the transformer encoder.

To simplify the description, for the  $s$ -th stage (*e.g.*  $s = 0, 1, 2$ ), we denote its input feature as  $\mathbf{F}_s$  with  $\mathbf{F}_s \in \mathbb{R}^{TH_sW_s \times C_s}$ , where  $H_s$  and  $W_s$  are the spatial height and width of the feature, and  $C_s$  is the channel number. Thus, the input of stage 0 (*i.e.*  $\mathbf{F}_0$ ) is  $\mathbf{F}^c$ , which is a combination of a set of  $T$  task-specific features as introduced in Sec. 3.3. For the output of InvPT decoder, we add up the feature maps from different stages after the alignment of feature resolution and channel dimension with bilinear interpolation and a linear projection layer, and then pass it to a Conv-BN-ReLU block to produce  $T$  upsampled and refined task-specific features for the final prediction.

**Task-Specific Reshaping and Upsampling** The transformer computing block typically operates on 2D feature-token sequences with the spatial structure of the feature map broken down, while the spatial structure is critical for dense prediction tasks, and it is not straightforward to directly perform upsampling on the feature sequence with the consideration of spatial structure. Another issue is



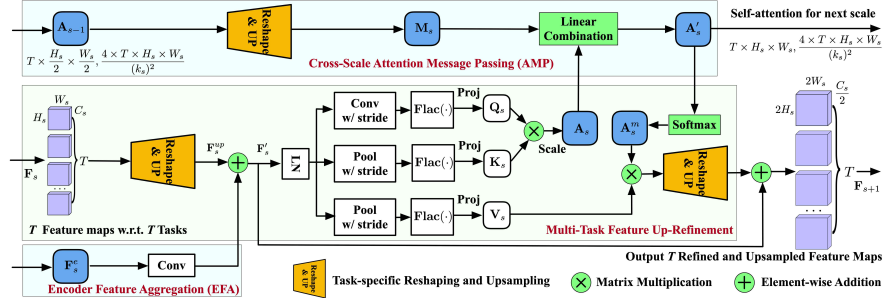


Fig. 4: Illustration of the proposed UP-Transformer Block. **Input:**  $\mathbf{A}_{s-1}$  is the attention score matrix of the  $(s-1)$ -th stage,  $\mathbf{F}_s$  is the input multi-task feature sequence of the  $s$ -th stage, and  $\mathbf{F}_s^e$  is the feature passed from the transformer encoder. **Output:**  $\mathbf{A}_s$  is the enhanced attention score matrix passing to the next stage and  $\mathbf{F}_{s+1}$  is the refined and upsampled feature sequence.

that the input features into the InvPT decoder consist of multiple different task-specific features, we need to perform feature upsampling and refinement for each single task separately to avoid feature corruption by other tasks. To address these issues, we design a task-specific reshaping and upsampling computation block for the InvPT decoder, *i.e.* the Reshape & UP module as illustrated in Fig. 3 (b) and Fig. 4. Reshape & Up first splits the feature tokens of  $\mathbf{F}_s \in \mathbb{R}^{TH_s W_s \times C_s}$  in the first dimension into  $T$  (*i.e.* the number of tasks) groups of features with tensor slicing, and reshapes each of them back to a spatial feature map with a shape of  $\mathbb{R}^{H_s \times W_s \times C_s}$ . Then a bilinear interpolation with a scaling factor of 2 is performed to enlarge the height and width by  $2\times$ . Each scaled task-specific feature map is further fed into a Conv-BN-Relu block to perform feature fusion and reduction of channel dimension if necessary. Finally, The  $T$  feature maps are reshaped back to token sequences and finally concatenated as an upsampled multi-task token sequence.

