

# TSP-Transformer: Task-Specific Prompts Boosted Transformer for Holistic Scene Understanding

Shuo Wang<sup>1</sup> Jing Li<sup>2</sup> Zibo Zhao<sup>1</sup> Dongze Lian<sup>3</sup>  
Binbin Huang<sup>1</sup> Xiaomei Wang<sup>4</sup> Zhengxin Li<sup>1</sup> Shenghua Gao<sup>1,5,6†</sup>  
<sup>1</sup>ShanghaiTech University <sup>2</sup>Xiaohongshu Inc. <sup>3</sup>National University of Singapore  
<sup>4</sup>Fudan University <sup>5</sup>Shanghai Engineering Research Center of Intelligent Vision and Imaging  
<sup>6</sup>Shanghai Engineering Research Center of Energy Efficient and Custom AI IC  
{wansghuo2022, zhaozb, liandz, huangbb, lizhx, gaoshh}@shanghaitech.edu.cn  
lijing1@alumni.shanghaitech.edu.cn, 17110240025@fudan.edu.cn

## Abstract

Holistic scene understanding includes semantic segmentation, surface normal estimation, object boundary detection, depth estimation, etc. The key aspect of this problem is to learn representation effectively, as each subtask builds upon not only correlated but also distinct attributes. Inspired by visual-prompt tuning, we propose a **Task-Specific Prompts Transformer**, dubbed TSP-Transformer, for holistic scene understanding. It features a vanilla transformer in the early stage and tasks-specific prompts transformer in the lateral stage, where tasks-specific prompts are augmented. By doing so, the transformer layer learns the generic information from the shared parts and is endowed with task-specific capacity. First, the tasks-specific prompts serve as induced priors for each task effectively. Moreover, the task-specific prompts can be seen as switches to favor task-specific representation learning for different tasks. Extensive experiments on NYUD-v2 and PASCAL-Context show that our method achieves state-of-the-art performance, validating the effectiveness of our method for holistic scene understanding. We also provide our code in the following link <sup>1</sup>.

## 1. Introduction

Understanding a scene [12, 15, 34] both semantically and geometrically is promising due to its lots of applications, such as autonomous driving, virtual reality, robotics, etc. The holistic understanding of the scene includes a wide variety of tasks, e.g., semantic segmentation, depth estimation, object boundary detection, surface normal estimation, etc. It is desirable to train a network that can deal with a

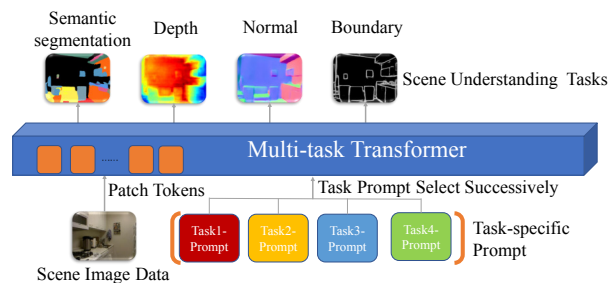


Figure 1. Overview of our proposed Task-Specific Prompts Transformer Network framework. We develop task-specific prompts for each task, which successively interact with the encoder to generate features relevant to the respective tasks. Task-specific prompts boost Multi-task Transformer with task-specific features and output the final multi-task predictions.

series of tasks jointly. By leveraging the task association and task specialty, impressive results on scene understanding are achieved by employing Convolution Neural Networks (CNNs) [4, 13, 33, 39]. Recently, transformer models [1, 36, 41, 43] have been introduced to model the multiple scene understanding tasks.

Previous works have designed different modules to learn task-specific and task-unified knowledge and interact with the information of these various tasks. For example, lots of works have designed special modules in the encoder [9, 23, 26] to learn task-specific representations and incorporate cross-task information interactions by means of manually designed structures. While the decoder [16, 33, 41] serves to decouple the characteristics of different tasks and facilitate cross-task interaction, some parameters are shared across all tasks. However, these methods are not fine enough for decoupling task-specific features and require carefully designed modules for multi-task dense learning.

Recently, building a unified network to perform multiple tasks has received more attention due to its potential for

<sup>1</sup><https://github.com/tb2-sy/TSP-Transformer>

general intelligence. A large number of transformer-based frameworks [3, 7, 11] have been constructed in the field of natural language processing to handle many different but related tasks. In computer vision, some recent approaches [5] design unified models [2] to obtain the output results based on the task description by implicitly learning the relevance of tasks, which show the feasibility of a unified model given instructions to perform multiple vision tasks. Inspired by these unified models [5, 27], we propose to utilize multi-task visual prompts to improve the multi-task scene understanding. Each prompt is dedicated to a specific task, which can be regarded as task-specific. These multi-task visual prompts are inserted into the input of the transformer layer along with the image tokens to perform self-attention such that the transformer layer learns the task-specific information. Such a method decouples characteristics of different tasks in the layer level, which is finer and more flexible.

Different from the previous unified models [5], where the explicit task description is inserted into the transformer to specify the task name, *e.g.*, ‘detect’ or ‘segment’, our task-specific prompts are learnable and updated independently for each task during training. Thanks to these learned task-specific prompts, the proposed model demonstrates the ability to produce task-specific output, which provides tailored results for specific requirements and is better adapted to the data distribution, and largely improves the performance. Furthermore, we also introduce an efficient task-specific feature fusion module to aggregate the features from the encoder of different tasks and pass them to the decoder.

We conduct extensive experiments on NYUD-v2 and PASCAL-Context for performance evaluation. Our experiments reveal that when task-specific features are inserted into the deeper layers of the encoder, then different tasks tend to have similar low-level features, and as we gradually introduce more and more task-specific prompts in deeper encoder layers, features of various tasks become less and less correlated, which helps task-specific features learning. Further, although tasks are closely correlated, our experiments show that introducing unified prompts is adverse to performance. The experiment results demonstrate that our proposed method is superior to other baseline methods, which validates the effectiveness of our proposed approach. In summary, our main contributions are as follows:

- We are introducing a novel approach, Task-Specific Prompts Transformer, to include task-specific prompts and various task features in multi-task training of scene understanding.
- We explore a new task-specific feature decoupling method; task-specific information can be finely decoupled with task-specific prompts. And we also design an efficient task-specific encoder feature fusion block to aggregate the encoder features of different tasks.

