

Multi-Task Dense Prediction via Mixture of Low-Rank Experts

Yuqi Yang^{1,2*} Peng-Tao Jiang^{2*} Qibin Hou^{1,3†} Hao Zhang² Jinwei Chen² Bo Li²

¹VCIP, CS, Nankai University ²vivo Mobile Communication Co., Ltd ³NKIARI, Shenzhen Futian

yangyq2000@mail.nankai.edu.cn, pt.jiang@vivo.com, andrewhoux@gmail.com

Abstract

Previous multi-task dense prediction methods based on the Mixture of Experts (MoE) have received great performance but they neglect the importance of explicitly modeling the global relations among all tasks. In this paper, we present a novel decoder-focused method for multi-task dense prediction, called Mixture-of-Low-Rank-Experts (MLoRE). To model the global task relationships, MLoRE adds a generic convolution path to the original MoE structure, where each task feature can go through this path for explicit parameter sharing. Furthermore, to control the parameters and computational cost brought by the increase in the number of experts, we take inspiration from LoRA and propose to leverage the low-rank format of a vanilla convolution in the expert network. Since the low-rank experts have fewer parameters and can be dynamically parameterized into the generic convolution, the parameters and computational cost do not change much with the increase of experts. Benefiting from this design, we increase the number of experts and its reception field to enlarge the representation capacity, facilitating multiple dense tasks learning in a unified network. Extensive experiments on the PASCAL-Context and NYUD-v2 benchmarks show that our MLoRE achieves superior performance compared to previous state-of-the-art methods on all metrics. Our code is available at <https://github.com/YuqiYang213/MLoRE>.

1. Introduction

Computer vision tasks, such as semantic segmentation [5, 29, 32, 51] and depth estimation [2, 37], have been significantly facilitated by deep learning techniques. Each vision task has its own elaborated deep models that usually follow a similar pipeline, *i.e.*, feature extraction and prediction. Besides, some tasks also share relations. These facts motivate researchers to study Multi-Task Learning (MTL) that is able to unify different task models into a single one.

*The first two authors contributed equally to this paper. Work was done when Yuqi Yang was an intern at vivo.

†Qibin Hou is the corresponding author.

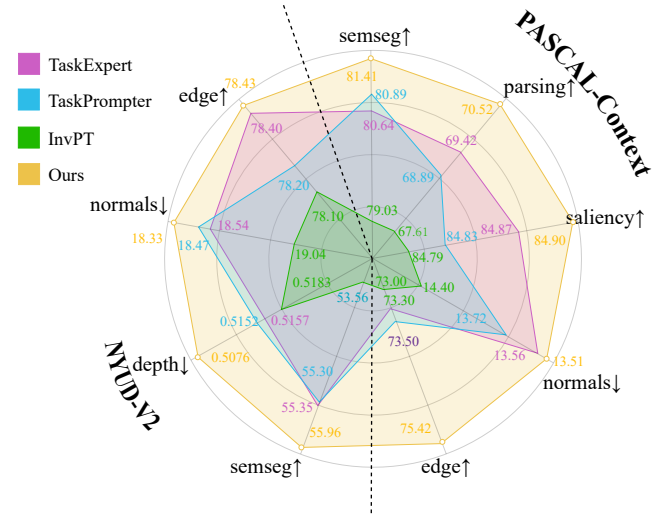


Figure 1. Performance comparison with state-of-the-art methods. Our MLoRE based on the proposed mixture of low-rank experts achieves superior performance on all tasks. ↑ denotes higher is better. ↓ denotes lower is better.

The significant advantage of multi-task learning lies in that it can improve training and inference efficiency while keeping commensurate performance to each task model. Due to such an advantage, MTL models have been applied to several directions, including autonomous driving [23, 28, 52] and scene understanding [46, 48], *etc.*

In this paper, we focus on multi-task learning for dense scene understanding, in which each task predicts pixel-wise results. A research line of early works [4, 16, 33, 35, 45, 46, 48] focuses on designing delicate network architectures, including encoder-focused and decoder-focused methods. The encoder-focused methods [16, 33, 35] design hand-crafted modules to share across task-specific encoders to construct task-generic features, while the decoder-focused methods [4, 45, 46, 48] propose tailored decoders to learn discriminative task-specific representations and build cross-task relationships.