**Multi-task UP-Transformer Block** The multi-task UP-Transformer block (see Fig. 4) is employed in both stage 1 and 2 (*i.e.*  $s = 1, 2$ ) and learns to gradually increase the spatial resolution of the multi-task feature and also perform feature interaction and refinement among *all* the tasks in a global manner. As illustrated in Fig. 4, after the Reshape & Up module with a  $3 \times 3$  Conv-BN-ReLU computation, we reduce the feature channel dimension by  $2\times$  and produce a learned upsampled feature-token sequence  $\mathbf{F}_s^{up} \in \mathbb{R}^{4TH_s W_s \times (C_s/2)}$ .  $\mathbf{F}_s^{up}$  is first added by the feature sequence passed from the transformer encoder, which we introduce later in this section, and a layer normalization (LN) [1] is performed on the combined features to produce  $\mathbf{F}'_s \in \mathbb{R}^{4TH_s W_s \times (C_s/2)}$ , which serves as the input for the self-attention calculation.

To calculate a global self-attention from the upsampled high-resolution multi-task feature sequence  $\mathbf{F}'_s$ , the memory footprint is prohibitively large. Thus, we first reduce the size of query  $\mathbf{Q}_s$ , key  $\mathbf{K}_s$ , and value  $\mathbf{V}_s$  matrices for self-attention



computation following works [44,46]. Specifically, we first split and reshape the feature-token sequence  $\mathbf{F}'_s$  into  $T$  groups of spatial feature maps corresponding to the  $T$  tasks, each with a shape of  $\mathbb{R}^{2H_s \times 2W_s \times (C_s/2)}$ . To generate the query embedding from each feature map, we use a convolution with a fixed kernel size  $k_c = 3$  and a stride 2 for each group, denoted as  $\text{Conv}(\cdot, k_c)$ . To generate the key and value embeddings, we use an average pooling operation  $\text{Pool}(\cdot, k_s)$  with a kernel size  $k_s$  ( $k_s = 2^{s+1}$  for the  $s$ -th stage). By controlling the kernel parameters of the convolution and pooling operations, we can largely improve the memory and computation efficiency of the global self-attention calculation, which makes the utilization of multiple consecutive UP-Transformer blocks possible. Then, we define  $\mathbf{Flac}(\cdot)$  as a function that first flattens the  $T$  groups of spatial feature maps and then performs concatenation to produce a multi-task feature-token sequence. Then we can perform global interaction among all the tasks with a multi-task self-attention. Let  $\mathbf{W}_s^q$ ,  $\mathbf{W}_s^k$ , and  $\mathbf{W}_s^v$  be the parameter matrices of three linear projection layers, the calculation of  $\mathbf{Q}_s$ ,  $\mathbf{K}_s$ , and  $\mathbf{V}_s$  can be formulated as follows:

$$\begin{aligned}\mathbf{Q}_s &= \mathbf{W}_s^q \times \mathbf{Flac}(\text{Conv}(\mathbf{F}'_s, k_c)), \mathbf{Q}_s \in \mathbb{R}^{TH_s W_s \times \frac{C_s}{2}}, \\ \mathbf{K}_s &= \mathbf{W}_s^k \times \mathbf{Flac}(\text{Pool}(\mathbf{F}'_s, k_s)), \mathbf{K}_s \in \mathbb{R}^{\frac{4TH_s W_s}{(k_s)^2} \times \frac{C_s}{2}}, \\ \mathbf{V}_s &= \mathbf{W}_s^v \times \mathbf{Flac}(\text{Pool}(\mathbf{F}'_s, k_s)), \mathbf{V}_s \in \mathbb{R}^{\frac{4TH_s W_s}{(k_s)^2} \times \frac{C_s}{2}}.\end{aligned}\quad (1)$$

With  $\mathbf{Q}_s$  and  $\mathbf{K}_s$ , the self-attention score matrix  $\mathbf{A}_s$  for the  $s$ -th stage can be calculated as:

$$\mathbf{A}_s = \frac{\mathbf{Q}_s \mathbf{K}_s^T}{\sqrt{C'_s}}, \mathbf{A}_s \in \mathbb{R}^{TH_s W_s \times \frac{4TH_s W_s}{(k_s)^2}}, \quad (2)$$

where  $\sqrt{C'_s}$  with  $C'_s = \frac{C_s}{2}$  is a scaling factor to address the magnitude problem [43]. In vanilla transformers, the self-attention map is directly calculated with a softmax function on  $\mathbf{A}_s$ . We propose a mechanism of Attention Message Passing (AMP) to enhance the attention  $\mathbf{A}_s$  before the Softmax normalization, using multi-scale information from different transformer stages.