- We conduct extensive experiments on two benchmarks: NYUD-v2 and PASCAL-Context. The experiment results demonstrate that our proposed method is superior to other baseline methods.

## 2. Related Work

**Multi-task learning for dense scene understanding.** Multi-task learning (MTL) is a promising area [12, 18, 34] of research to improve generalization performance by using domain knowledge contained in supervised signals of associated tasks. Many previous works [4, 33, 39, 41] have explored lots of possibilities in this field. In particular, PAD-Net [39] proposes a new multi-task-oriented predictive distillation network structure that supplies abundant multi-modal data for learning the target task. MTI-Net [33] explicitly considers task interactions at multiple scales by designing the network structure and utilizing multi-scale information. Multi-scale processing preserves the original image’s features across various scales, such as the patch level. Neural architecture search (NAS) [19, 22, 38] techniques have also been employed in the field of multi-task learning, such as ATRC [4], with the help of knowledge distillation techniques to obtain richer information in the case of multi-task information sharing. The above methods are almost implemented based on CNNs. Owing to the exceptional performance of transformers in the visual domain, many researches [41] are being explored using Transformers [35]. Based on the structure of the Vision Transformer, InvPT [41] improves the decoder to expand the receptive field and realizes information interaction at higher resolutions. In this paper, our work explores the role of the Visual Prompt Transformer in multi-task learning and proposes a new method to use Visual Prompt Learning [10] in multi-task learning.

**Visual prompt learning.** Prompt learning first emerged in the field of natural language processing [14, 17, 24, 28, 31] for parameter-efficient tuning, which enables the model to better understand tasks and improve performance. Some large models in the field of natural language processing, such as GPT-3 [3], have demonstrated the ability to transfer under small sample conditions under the blessing of the prompt. Recent work has extended the prompt learning or parameter-efficient learning to the field of computer vision [10, 20, 29, 37, 40, 45], and the modality of the prompt is not limited to text. VPT [10] inserts some learnable parameters into the Vision Transformer Encoder Layer, and these learnable token parameters interact with the input image token, thereby affecting downstream tasks. However, the prompts method [10] mentioned above in the visual field are limited to fine-tune classification tasks and have not been tried in other settings, such as multi-task learning.

Recently, we noticed that there are some other concurrent works with similar ideas yet distinct meth-

ods, such as [21, 42]. Although both our model and TaskPrompter [42] enhance model performance for multi-task scene understanding with the proposed prompts to guide the learning process, our implementation methods and TaskPrompter have their own advantages and disadvantages. TaskPrompter uses two prompting paradigms in spatial and channel, called Spatial Task Prompting and Channel Task Prompting. They design a set of spatial-channel task prompts and learn their spatial and channel interactions with the shared image tokens in each transformer layer with attention mechanism. However, our method uses the standard visual prompt paradigm [10]. The self-attention mechanism generates task-specific features by sequentially concatenating the learnable prompts of different tasks and image patch tokens. Task-specific features and task fusion features are used as the input of the decoder stage.

### 3. Method

#### 3.1. Overview

In this section, we introduce our Task-Specific Prompts Transformer Network. The overall framework can be shown in Fig. 2. Our method is implemented on [41], consisting of three core parts, Task-Specific Prompts Transformer Encoder in Sec. 3.2, Efficient Multi-task Feature Fusion block in Sec. 3.3 and Multi-head Transformer Decoder in Sec. 3.4. The details of these parts are introduced as follows. Task-Specific Prompts Transformer Encoder consists of vanilla transformer encoder layers [35] in the early stage and our proposed task-specific prompts transformer encoder layers in the late stage. A few shared vanilla transformer encoder layers in the early stage can reduce the number of parameters and computations and learn general knowledge. Task-specific prompts transformer encoder layers decouple different task features and learn task-specific information. Task-specific prompts are introduced to the task-specific prompts transformer encoder layers interacts with the image feature to learn corresponding task dependencies. Efficient Multi-task Feature Fusion block aggregates fundamental encoder features from different tasks to pass to the decoder stage, which is vital for some low-level tasks. Multi-head Transformer decoder learns to produce refined scene representations within global spatial and task information, which are further used to produce the final predictions with task-specific output head.

#### 3.2. Task-Specific Prompts Transformer Encoder

As shown in Fig. 2, our Task-Specific Prompts Transformer Encoder mainly consists of three parts, patch embedding, shared vanilla transformer encoder layers in the early stage, and task-specific prompts transformer encoder layers in the late stage. A few shared vanilla transformer layers in the early stage of the encoder, and task-specific

prompts are introduced from particular layers in our default model. A few shared encoder layers in the early stage can reduce the number of parameters and computations and learn general knowledge. Task-specific prompts encoder layers in the late stage decouple different task features and learn task-specific information, and this decoupling is finely performed at the layer level.

**Patch embedding.** For a vanilla Vision Transformer (ViT) architecture with  $N$  layers, an input image  $I \in \mathbb{R}^{H \times W \times 3}$  is divided into  $m$  fixed-sized patches  $I_j$ . Each image patch  $I_j$  is then first embedded into a  $D$ -dimensional embedding token with positional encoding:

$$\mathbf{e}_0^j = \text{Embed}(I_j) \quad \mathbf{e}_0^j \in \mathbb{R}^D, j = 1, 2, \dots, m \quad (1)$$

The embeddings of the  $j$ -th image patch after embedded are denoted as  $\mathbf{e}_0^j$  ( $\mathbf{e}_0^j \in \mathbb{R}^D, j = 1, 2, \dots, m$ ). We use  $\mathbf{E}_i = \text{Concat}(\mathbf{e}_i^1, \dots, \mathbf{e}_i^m)$  as the collection of image patch token embeddings.