Unlike the above methods focusing on designing static network architectures, some methods [6, 10, 14] introduce the mixture-of-experts (MoE) technique that sheds light on

Table 1. Parameters and FLOPs of different settings in the standard MoE and the proposed MLoRE. The left setting presents the number of experts and the kernel size of the convolutions in the expert network.

Settings	Params (M)		FLOPs (G)	
	MoE	MLoRE	MoE	MLoRE
5 experts, [1×1, 1×1]	3.1	1.2	3.00	1.49
10 experts, [1×1, 1×1]	4.7	1.6	4.49	1.58
15 experts, [1×1, 1×1]	6.3	1.9	5.99	1.66
5 experts, [3×3, 1×1]	14.9	3.4	14.24	7.12
10 experts, [3×3, 1×1]	22.4	4.7	21.37	7.21
15 experts, [3×3, 1×1]	29.9	6.0	28.49	7.29

learning dynamic and automatic manner of task specialization and cooperations [10]. They utilize MoE to design the encoder blocks and dynamically select network paths for different tasks and different samples. However, compared to these encoder-focused methods, the utilization of MoE in the decoder is less studied, where only Ye *et al.* [50] first applied MoE to the decoder recently. It decodes task-specific features by dynamically assembling task-generic features from different experts and achieves superior performance than previous decoder-focused methods. This motivates us to delve deep into the MoE-based MTL decoder.

Benefiting from the dynamic routing process, these MoE-based methods can significantly improve the diversity in both parameters and features, resulting in more discriminative task-specific features. However, such a paradigm still has some limitations. First, although MoE-based methods can build connections in a subset of tasks by sharing the same expert in the dynamic routing process, the chances of expert sharing among all the tasks are very low, potentially hindering the router from building global relationships among all tasks. Nevertheless, the global relationship modeling among all tasks has been proven useful for dense MTL in [48, 49]. Based on this fact, we argue it is essential to explicitly model global relationship among all tasks in MoE. Furthermore, the capacity of the task-generic feature spaces is highly related to the number of experts. As empirically demonstrated in [11, 15], increasing the number of experts has the potential to facilitate multi-task learning in a unified network. However, for existing dense multi-task learning, adding more expert networks will introduce more parameters and computational costs, which is a heavy burden for the whole model.

To address the above issues, we propose a novel decoder-focused approach, which we call it Mixture of Low-Rank Experts (MLoRE). The core idea of the MLoRE framework is to explicitly model global relationships among all tasks in the MoE and free the MoE from heavily dense computation burdens when increasing the number of experts to enlarge the task-generic feature spaces and context. To address the

first issue, MLoRE builds upon the basic MoE structure and introduces a task-sharing generic convolution path that parallels with the MoE. Specifically, the backbone features are first projected to different task features and then all of them are fed into the generic convolution path and the original MoE’s expert networks. By simply sharing the same generic path among all tasks, different tasks can be globally correlated with each other. Furthermore, to enhance the discrimination of task-specific features, we exclude some experts from the dynamic routing process and enable them to serve specific tasks.

To increase the number of experts while not bringing too many parameters and FLOPs, we take inspiration from LoRA that the basic models adapted to different tasks only needs low-rank weight updates. To be specific, we transform the expert networks of the MoE to the different low-rank formats of a vanilla convolution, which saves more than 60% of the parameters compared to the standard MoE module, as shown in Tab. 1. In addition, to control the computational cost brought by the increasing number of expert networks, we do not use any non-linear activation functions in all the expert networks and the generic convolution path, which enables re-parameterization during inference. Through re-parameterization, the knowledge of experts can be injected into the generic convolution path, reducing the computational cost for dense tasks. To our knowledge, we are the first to use linear experts in MoE for multi-task dense prediction. To verify the effectiveness of our method, we conduct comprehensive experiments on the PASCAL-Context and NYUD-v2 datasets. Our method achieves new performance records on all tasks as shown in Fig. 1.

In summary, our contributions are three-fold:

- We analyze the issues of MoE when applied in multi-task learning and propose a novel decoder-focused framework, MLoRE, which can explicitly model global relationships among all tasks and enlarge the capacity of feature representations without increasing the model size too much.
- We introduce a simple task-sharing generic path to the MoE structure and propose linear and low-rank expert networks inspired by LoRA. The generic convolution path and low-rank expert paths can be linearly combined, enabling re-parameterization at inference.
- Experiments on PASCAL-Context and NYUDv2 show that the proposed method clearly outperforms previous state-of-the-art MTL methods on all tasks.