**Cross-Scale Self-Attention Message Passing** To enable the InvPT decoder to model the cross-task interaction at different scales more effectively, we pass attention message to the attention score matrix  $\mathbf{A}_s$  at the current  $s$ -th stage from  $\mathbf{A}_{s-1} \in \mathbb{R}^{\frac{TH_s W_s}{4} \times \frac{4TH_s W_s}{(k_s)^2}}$ , which is computed at the  $(s-1)$ -th stage. It can be noted that the second dimension of  $\mathbf{A}_s$  at different stages maintain the same size due to the design of the kernel size  $k_s$ . An illustration of AMP is shown in Fig. 4. Specifically, we perform ‘Reshape & Up’ operation to first adjust the shape of  $\mathbf{A}_{s-1}$  to  $\mathbb{R}^{\frac{H_s}{2} \times \frac{W_s}{2} \times \frac{TH_s W_s}{(k_s)^2}}$ , and then perform  $2 \times$  bilinear interpolation in the first two spatial dimensions, and finally flatten it to have the attention message matrix  $\mathbf{M}_{s-1} \in \mathbb{R}^{TH_s W_s \times \frac{4TH_s W_s}{(k_s)^2}}$  which has the same dimension as  $\mathbf{A}_s$ . Finally, we perform a linear combination of the self-attention maps  $\mathbf{A}_s$  and  $\mathbf{M}_{s-1}$  to pass attention message from the  $(s-1)$ -th stage to the  $s$ -th stage as:

$$\begin{aligned}\mathbf{A}'_s &= \alpha_s^1 \mathbf{A}_s + \alpha_s^2 \mathbf{M}_{s-1}, \\ \mathbf{A}_s^m &= \text{Softmax}(\mathbf{A}'_s),\end{aligned}\quad (3)$$

where  $\alpha_s^1$  and  $\alpha_s^2$  are learnable weights for  $\mathbf{A}_s$  and  $\mathbf{M}_{s-1}$ , respectively;  $\text{Softmax}(\cdot)$  is a *row-wise* softmax function. After multiplying with the new attention map  $\mathbf{A}_s^m$  with the value matrix  $\mathbf{V}_s$ , we obtain the final multi-task feature  $\mathbf{F}_{s+1}$  as output, which is refined and upsampled, and can be further fed into the next UP-Transformer block. The whole process can be formulated as follows:

$$\mathbf{F}_{s+1} = \text{Reshape\_Up}(\mathbf{A}_s^m \times \mathbf{V}_s) + \mathbf{F}'_s, \quad (4)$$

where the function  $\text{Reshape\_Up}(\cdot)$  denotes the task-specific reshaping and up-sampling. It enlarges the feature map by  $2\times$  to be the same spatial resolution as  $\mathbf{F}'_s$  which is a residual feature map from the input.

**Efficient Multi-Scale Encoder Feature Aggregation** For dense scene understanding, some basic tasks such as object boundary detection require lower-level visual representation. However, it is tricky to efficiently use the multi-scale features in transformer as standard transformer structures have a quadratic computational complexity regarding the image resolution, and typically only operate on small-resolution feature maps. To gradually increase the feature map size and also incorporate multi-scale features is very challenging to GPU memory. Therefore, we design an efficient yet effective multi-scale encoder feature aggregation (EFA) strategy. As shown in Fig. 4, in each upsampling stage  $s$  of the InvPT decoder, we obtain a corresponding-scale feature sequence  $\mathbf{F}_s^e$  with channel number  $C_s^e$  from the transformer encoder. We reshape it to a spatial feature map and apply a  $3 \times 3$  convolution to produce a new feature map with channel dimension  $C_s$ , and then reshape and expand it  $T$  times to align with the dimension of  $\mathbf{F}_s^{up}$  and add to it. Benefiting from the proposed efficient UP-Transformer block, we can still maintain high efficiency even if at each consecutive stage of the decoder, the resolution is gradually enlarged by two times.

## 4 Experiments

We present extensive experiments to demonstrate the effectiveness of the proposed InvPT framework for multi-task dense prediction.