**Shared vanilla transformer encoder layer.** After image patch embedding,  $\mathbf{E}_i$  is then input to the  $(i + 1)$ -th Transformer encoder layer  $L_{i+1}$ :

$$[\mathbf{x}_{i+1}, \mathbf{E}_{i+1}] = L_{i+1}([\mathbf{x}_i, \mathbf{E}_i]) \quad i = 0, 1, 2, \dots, M - 1 \quad (2)$$

where  $\mathbf{x}_i \in \mathbb{R}^D$  denote learnable [CLS]'s embedding at  $L_i$  transformer encoder layer output space.

The shared vanilla transformer encoder layer  $L$  consists of two main blocks of multi-head self-attention (MSA) and multi-layer perceptron (MLP). Given an input  $\mathbf{E}_i \in \mathbb{R}^{m \times D}$ , where  $D$  is the embedding dimension, MSA first maps  $\mathbf{E}_i$  to queries  $\mathbf{Q} \in \mathbb{R}^{m \times d}$ , keys  $\mathbf{K} \in \mathbb{R}^{m \times d}$  and values  $\mathbf{V} \in \mathbb{R}^{m \times d}$  using three projection matrices,  $\mathbf{W}_q \in \mathbb{R}^{D \times d}$ ,  $\mathbf{W}_k \in \mathbb{R}^{D \times d}$  and  $\mathbf{W}_v \in \mathbb{R}^{D \times d}$ , where  $d$  denote the hidden layer dimension. And MSA computes the weighted sums over the values based on the self-attention between the queries and keys as follows:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (3)$$

where  $\frac{1}{\sqrt{d}}$  is a scaling factor.

**Task-specific prompts transformer encoder layer.** We concatenate  $n$  task-specific prompt tokens  $p \in \mathbb{R}^D$  for each task which are initialized with a fixed value or a random distribution to the Task-specific prompts transformer encoder Layer for each task to learn task-specific features. Task-specific prompts forward separately and interact with the patch token embedding features through MSA to learn dependencies for different tasks. Task-specific prompts are introduced at the input space of particular transformer layers, and distinct layers introduced may yield varying performance. For  $(M + 1)$ -th Layer  $L_{M+1}$ , we denote the collection of task  $t$  input learnable task-specific prompts as

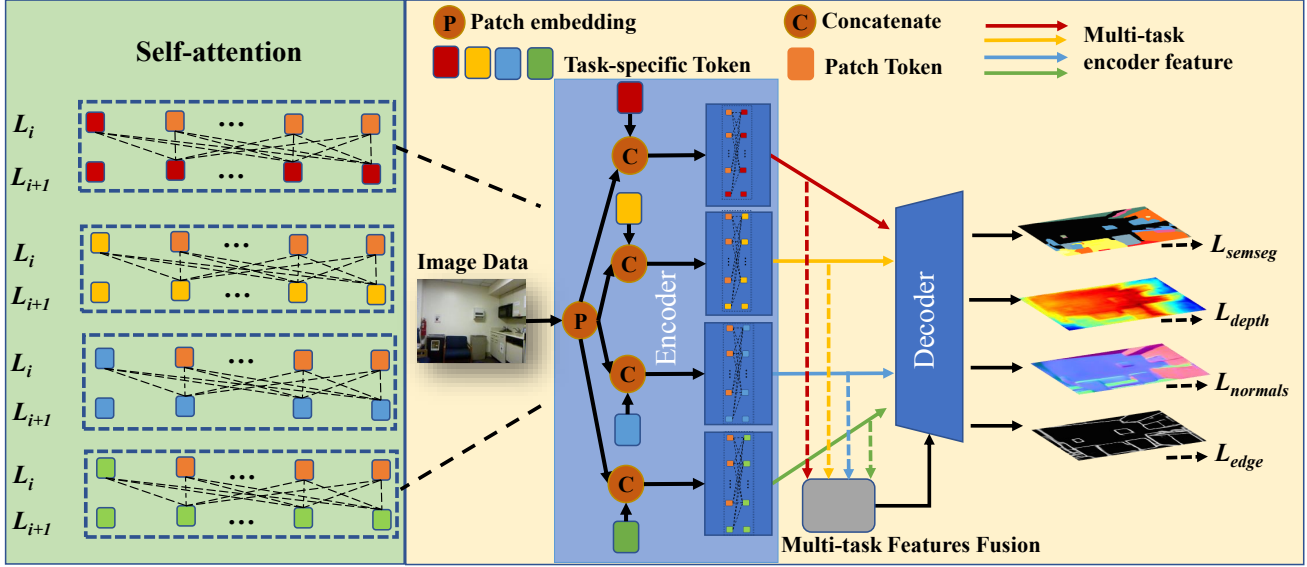


Figure 2. Pipeline of our proposed Task-Specific Prompts Transformer. It consists of three parts: Task-specific prompts, Encoder, and Decoder. Given an image and a specific task, the task prompts are adopted to concatenate with the image feature tokens and perform self-attention in the Encoder. Successively perform the above steps to generate multi-task features from the Encoder. The features output from the Encoder are fused and fed into the Decoder to perform multi-task prediction.

$\mathbf{P}_0^t = \{\mathbf{p}_0^k \in \mathbb{R}^D \mid k \in \mathbb{N}, 1 \leq k \leq n\}$ , denoted  $t = 1, 2, \dots, T$ , where  $T$  is the number of tasks. The task-specific prompts transform encoder layer is formulated as:

$$\begin{aligned} [\mathbf{x}_{M+1}^t, \mathbf{E}_{M+1}^t, -] &= L_{M+1}([\mathbf{x}_M, \mathbf{E}_M, \mathbf{P}_0^t]) \\ &\dots \\ [\mathbf{x}_N^t, \mathbf{E}_N^t, -] &= L_N([\mathbf{x}_{N-1}^t, \mathbf{E}_{N-1}^t, \mathbf{P}_{N-M-1}^t]) \end{aligned} \quad (4)$$

where ‘-’ means tokens are discarded and replaced by next-layer task-specific prompt tokens. If we choose to introduce task-specific prompt tokens  $\mathbf{P}_0^t$  in the  $i$ -th encoder layer, then  $\mathbf{P}_0^t$  concatenate of image patch token embedding  $E_{i-1}$  and learnable [CLS]’s embedding  $\mathbf{x}_{i-1}$ . Then, patch token embeddings after task-specific prompt interactions turn into task-specific features from encoder. In this way, task-specific prompts encoder layers achieve the decoupling of different task features and learn task-specific information, and this operation is finely performed at the layer level. Task-specific prompt tokens are learned by the corresponding task ground-truth back-propagation to update independently for each different task during training. The forward process of testing is the same as that of training. Task-specific prompt tokens for each task follow the same pipeline above, and interact with other shared parameters in the encoder separately.