2. Related Work

2.1. Dense Multi-Task Learning

In computer vision, multi-task learning (MTL) for dense prediction tasks has been widely studied. The previous methods can be divided into two categories, including optimization-based methods and architecture-based meth-

ods [44]. Optimization methods [8, 9, 17, 27, 55] facilitate MTL by utilizing different strategies to balance the influence of each task. Architecture-based methods aim to design a unified deep network for MTL. They can be further classified into two categories, encoder-focused methods, and decoder-focused methods. (i) Encoder-focused methods [3, 18, 33, 38] design multi-task backbones to extract features adapted to different tasks. Typical methods include cross-stitch networks [35], neural discriminative dimensionality reduction networks [16], multi-task attention networks [30], branched networks [43], and mixture-of-expert networks [10, 14]. (ii) Decoder-focused methods [4, 45, 46, 48–50, 53, 54, 56] share the same backbone and design delicate heads to extract task-specific features for each task and cross-task relationships. The advantage of decoder-focused methods lies in that they can benefit from powerful off-the-shelf vision backbones, such as DINOv2 [36]. Our method also falls into the decoder-focused category and studies how to produce task-specific features with the MoE technique.

2.2. Mixture-of-Experts

Mixture-of-Experts (MoE) [24, 25] learns multiple expert networks and a router network that controls the probability of each expert contributing to the final output. This technique is also used in multi-task learning, which can better adapt to the data diversity. Different expert networks learn different discriminative features. The router network learns hard/soft task-specific coefficients to dynamically assemble discriminative features for each task. Prior MoE-based multi-task methods [6, 10, 14] are mainly encoder-focused methods. They introduce MoE into the backbone blocks to sparsely activate different paths for different tasks in the inference stage. Recently, Ye *et al.* [50] first introduced the MoE technique to the decoder. They utilize the spatial context-aware gates to combine each pixel of the features from different expert networks. The above methods decompose the backbone feature into multiple generic feature spaces and assemble discriminative task-specific features from them. Unlike the above MoE-based MTL method, our method first builds global relationships among all tasks explicitly in the MoE structure rather than leaving this work implicitly done by the task-specific routers. Moreover, our proposed low-rank experts give MoE better efficiency compared to the naive MoE and this gap gradually gets larger as the number of experts increases.

2.3. Low-Rank Structure

The low-rank structure is often used in deep learning for its efficiency [22, 40, 42, 47]. Recently, many methods in parameter-efficient adaptation [21, 26, 31, 41], such as LoRA [21], utilize the low-rank structure and have shown impressive results. LoRA takes inspiration from Aghajanyan *et*

al. [1] that the difference in weights between the pre-trained model and the adapted model resides on low intrinsic rank. It learns an extra low-rank matrix instead of tuning the whole layer for adaptation. More related to our work, earlier MTL methods [40, 47] utilize low-rank structure to model task-generic features and generate task-specific features by linear combination. Unlike them, our method utilizes the low-rank structure to control the computation budget when increasing the number of experts in MoE.

3. Method

3.1. Overall Framework

The overall framework follows the multi-scale architecture as in previous works [49, 50]. We utilize an off-the-shelf vision transformer (ViT) as the encoder and collect multi-scale features from different layers. Formally, given an input image \mathbf{I} and a vision transformer \mathcal{F} , we can attain multi-scale feature sets $\{\mathbf{X}^l = \mathcal{F}^l(\mathbf{I})\}$ from different layers, where $\mathbf{X}^l \in \mathbb{R}^{L \times C}$. Here, l , $L = H \times W$, and C denote the layer index, the number of patches, and the feature dimension, respectively. $\mathcal{F}^l(\mathbf{I})$ is the output features of the l -th transformer layer. The multi-scale features are fed into the decoder which includes two stacked MLoRE modules at each scale. For each task, the output features of MLoRE from different scales are concatenated together to generate the final task features for dense predictions.