**Datasets** The experiments are conducted on two popular scene understanding datasets with multi-task labels, *i.e.* NYUD-v2 [37] and PASCAL-Context [9]. **NYUD-v2** contains various indoor scenes such as offices and living rooms with 795 training and 654 testing images. It provides different dense labels, including semantic segmentation, monocular depth estimation, surface normal estimation and object boundary detection. **PASCAL-Context** is formed from PASCAL dataset [15]. It has 4,998 images in the training split and 5,105 in the testing split, covering both indoor and outdoor scenes. This dataset provides pixel-wise labels for semantic segmentation, human parsing and object boundary detection. Additionally, [25] generates surface normal and saliency labels for this dataset. We perform experiments on *all* tasks in both datasets for evaluation.

**Evaluation** Semantic segmentation (Semseg) and human parsing (Parsing) are evaluated with mean Intersection over Union (mIoU); monocular depth estimation (Depth) is evaluated with Root Mean Square Error (RMSE); surface normal

Table 1: Ablation study on the InvPT decoder. The proposed InvPT and its components yield consistent improvement on different datasets and achieve clear overall improvement on each single task and the multi-task (MT) performance  $\Delta_m$ . The gain shows absolute performance point improvement. The different variants of the InvPT all use Swin-tiny as its encoder structure. ‘ $\downarrow$ ’ means lower better and ‘ $\uparrow$ ’ means higher better.

Model	NYUD-v2				PASCAL-Context				
	Semseg mIoU $\uparrow$	Depth RMSE $\downarrow$	Normal mErr $\downarrow$	Boundary odsF $\uparrow$	Semseg mIoU $\uparrow$	Parsing mIoU $\uparrow$	Saliency maxF $\uparrow$	Normal mErr $\downarrow$	Boundary odsF $\uparrow$
InvPT Baseline (ST)	43.29	0.5975	20.80	76.10	72.43	61.13	83.43	14.38	71.50
InvPT Baseline (MT)	41.06	0.6350	21.47	76.00	70.92	59.63	82.63	14.63	71.30
InvPT w/ UTB	43.18	0.5643	21.05	76.10	72.34	61.08	83.99	14.49	71.60
InvPT w/ UTB+AMP	43.64	0.5617	20.87	76.10	73.29	61.78	84.03	14.37	71.80
InvPT w/ UTB+AMP+EFA (Full)	44.27	0.5589	20.46	76.10	73.93	62.73	84.24	14.15	72.60
Gain on Each Task (vs. MT)	$\Delta 3.21$	$\Delta 0.0761$	$\Delta 1.01$	$\Delta 0.10$	$\Delta 3.01$	$\Delta 3.10$	$\Delta 1.61$	$\Delta 0.48$	$\Delta 1.30$
MT Performance $\Delta_m$ (vs. ST) [25]	+2.59				+1.76				




Fig. 5: (a): Qualitative comparison with the previous best method ATRC [5] on NYUD-v2. (b)&(c): Qualitative analysis of InvPT decoder on PASCAL-Context. Results of different model variants are shown by columns.

estimation (Normal) is evaluated by the mean error (mErr) of predicted angles; saliency detection (Saliency) is evaluated with maximal F-measure (maxF); object boundary detection (Boundary) is evaluated with the optimal-dataset-scale F-measure (odsF). To evaluate the average performance gain of multi-task models against single-task models, we adopt the “multi-task learning performance” (MT Performance) metric  $\Delta_m$  introduced in [25], which is an important metric to reflect the performance of a multi-task model (the higher the better).

**Implementation Details** For the ablation study, we adopt Swin-Tiny transformer [22] pre-trained on ImageNet-22K [13] as the transformer encoder. The models for different evaluation experiments are trained for 40,000 iterations on both datasets, with a batch size of 6. For the transformer models, Adam optimizer is adopted with a learning rate of  $2 \times 10^{-5}$ , and a weight decay rate of  $1 \times 10^{-6}$ . Polynomial learning rate scheduler is used. The output channel number of preliminary decoder is 768. More details are in supplemental materials.