### 3.3. Efficient Multi-task Features Fusion

Features from encoder are vital for some low-level tasks, such as boundary detection. To address the computa-

tional cost and quadratic complexity of the Transformer, a common strategy is to downsample the feature map to a lower resolution and then handle the multi-scale features from different transformer encoder layers. InvPT [41] decoder designs a multi-scale features from encoder aggregation method, but these aggregate features are just multi-scale instead of multi-task and multi-scale features from encoder. In our task-specific prompts transformer encoder, task-specific prompts encoder corresponding task features. Multi-task features fusion mechanism is necessary, which can play the role of information complementation between different tasks.

In order to better aggregate multi-task features, we apply aggregation weight matrix  $W \in \mathbb{R}^{4 \times T}$  (fixed or learnable) for different  $T$  tasks features  $F_i^t$  ( $t = 1, 2, \dots, T$ ) to obtain  $i$ -th scale multi-task aggregation features  $F_i$  as follows:

$$\mathbf{F}_i = \sum_{t=1}^T (\mathbf{W}_{t,i} \otimes \mathbf{F}_i^t) \quad (5)$$

where  $i=1,2,3,4$ , and  $\otimes$  means element-wise multiplication after broadcast operation. Our experimental results show that aggregation weight matrix full of fixed-weight  $\frac{1}{T}$  achieves a reasonably high level of performance. We carried out an enumeration experiment focusing on fixed weight cases and observed that the absence of encoder information for any task leads to performance drops. Furthermore, the complementarity of information between tasks is crucial, and more details can be found in Sec. 4.3.



### 3.4. Multi-head Transformer Decoder

Our MTL model uses multi-task features from encoder as input for Multi-head Transformer Decoder. In the encoder stage, we obtain enough task-specific information. The purpose of the decoder is to interact and refine task-specific features.

For the decoder part, we adopted the design structure of Inverted Pyramid Transformer Decoder [41], which is mainly composed of a multi-task up-transformer block that is applied to gradually increases the spatial resolution of the feature map and also perform multi-task features interaction to refine all tasks.

### 3.5. Loss Function

After the multi-task decoder heads output the prediction results, we apply task-specific loss functions for each task to computer loss with ground truth. For pixel-by-pixel classification tasks: semantic segmentation, human parsing, saliency detection, and boundary detection, the cross-entropy loss is uniformly used for optimization. L1-Loss is used for depth estimation and surface normals estimation. The whole model can be end-to-end optimized. We set  $\lambda_{seg}=1.0$ ,  $\lambda_{depth}=1.0$ ,  $\lambda_{normals}=10.0$ ,  $\lambda_{edge}=50.0$ ,  $\lambda_{partseg}=2.0$ , and  $\lambda_{sal}=5.0$ .  $\lambda$  represents the weight of the loss corresponding to each task:

$$\mathcal{L} = \lambda_{seg} \mathcal{L}_{seg} + \lambda_{depth} \mathcal{L}_{depth} + \lambda_{normals} \mathcal{L}_{normals} + \lambda_{edge} \mathcal{L}_{edge} + \lambda_{partseg} \mathcal{L}_{partseg} + \lambda_{sal} \mathcal{L}_{sal} \quad (6)$$

We use the exact same loss function weights as those in InvPT [41], without making any adjustments.

## 4. Experiments

### 4.1. Experimental Setup

**Datasets.** Following previous works, we use NYUD-v2 and PASCAL-Context datasets for performance evaluation. NYUD-v2 [32] contains 795 training images and 654 testing images for indoor scenes, including four tasks: semantic segmentation, monocular depth estimation, surface normal estimation, and object boundary detection. PASCAL-Context [8] contains 4998 training images and 5105 testing images. It labels for five tasks: semantic segmentation, human parts parsing, monocular depth estimation, surface normal estimation, and object boundary detection. We conduct experiments on both datasets for performance evaluation.

**Evaluation metric.** Semantic segmentation and human parts parsing tasks both are segmentation tasks and are evaluated with the mean Intersection over Union (IoU) and monocular depth estimation task using the Root Mean Square Error (RMSE) as the evaluation metric. The surface normal estimation task uses mean angular error (mErr) as the evaluation metric, saliency detection is evaluated with

the maximum F-measure (maxF), and the boundary detection task uses the optimal dataset-scale F-measure (odsF) score as the evaluation metric.

**Implementation details.** We mainly perform ablation studies using Vision Transformer(including Vit-L, Vit-B) pre-trained on ImageNet-21K [6] as the transformer encoder on the NYUD-v2 dataset. We set the training iterations of the model to 40k on the NYUD-v2 dataset and the PASCAL-Context dataset, both with a batch size of 6 using NVIDIA A40 GPU. For the optimization, Adam optimizer is used, and the learning rate is set to  $2 \times 10^{-5}$  with a weight decay rate of  $1 \times 10^{-6}$ . A polynomial learning rate scheduler is used.

Method	Semseg (IoU)↑	Depth (RMSE)↓	Normal (mErr)↓	Boundary (odsF)↑
Cross-Stitch [26]	36.34	0.6290	20.88	76.38
PAP [44]	36.72	0.6178	20.82	76.42
PSD [46]	36.69	0.6246	20.87	76.42
PAD-Net [39]	36.61	0.6270	20.85	76.38
MTI-Net [33]	45.97	0.5365	20.27	77.86
ATRC [4]	46.33	0.5363	20.18	77.94
InvPT [41]	53.56	0.5183	19.04	78.10
TaskPrompter [42]	55.30	0.5152	18.47	<b>78.20</b>
<b>Ours</b>	<b>55.39</b>	<b>0.4961</b>	<b>18.44</b>	77.50

(a) State-of-the-art comparison on NYUD-v2.

Method	Semseg (IoU)↑	Parsing (IoU)↑	Saliency (maxF)↑	Normal (mErr)↓	Boundary (odsF)↑
ASTMT [25]	68.00	61.10	65.70	14.70	72.40
PAD-Net [39]	53.60	59.60	65.80	15.30	72.50
MTI-Net [33]	61.70	60.18	84.78	14.23	70.80
ATRC [4]	62.69	59.42	84.70	14.20	72.96
ATRC-ASPP [4]	63.60	60.23	83.91	14.30	70.86
ATRC-BMTAS [4]	67.67	62.93	82.29	14.24	72.42
InvPT [41]	79.03	67.61	84.81	14.15	73.00
TaskPrompter [42]	80.89	68.89	84.83	13.72	73.50
<b>Ours</b>	<b>81.48</b>	<b>70.64</b>	<b>84.86</b>	<b>13.69</b>	<b>74.80</b>

(b) State-of-the-art comparison on PASCAL-Context.

Table 1. State-of-the-art comparison on NYUD-v2 (a) and PASCAL-Context (b). Our method significantly outperforms the previous state-of-the-arts methods.

### 4.2. Comparisons With State-of-the-art Methods

We compare our method with the state-of-the-art methods, including PAD-Net, MTI-Net, ATRC, InvPT and TaskPrompter. The results on the NYUD-v2 and PASCAL-Context datasets are reported in Table 1 (a) and Table 1 (b), respectively. On NYUD-v2, the performance improvement for semantic segmentation, depth estimation, and surface normal estimation is significant, while for boundary detection, it is comparable with state-of-the-art methods. All tasks, except boundary detection, achieve state-of-the-art results on the NYUD-v2 dataset in our experimental results. As discussed in previous works [12,30], one possible reason

Task Prompt	Semseg (IoU) $\uparrow$	Depth (RMSE) $\downarrow$	Normal (mErr) $\downarrow$	Boundary (odsF) $\uparrow$
	53.56	0.5183	19.04	<b>78.10</b>
✓	<b>55.39</b>	<b>0.4961</b>	<b>18.44</b>	77.50

(a) Task prompt ablation on NYUD-v2.

Task Prompt	Semseg (IoU) $\uparrow$	Parsing (IoU) $\uparrow$	Saliency (maxF) $\uparrow$	Normal (mErr) $\downarrow$	Boundary (odsF) $\uparrow$
	79.03	67.61	84.81	14.15	73.00
✓	<b>81.48</b>	<b>70.64</b>	<b>84.86</b>	<b>13.69</b>	<b>74.80</b>

(b) Task prompt ablation on PASCAL-Context.

Table 2. Task prompt ablation on on NYUD-v2 (a) and PASCAL-Context (b). Our proposed task-specific prompt significantly enhances multi-task performance.

for the difference in performance gain is the task competition problem in training. The reason for task competition is also confirmed in our supplementary material. When each task does not share the encoder, the performance of boundary detection is almost the same as that of InvPT [41]. Although the performance of boundary detection in Table 1 has declined, overall, our results are still state-of-the-art. To enhance the overall performance, we present the performance in Table 1. On Pascal-Context, the improvements in semantic segmentation, human parsing, saliency detection, surface normal estimation, and boundary detection are all significant.

### 4.3. Ablation Studies

**The effectiveness of task-specific prompts.** To verify the effectiveness of our proposed task prompt learning method, we conduct ablation studies on both NYUD-v2 and PASCAL-Context. The introduction of task prompts significantly enhances multi-task performance, as demonstrated in Table 2 (a) and Table 2 (b). The performance improves because task-specific prompt tokens facilitate learning independent embedding spaces for each task. Such a method decouples the characteristics of different tasks at the layer level, providing a finer and more flexible representation of various tasks. In the following section, we conduct comprehensive experiments to evaluate the effectiveness of task-specific prompts and derive numerous valuable insights, such as the positions of task-specific prompts hold greater significance than the number of prompt tokens per task in each layer and introducing shared prompt tokens results in unfavorable performance, among other conclusions.

**The effectiveness of efficient multi-task features fusion.**

As shown in Table 3, we explore a variety of efficient multi-task encoder features fusion strategies. In theory, employing learnable weights and utilizing cross-task attention offers a broader latent space than a fixed-weight strategy, potentially leading to better performance through an optimal solution. However, our experimental results indicate that

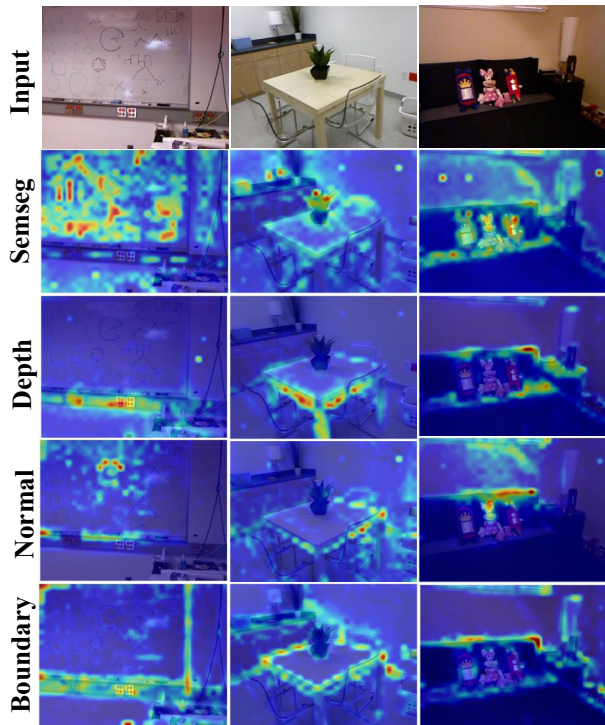


Figure 3. Visualizing the attention maps for task-specific prompts demonstrates their ability to focus on distinct spatial positions within patch tokens. This observation indicates that task prompts effectively learn task-specific representations through their interactions with image patch tokens.

more complex fusion strategies tend to yield poorer performance, suggesting that the role of this block is both critical and sensitive. Complex fusion strategies can increase instability in the model optimization process for our multi-objective optimization problem [30], leading to suboptimal performance. Moreover, multi-task features from encoder are essential, particularly for low-level tasks such as boundary detection. We conducted an enumeration experiment concentrating on fixed-weight scenarios and observed that the lack of encoder information for any task results in performance decreasing. Additionally, the complementarity of information between tasks is vital.