#### 4.1 Model Analysis

**Baselines and Model Variants** To have a deep analysis of the proposed InvPT framework, we first define several model baselines and variants (see Table 1): (i) “InvPT Baseline (MT)” denotes a strong multi-task baseline model of the proposed InvPT framework. It uses Swin-tiny encoder and two  $3 \times 3$  Conv-BN-




Fig. 6: Investigation of the number of stages in the InvPT decoder.

Encoder	Semseg mIoU $\uparrow$	Parsing mIoU $\uparrow$	Saliency maxF $\uparrow$	Normal mErr $\downarrow$	Boundary odsF $\uparrow$
Swin-T	73.93	62.73	<b>84.24</b>	<b>14.15</b>	72.60
Swin-B	77.50	66.83	83.65	14.63	73.00
Swin-L	<b>78.53</b>	<b>68.58</b>	83.71	14.56	<b>73.60</b>
Vit-B	77.33	66.62	<b>85.14</b>	<b>13.78</b>	<b>73.20</b>
Vit-L	<b>79.03</b>	<b>67.61</b>	84.81	14.15	73.00

Table 2: Performance comparison of using different transformer encoder structures in InvPT on PASCAL-Context.

ReLU blocks as decoder for each task, which is equivalent to the preliminary decoder in InvPT. The encoder feature map is upsampled by  $8\times$  before the final prediction. It also combines multi-scale features from the encoder to help boost performance. This is a typical multi-task baseline structure as in previous works [42, 5]. (ii) “InvPT Baseline (ST)” has a similar structure as “InvPT Baseline” but it is trained under single-task setting. (iii) “InvPT w/ UTB” indicates adding the proposed UP-Transformer block upon “InvPT Baseline (MT)”; Similarly, “InvPT w/ UTB + AMP” indicates further adding the cross-scale Attention Message Passing, and “InvPT w/ UTB + AMP + EFA” denotes the full model by further adding multi-scale Encoder Feature Aggregation.

**Effectiveness of InvPT Decoder** In this part, we investigate the effectiveness of the proposed three modules to demonstrate the proposed InvPT decoder, *i.e.* UP-Transformer Block (UTB), cross-scale Attention Message Passing (AMP), and Encoder Feature Aggregation (EFA), on both datasets. The experimental results for this investigation are shown in Table 1. It can be observed that the UTB, AMP and EFA modules all achieve clear improvement. Specifically, as a core module of the proposed InvPT framework, UTB significantly improves the task Semseg by 2.12 (mIoU), Depth by 0.0707 (RMSE) and Normal by 0.42 (mErr) on NYUD-v2. Finally, the full model of InvPT achieves remarkable performance gain compared against both the Single-task (ST) baseline (see MT performance  $\Delta_m$ ) and the multi-task (MT) baseline (see the Gain on each task), clearly verifying the effectiveness of the proposed InvPT decoder.

For qualitative comparison, in Fig. 5, we show prediction examples generated by different model variants which add these modules one by one on PASCAL-Context. It is intuitively to observe that the proposed UTB, AMP and EFA all help produce visually more accurate predictions to against the baseline.

**Multi-task Improvement against Single-task Setting** To validate the effectiveness of the proposed multi-task model, we compare it with its single-task variant “InvPT Baseline (ST)” on both datasets in Table 1. Our full model achieves strong performance improvement against the single-task model, yielding **2.59%** multi-task performance on NYUD-v2 and **1.76%** on PASCAL-Context.

**Effect of Different Transformer Encoders** We also compare the model performance using two families of transformer encoders, *i.e.* Swin Transformer (Swin-T, Swin-B, and Swin-L) [22] and ViT (ViT-B and ViT-L) [14], used as our InvPT encoder. The results on PASCAL-Context are shown in Table 2.




Fig. 7: (a): Statistics of learned features with t-SNE [23] of all 20 classes on Pascal-Context. (b): Visualization of examples of learned feature of semantic segmentation. From both (a) and (b), we observe that features learned by InvPT is effectively improved and is more discriminative compared to the baseline.

We observe that the models with higher capacity in the same family can generally obtain consistent performance gain on semantic tasks including Semseg and Parsing, while on other tasks (*e.g.* Saliency) the improvement is however not clear. One possible reason for the difference of performance gain is the task competition problem in training as discussed in previous works [36,19].