**Task-specific prompts vs. task-unified prompts.** In our approach, we propose that task-specific prompts encode the induced prior knowledge for each task, promoting task-specific representation learning. Besides using task-specific prompts, we also design experiments by introducing task-unified prompts, which are the same for all tasks and do not update independently. The task-specific prompts and task-unified prompts can be used in the following ways: 1) task-unified prompts only; 2) task-unified prompts and task-specific prompts are concatenated together; 3) task-unified prompts and task-specific prompts are blended with an element-wise summation; 4) task-specific prompt only. 5) task-unified prompts generated

Method	Semseg (IoU) $\uparrow$	Depth (RMSE) $\downarrow$	Normal (mErr) $\downarrow$	Boundary (odsF) $\uparrow$
w/o fusion	53.47	0.5113	19.59	69.50
fixed weight	<b>55.39</b>	<b>0.4961</b>	<b>18.44</b>	<b>77.50</b>
learnable weight	55.18	0.4998	18.59	<b>77.50</b>
cross-task attention	55.16	0.5066	18.60	77.30

(a) Ablation for efficient multi-task features fusion strategy on NYUD-v2.

Fixed weight	Semseg (IoU) $\uparrow$	Depth (RMSE) $\downarrow$	Normal (mErr) $\downarrow$	Boundary (odsF) $\uparrow$
0.25, 0.25, 0.25, 0.25	<b>55.39</b>	<b>0.4961</b>	<b>18.44</b>	<b>77.50</b>
0.00, 0.25, 0.25, 0.25	54.59	0.4993	18.72	77.00
0.25, 0.00, 0.25, 0.25	55.18	0.4980	18.70	77.00
0.25, 0.25, 0.00, 0.25	55.02	0.5019	18.65	77.40
0.25, 0.25, 0.25, 0.00	55.14	0.5031	18.68	77.20
0.25, 0.25, 0.00, 0.00	54.86	0.5000	18.63	77.10
0.25, 0.00, 0.00, 0.00	54.59	0.5012	18.74	77.30
0.00, 0.00, 0.00, 0.00	53.47	0.5113	19.59	69.50

(b) Ablation for efficient multi-task features fusion fixed weight on NYUD-v2. For the NYUD-v2 dataset consisting of four tasks, using a fixed weight of  $\frac{1}{4} = 0.25$  and arranged in the sequence of semantic segmentation, depth estimation, surface normal estimation, and boundary detection.

Table 3. Ablation for efficient multi-task features fusion module. We explore a variety of efficient multi-task encoder features fusion strategies (a) and achieve the best performance with fixed values. Concentrating on fixed-weight scenarios (b) and observed that the lack of encoder information for any task results in performance decreasing.

by fusing task-specific prompts with cross-prompt attention operation; task-specific prompts and fused task-unified prompts are blended with an element-wise summation. It is worth noting that 3) will also lead to task-specific prompts, but the optimization in 3) is different from that in 4) because the task-unified prompts are supervised by four different tasks while the task-specific prompts are supervised by the corresponding task only. Results in Table 4 show that task-unified prompts constantly undermine the overall performance, compared with our method with task-specific prompts only, where  $n$  and  $m$  mean the number of task-unified prompts and task-specific prompts. There are two possible reasons for the poor performance of task-unified prompts: i) it may be possible to exist the task-unified prompts for fewer tasks, but difficult to exist the task-unified prompts for all four tasks; ii) the encoder aims at explicitly decoupling the features of different tasks. Introducing the task-unified prompts in the encoder goes against the feature decoupling for different tasks.

**The positions of task-specific prompts.** In our implementation, the transformer encoder contains 24 layers in total. We denote the input as layer 1. In Table 5, we report the results with prompts inserted at multiple layers where  $i - j$  means the task-specific prompts are used from the  $i$ -th layer to the  $j$ -th layer. First, we can see that inserting task-specific prompts into more layers always boosts the performance for semantic segmentation and depth estimation, but for surface normal estimation and boundary detection, more lay-

$m$	$n$	method	Semseg (IoU) $\uparrow$	Depth (RMSE) $\downarrow$	Normal (mErr) $\downarrow$	Boundary (odsF) $\uparrow$
5	0	/	<b>55.39</b>	0.4961	<b>18.44</b>	77.50
5	1	concat	55.23	<b>0.4935</b>	18.65	76.90
5	5	concat	54.45	0.4940	18.63	77.40
5	5	add	54.82	0.5001	18.58	77.40
0	5	/	53.45	0.5042	18.81	77.30
0	20	/	53.43	0.5055	18.72	77.40
5	5	cross-prompt attention	55.38	0.5016	18.56	<b>77.80</b>

Table 4. Performance by different combinations of the task-unified and task-specific prompts on NYUD-v2. The task-specific prompts and task-unified prompts can be blended in the above ways where  $n$  and  $m$  mean the number of task-unified prompts and task-specific prompts. Results show that task-unified prompts always undermine the overall performance, compared with task-specific prompts only.

ers with visual prompts do not necessarily bring about the performance gain. Considering the trade-off between performance and computational costs, we use prompts for 12 layers. Further, we propose inserting prompts at different positions, and it shows that the performance is better by inserting the visual prompts at deeper layers of the encoder on NYUD-v2. Thus we choose to insert the visual prompts for layers between 13-24 in all our ablation experiments on NYUD-v2.

Layers with prompts	Semseg (IoU) $\uparrow$	Depth (RMSE) $\downarrow$	Normal (mErr) $\downarrow$	Boundary (odsF) $\uparrow$
w/o prompt	53.56	0.5183	19.04	<b>78.10</b>
1-6	54.24	0.4989	18.72	77.40
1-12	54.96	0.4948	18.72	77.40
1-18	55.45	0.4929	18.51	77.60
1-24	<b>55.80</b>	<b>0.4898</b>	18.63	77.60
1-12	54.96	0.4948	18.72	77.40
13-24	55.39	0.4961	<b>18.44</b>	77.50

Table 5. The performance with different prompts inserting positions on NYUD-v2. Inserting task-specific prompts into more layers always boosts the performance for semantic segmentation and depth estimation, but for surface normal estimation and boundary detection, more layers with prompts do not necessarily bring about the performance gain.

**The number of task-specific prompts per layer.** As shown in Table 6, the performance is not always improved when the number of prompt tokens increases. Performance decreases as the number of prompt tokens increases to a certain threshold, and the computational demand rises significantly. Moreover, each task appears to favor a distinct number of task-specific prompts. Depth, normal and boundary accuracy achieved the best results when token number  $N=5$ . Unlike these tasks, semantic segmentation achieves the best performance at  $N=50$ . However, the change in performance is relatively minor. For all experiments with task-specific prompts, even with only one prompt for each task, the performance is always better than without task-specific prompts.

Prompt Number	Semseg (IoU)↑	Depth (RMSE)↓	Normal (mErr)↓	Boundary (odsF)↑
w/o prompt	53.56	0.5183	19.04	<b>78.10</b>
1	55.27	0.4990	18.81	77.00
5	55.39	<b>0.4961</b>	<b>18.44</b>	77.50
10	55.19	<b>0.4961</b>	18.58	77.60
50	<b>55.61</b>	0.4972	18.66	77.10
100	55.18	0.5001	18.61	77.40
200	55.38	0.4971	18.69	77.30

Table 6. Ablation for Prompt Number on NYUD-v2. As the number of prompt tokens increases, the performance does not always get better.

**Influences of different backbones.** To validate the generalization of our task-specific prompts, we also use different backbones, and the results are shown in Table 7. We can see that with different backbones, our task-specific prompts always boost the performance for various scene-understanding tasks. So our task-specific prompts strategy can be compatible with other backbones.

Backbone	Semseg (IoU)↑	Depth (RMSE)↓	Normal (mErr)↓	Boundary (odsF)↑
InvPT(Vit-B)	50.30	0.5367	19.00	77.60
Ours(Vit-B)	51.22	0.5301	18.78	76.90
InvPT(Vit-L)	53.56	0.5183	19.04	<b>78.10</b>
Ours(Vit-L)	<b>55.39</b>	<b>0.4961</b>	<b>18.44</b>	77.50

Table 7. Ablation for Backbone on NYUD-v2. With different backbones, our task-specific prompts always boost the performance for different scene-understanding tasks.

**The task specialty of the different task-specific prompts.** To show the specialty of the different task-specific prompts, we replace the original corresponding task prompt tokens with the same trained task prompt tokens in inference. Results in Table 8 show that the performance drops by replacing the task-specific prompts with prompt tokens learned for the other tasks. The reason is that different task-specific prompts are pretty different from each other. It is worth noting that when we replace all prompts with task-specific prompts, such as those for semantic segmentation, the corresponding performance for the given task also changes due to the cross-task feature fusion module. To show whether the task prompts learn task-specific representation on patch tokens more intuitively, we visualize the attention map values between task prompts and patch tokens in the Encoder module as shown in Fig. 3. The attention map values are highly related to each task’s particularity, demonstrating that the task prompts can effectively encode task-specific representations.

**The correlation of features corresponding to different tasks.** In our experiments, we introduce the task-specific prompts from the 13th to the 24th layer. We also calculate the average correlation for features corresponding to different tasks. We denote the out space features corresponding

Model	Semseg (IoU)↑	Depth (RMSE)↓	Normal (mErr)↓	Boundary (odsF)↑
task-specific prompt	<b>55.39</b>	0.4961	18.44	77.50
segmentation prompt	55.35	0.5596	20.58	77.50
depth prompt	48.30	<b>0.4950</b>	18.76	77.00
normal prompt	40.85	0.5533	<b>18.43</b>	76.60
edge prompt	52.88	0.5497	20.15	<b>77.60</b>

Table 8. Ablation for prompt task-specialty. The performance drops by replacing the task-specific prompts with prompt tokens learned for other tasks.




Figure 4. Investigating the correlation of features corresponding to different tasks in {13, 18, 24} – th layers output space. Such as semseg-depth to indicate the feature correlation between segmentation and depth estimation.

to the  $k$ -th layer as  $F_k^A$  for task-A and denote the features of the  $k$ -th of the same image as  $F_k^B$ , then we calculate the cosine distance between  $F_k^A$  and  $F_k^B$  and do the average over all the test images. The results are shown in Fig. 4. We can see that as the layer goes deeper, as we introduce more and more task-specific prompts, the features corresponding to different tasks become more and more uncorrelated.

## 5. Conclusion

This paper presents a task-specific prompts-boosted transformer for holistic scene understanding, where the task-specific prompts are fed into different input layers of the transformer encoder for different tasks while the parameters of the transformer layer are shared across different tasks. Then we fuse task-specific features of different tasks with a task-fusion module and feed the fused feature into the transformer decoder scene prediction. Extensive experiments show that our method achieves state-of-the-art performance, which validates the effectiveness of our approach.

**Acknowledgement.** The work was supported by NSFC #61932020, #62172279, Program of Shanghai Academic Research Leader, and “Shuguang Program” supported by Shanghai Education Development Foundation and Shanghai Municipal Education Commission.