**Effect of the Number of Stages** Our UP-Transformer block typically consists of three stages. In Fig. 6, we show how the number of stages of InvPT decoder influences the performance of the different tasks on PASCAL-Context. It can be observed that using more stages can help the InvPT decoder learn better predictions for all tasks. Our efficient design makes the multi-task decoding feature maps with gradually increased resolutions possible.

**Qualitative Study of learned Features with InvPT** In Fig. 7, we show visualization comparison of the learned final features between the transformer baseline (*i.e.* InvPT Baseline (MT)) and our InvPT full model, to further demonstrate how the features are improved using our proposed InvPT model. The statistics of the learned feature points is visualized with t-SNE [24] on all 20 semantic classes of Pascal-Context dataset. It is obvious that our model helps learn more discriminative features, thus resulting in higher quantitative results. The generated spatial feature maps for segmentation are also intuitively better.

**Generalization Performance** To qualitatively study the generalization performance of the proposed multi-task transformer for dense scene understanding, we compare it with the best performing methods, including ATRC [5] and PAD-Net [48], on the challenging DAVIS video segmentation Dataset [30]. The results are shown in the video demo in the github page. All the models are trained on PASCAL-Context with all the five tasks. Then the models are directly tested on DAVIS to generate multi-task predictions on video sequences. One example frame is shown in Fig. 8, which clearly shows our advantage on this perspective.

## 4.2 State-of-the-art Comparison

Table 3 shows a comparison of the proposed InvPT method against existing state-of-the-arts, including PAD-Net [48], MTI-Net [42] and ATRC [5], on both



Fig. 8: Study on generalization performance. Models are all trained on PASCAL-Context [9] and tested on DAVIS video dataset [30]. The proposed method yields better generalization performance compared to PAD-Net [48] and ATRC [5].

Table 3: State-of-the-art comparison on NYUD-v2 (*left*) and PASCAL-Context (*right*). Our InvPT significantly outperforms the previous state-of-the-arts by a large margin. ‘↓’ means lower better and ‘↑’ means higher better.

Model	Semseg mIoU ↑	Depth RMSE ↓	Normal mErr ↓	Boundary odsF ↑	Model	Semseg mIoU ↑	Parsing mIoU ↑	Saliency maxF ↑	Normal mErr ↓	Boundary odsF ↑
Cross-Stitch [27]	36.34	0.6290	20.88	76.38	ASTMT [25]	68.00	61.10	65.70	14.70	72.40
PAP [52]	36.72	0.6178	20.82	76.42	PAD-Net [48]	53.60	59.60	65.80	15.30	72.50
PSD [55]	36.69	0.6246	20.87	76.42	MTI-Net [42]	61.70	60.18	84.78	14.23	70.80
PAD-Net [48]	36.61	0.6270	20.85	76.38	ATRC [5]	62.69	59.42	84.70	14.20	70.96
MTI-Net [42]	45.97	0.5365	20.27	77.86	ATRC-ASPP [5]	63.60	60.23	83.91	14.30	70.86
ATRC [5]	46.33	0.5363	20.18	77.94	ATRC-BMTAS [5]	67.67	62.93	82.29	14.24	72.42
InvPT (ours)	<b>53.56</b>	<b>0.5183</b>	<b>19.04</b>	<b>78.10</b>	InvPT (ours)	<b>79.03</b>	<b>67.61</b>	<b>84.81</b>	<b>14.15</b>	<b>73.00</b>

NYUD-v2 and PASCAL-Context. On all the 9 metrics from these two benchmarks, the proposed InvPT achieves clearly superior performance, especially for higher-level scene understanding tasks such as Semseg and Parsing. Notably, on NYUD-v2, our InvPT surpasses the previous best performing method (*i.e.* ATRC) by **+7.23** (mIoU) on Semseg, while on PASCAL-Context, InvPT outperforms ATRC by **+11.36** (mIoU) and **+4.68** (mIoU) on Semseg and Parsing, respectively. A qualitative comparison with ATRC is shown in Fig. 5.