## References

- [1] Yutong Bai, Jieru Mei, Alan L Yuille, and Cihang Xie. Are transformers more robust than cnns? *Advances in Neural Information Processing Systems*, 34:26831–26843, 2021. 1
- [2] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 2
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 2
- [4] David Brüggenmann, Menelaos Kanakis, Anton Obukhov, Stamatios Georgoulis, and Luc Van Gool. Exploring relational context for multi-task dense prediction. In *Proc. IEEE/CVF Int. Conf. Computer Vision*, pages 15869–15878, 2021. 1, 2, 5
- [5] Ting Chen, Saurabh Saxena, Lala Li, Tsung-Yi Lin, David J Fleet, and Geoffrey Hinton. A unified sequence interface for vision tasks. *arXiv preprint arXiv:2206.07669*, 2022. 2
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conf. Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009. 5
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [8] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010. 5
- [9] Yuan Gao, Jiayi Ma, Mingbo Zhao, Wei Liu, and Alan L Yuille. Nddr-cnn: Layerwise feature fusing in multi-task cnns by neural discriminative dimensionality reduction. In *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, pages 3205–3214, 2019. 1
- [10] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. *arXiv preprint arXiv:2203.12119*, 2022. 2, 3
- [11] Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438, 2020. 2
- [12] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 7482–7491, 2018. 1, 2, 5
- [13] Iasonas Kokkinos. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 6129–6138, 2017. 1
- [14] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 2
- [15] Wei-Hong Li, Xialei Liu, and Hakan Bilen. Learning multiple dense prediction tasks from partially annotated data. In *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, pages 18879–18889, 2022. 1
- [16] Wei-Hong Li, Xialei Liu, and Hakan Bilen. Universal representations: A unified look at multiple task and domain learning. *arXiv preprint arXiv:2204.02744*, 2022. 1
- [17] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 2
- [18] Dongze Lian, Lina Hu, Weixin Luo, Yanyu Xu, Lixin Duan, Jingyi Yu, and Shenghua Gao. Multiview multitask gaze estimation with deep convolutional neural networks. *IEEE transactions on neural networks and learning systems*, 30(10):3010–3023, 2018. 2
- [19] Dongze Lian, Yin Zheng, Yintao Xu, Yanxiong Lu, Leyu Lin, Peilin Zhao, Junzhou Huang, and Shenghua Gao. Towards fast adaptation of neural architectures with meta learning. In *International Conference on Learning Representations (ICLR)*, 2020. 2
- [20] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2
- [21] Xiwen Liang, Minzhe Niu, Jianhua Han, Hang Xu, Chun-jing Xu, and Xiaodan Liang. Visual exemplar driven task-prompting for unified perception in autonomous driving. *arXiv preprint arXiv:2303.01788*, 2023. 3
- [22] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. 2
- [23] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1871–1880, 2019. 1
- [24] Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*, 2021. 2
- [25] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive single-tasking of multiple tasks. In *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, pages 1851–1860, 2019. 5
- [26] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 3994–4003, 2016. 1, 5
- [27] Jia Ning, Chen Li, Zheng Zhang, Zigang Geng, Qi Dai, Kun He, and Han Hu. All in tokens: Unifying output space of visual tasks via soft token. *arXiv preprint arXiv:2301.02229*, 2023. 2
- [28] Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*, 2019. 2

- [29] Mark Sandler, Andrey Zhmoginov, Max Vladymyrov, and Andrew Jackson. Fine-tuning image transformers using learnable memory. In *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, pages 12155–12164, 2022. [2](#)
- [30] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31, 2018. [5](#), [6](#)
- [31] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020. [2](#)
- [32] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *European Conf. Computer Vision*, pages 746–760. Springer, 2012. [5](#)
- [33] Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Mti-net: Multi-scale task interaction networks for multi-task learning. In *European Conf. Computer Vision*, pages 527–543. Springer, 2020. [1](#), [2](#), [5](#)
- [34] Simon Vandenhende, Stamatios Georgoulis, Wouter Van Gansbeke, Marc Proesmans, Dengxin Dai, and Luc Van Gool. Multi-task learning for dense prediction tasks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 2021. [1](#), [2](#)
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [2](#), [3](#)
- [36] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proc. IEEE/CVF Int. Conf. Computer Vision*, pages 568–578, 2021. [1](#)
- [37] Chen Henry Wu, Saman Motamed, Shaunak Srivastava, and Fernando D De la Torre. Generative visual prompt: Unifying distributional control of pre-trained generative models. *Advances in Neural Information Processing Systems*, 35:22422–22437, 2022. [2](#)
- [38] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: Stochastic neural architecture search. *arXiv preprint arXiv:1812.09926*, 2018. [2](#)
- [39] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 675–684, 2018. [1](#), [2](#), [5](#)
- [40] Shicheng Xu, Liang Pang, Huawei Shen, and Xueqi Cheng. Match-prompt: Improving multi-task generalization ability for neural text matching via prompt learning. In *Proc. 31st ACM Int. Conf. Information & Knowledge Management*, pages 2290–2300, 2022. [2](#)
- [41] Hanrong Ye and Dan Xu. Inverted pyramid multi-task transformer for dense scene understanding. *arXiv preprint arXiv:2203.07997*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [42] Hanrong Ye and Dan Xu. Taskprompter: Spatial-channel multi-task prompting for dense scene understanding. In *The Eleventh Int. Conf. Learning Representations*, 2023. [3](#), [5](#)
- [43] Yuhui Yuan, Rao Fu, Lang Huang, Weihong Lin, Chao Zhang, Xilin Chen, and Jingdong Wang. Hrformer: High-resolution transformer for dense prediction. *arXiv preprint arXiv:2110.09408*, 2021. [1](#)
- [44] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Yan Yan, Nicu Sebe, and Jian Yang. Pattern-affinitive propagation across depth, surface normal and semantic segmentation. In *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, pages 4106–4115, 2019. [5](#)
- [45] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. [2](#)
- [46] Ling Zhou, Zhen Cui, Chunyan Xu, Zhenyu Zhang, Chaoqun Wang, Tong Zhang, and Jian Yang. Pattern-structure diffusion for multi-task learning. In *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, pages 4514–4523, 2020. [5](#)