## 5 Conclusion

This paper presented a novel transformer framework, Inverted Pyramid Multi-task Transformer (InvPT), for the multi-task dense prediction for visual scene understanding. InvPT is able to effectively learn the long-range interaction in both spatial and all-task contexts on the multi-task feature maps with gradually increased spatial resolution for dense prediction. Extensive experiments demonstrated the effectiveness of the proposed method, and also showed its significantly better performance on two popular benchmarks compared to the previous state-of-the-art methods.

**Acknowledgements.** This research is supported in part by the Early Career Scheme of the Research Grants Council (RGC) of the Hong Kong SAR under grant No. 26202321 and HKUST Startup Fund No. R9253.

## References

1. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016) 8
2. Bai, Y., Mei, J., Yuille, A., Xie, C.: Are transformers more robust than cnns? In: NeurIPS (2021) 4
3. Bello, I., Zoph, B., Le, Q., Vaswani, A., Shlens, J.: Attention augmented convolutional networks. In: ICCV (2019) 2
4. Bhojanapalli, S., Chakrabarti, A., Glasner, D., Li, D., Unterthiner, T., Veit, A.: Understanding robustness of transformers for image classification. In: ICCV (2021) 4
5. Bruggemann, D., Kanakis, M., Obukhov, A., Georgoulis, S., Van Gool, L.: Exploring relational context for multi-task dense prediction. In: ICCV (2021) 2, 4, 11, 12, 13, 14
6. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020) 4
7. Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. In: CVPR (2021) 4
8. Chen, T., Saxena, S., Li, L., Fleet, D.J., Hinton, G.: Pix2seq: A language modeling framework for object detection. arXiv preprint arXiv:2109.10852 (2021) 4
9. Chen, X., Mottaghi, R., Liu, X., Fidler, S., Urtasun, R., Yuille, A.: Detect what you can: Detecting and representing objects using holistic models and body parts. In: CVPR (2014) 10, 14
10. Chen, Z., Badrinarayanan, V., Lee, C.Y., Rabinovich, A.: Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In: ICML (2018) 3
11. Chen, Z., Ngiam, J., Huang, Y., Luong, T., Kretschmar, H., Chai, Y., Anguelov, D.: Just pick a sign: Optimizing deep multitask models with gradient sign dropout. In: NeurIPS (2020) 3
12. d’Ascoli, S., Touvron, H., Leavitt, M., Morcos, A., Biroli, G., Sagun, L.: Convit: Improving vision transformers with soft convolutional inductive biases. In: ICML (2021) 5
13. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009) 11
14. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021) 2, 4, 6, 7, 12
15. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. IJCV 88(2), 303–338 (2010) 10
16. Gao, Y., Bai, H., Jie, Z., Ma, J., Jia, K., Liu, W.: Mtl-nas: Task-agnostic neural architecture search towards general-purpose multi-task learning. In: CVPR (2020) 2
17. Gao, Y., Ma, J., Zhao, M., Liu, W., Yuille, A.L.: Nddr-cnn: Layerwise feature fusing in multi-task cnns by neural discriminative dimensionality reduction. In: CVPR (2019) 3, 4
18. Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. In: NeurIPS (2021) 4
19. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: CVPR (2018) 2, 3, 4, 13



20. Liang, M., Yang, B., Chen, Y., Hu, R., Urtasun, R.: Multi-task multi-sensor fusion for 3d object detection. In: CVPR (2019) [3](#)
21. Liu, Y., Sangineto, E., Bi, W., Sebe, N., Lepri, B., De Nadai, M.: Efficient training of visual transformers with small-size datasets. arXiv preprint arXiv:2106.03746 (2021) [5](#)
22. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021) [5](#), [6](#), [11](#), [12](#)
23. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research **9**(11) (2008) [13](#)
24. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. JMLR **9**(11) (2008) [13](#)
25. Maninis, K.K., Radosavovic, I., Kokkinos, I.: Attentive single-tasking of multiple tasks. In: CVPR (2019) [3](#), [4](#), [10](#), [11](#), [14](#)
26. Mao, J., Xue, Y., Niu, M., Bai, H., Feng, J., Liang, X., Xu, H., Xu, C.: Voxel transformer for 3d object detection. In: ICCV (2021) [4](#)
27. Misra, I., Shrivastava, A., Gupta, A., Hebert, M.: Cross-stitch networks for multi-task learning. In: CVPR (2016) [3](#), [4](#), [14](#)
28. Pan, X., Xia, Z., Song, S., Li, L.E., Huang, G.: 3d object detection with point-former. In: CVPR (2021) [4](#)
29. Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Shazeer, N., Ku, A., Tran, D.: Image transformer. In: ICML (2018) [4](#)
30. Perazzi, F., Pont-Tuset, J., McWilliams, B., Gool, L.V., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: CVPR (2016) [13](#), [14](#)
31. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. arXiv preprint arXiv:2103.00020 (2021) [4](#)
32. Ramachandran, P., Parmar, N., Vaswani, A., Bello, I., Levskaya, A., Shlens, J.: Stand-alone self-attention in vision models. In: NeurIPS (2019) [5](#)
33. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: ICML (2021) [4](#)
34. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: ICCV (2021) [2](#), [5](#)
35. Ruder, S., Bingel, J., Augenstein, I., Søgaard, A.: Latent multi-task architecture learning. In: AAAI (2019) [3](#)
36. Sener, O., Koltun, V.: Multi-task learning as multi-objective optimization. In: NeurIPS (2018) [13](#)
37. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgb-d images. In: ECCV (2012) [10](#)
38. Srinivas, A., Lin, T.Y., Parmar, N., Shlens, J., Abbeel, P., Vaswani, A.: Bottleneck transformers for visual recognition. In: CVPR (2021) [5](#)
39. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: ICML (2021) [5](#)
40. Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., Jégou, H.: Going deeper with image transformers. In: ICCV (2021) [5](#)
41. Vandenhende, S., Georgoulis, S., Van Gansbeke, W., Proesmans, M., Dai, D., Van Gool, L.: Multi-task learning for dense prediction tasks: A survey. TPAMI (2021) [1](#), [3](#)

42. Vandenhende, S., Georgoulis, S., Van Gool, L.: Mti-net: Multi-scale task interaction networks for multi-task learning. In: ECCV (2020) [3](#), [4](#), [7](#), [12](#), [13](#), [14](#)
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017) [4](#), [9](#)
44. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: ICCV (2021) [2](#), [5](#), [7](#), [9](#)
45. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: CVPR (2018) [2](#)
46. Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L.: Cvt: Introducing convolutions to vision transformers. arXiv preprint arXiv:2103.15808 (2021) [5](#), [7](#), [9](#)
47. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. In: NeurIPS (2021) [5](#)
48. Xu, D., Ouyang, W., Wang, X., Sebe, N.: Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In: CVPR (2018) [1](#), [2](#), [3](#), [4](#), [13](#), [14](#)
49. Yang, J., Li, C., Zhang, P., Dai, X., Xiao, B., Yuan, L., Gao, J.: Focal self-attention for local-global interactions in vision transformers. In: NeurIPS (2021) [5](#)
50. Yuan, Y., Fu, R., Huang, L., Lin, W., Zhang, C., Chen, X., Wang, J.: Hrformer: High-resolution transformer for dense prediction. In: NeurIPS (2021) [2](#), [5](#), [7](#)
51. Zhang, X., Zhou, L., Li, Y., Cui, Z., Xie, J., Yang, J.: Transfer vision patterns for multi-task pixel learning. In: ACMMM (2021) [3](#), [4](#)
52. Zhang, Z., Cui, Z., Xu, C., Yan, Y., Sebe, N., Yang, J.: Pattern-affinitive propagation across depth, surface normal and semantic segmentation. In: CVPR (2019) [2](#), [3](#), [4](#), [14](#)
53. Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: ICCV (2021) [4](#)
54. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: CVPR (2017) [7](#)
55. Zhou, L., Cui, Z., Xu, C., Zhang, Z., Wang, C., Zhang, T., Yang, J.: Pattern-structure diffusion for multi-task learning. In: CVPR (2020) [2](#), [3](#), [4](#), [14](#)