3.2. Preliminaries: Mixture of Experts

Before describing the details of the proposed Mixture-of-Low-Rank-Experts (MLoRE) module, we first introduce the basic form of MoE $f_{moe}(\cdot)$. Formally, suppose MoE contains N experts and T router networks, denoted as $\mathbb{E} = \{E_1, E_2, \dots, E_N\}$ and $\mathbb{G} = \{G_1, G_2, \dots, G_T\}$, respectively. N and T are the number of experts and the number of tasks, respectively. The backbone feature \mathbf{X}^l from the l -th layer is fed into the networks of N experts and T router networks, respectively. For convenience, the superscript l is omitted in the following. For the n -th expert, the discriminative output feature is generated by $\mathbf{X}_n = E_n(\mathbf{X})$. In the meanwhile, MoE learns gates from the task-specific router networks for different tasks. For task t , the gate value for each expert generated by the router network can be represented as $\mathbf{g}^t = G_t(\mathbf{X})$, where $\mathbf{g}^t \in \mathbb{R}^N$. Finally, MoE generates task-specific feature \mathbf{S}^t for task t by utilizing gates to combine the expert features, which can be formulated as

$$\mathbf{S}^t = f_{moe}(\mathbf{X}) = \sum_{n=1}^N \mathbf{g}_n^t \mathbf{X}_n. \quad (1)$$

For M³ViT [14] and Mod-Squad [10], they deactivate some experts according to the corresponding gate values for one inference and select the top- k experts. The task-specific features are used to make predictions for each task.




Figure 2. Overall framework of the proposed method. The MLoRE modules are equipped at different layers, where the backbone features from different layers are fed into the MLoRE modules, respectively. At each selected layer, the backbone feature is first projected to different task features and then sent to the task-sharing convolution, task-sharing low-rank expert networks followed by the task-specific router network and task-specific low-rank expert networks. The outputs of these branches are accumulated to generate task-specific features. At each selected layer, we stack two MLoRE modules.

The advantages of MoE lie in that it enables the dynamic encoding of features for each sample and each task and increases feature encoding diversity via multiple experts. However, when applying the MoE technique to build an MTL decoder, we find that it struggles with building global task relationships. Moreover, when increasing the number of experts to enlarge the capacity of feature representations and the context of the expert network, the parameters and computational cost increase accordingly. To address the above issues, we propose the Mixture-of-Low-Rank-Expert (MLoRE) module.

3.3. Mixture of Low-Rank Experts

The overall pipeline of the proposed Mixture-of-Low-Rank-Expert (MLoRE) module is illustrated in Fig. 2. Considering building cross-task relations across the task-specific features, we first project the backbone features into task features using several lightweight convolutional layers. Then, the feature of each task is delivered to the task-sharing generic path and multiple task-sharing expert networks with different low-rank convolutions. The low-rank experts are selected according to the predictions of the task-specific router network for each task. Furthermore, except for building task-specific features based on task-specific routers, we introduce additional task-specific low-rank expert networks to assist in building more discriminative task-specific features. The features from the task-sharing generic path, task-sharing low-rank expert networks selected by task-specific routers and the task-specific expert networks are summed up to generate discriminative task-specific features. We introduce linearity into the MLoRE module and do not utilize any activation function in all paths, enabling reparameterization to reduce computational cost.

To be specific, the backbone feature \mathbf{X} at the l -th layer is first projected to multiple task splits corresponding to each task using 1×1 convolutions, which can be denoted as $\{\mathbf{X}^t = f_{t,1 \times 1}(\mathbf{X}), t \in [1, \dots, T]\}$, where $\mathbf{X}^t \in \mathbb{R}^{C \times H \times W}$. Then, the task feature is sent to three paths, *i.e.*, task-sharing generic path $f_g(\cdot)$, task-sharing low-rank expert path $f_{lre}(\cdot)$ with task-specific router network $f_{sr}^t(\cdot)$, and task-specific low-rank expert path $f_{se}(\cdot)$. The output task-specific feature \mathbf{S}^t is obtained by

$$\mathbf{S}^t = f_g(\mathbf{X}^t) + f_{sr}^t(f_{lre}(\mathbf{X}^t)) + f_{se}^t(\mathbf{X}^t). \quad (2)$$

In the following, we introduce the network details of all these paths in the MLoRE module.

Task-sharing generic path contains a simple 3×3 convolutional layer with a weight matrix $\mathbf{W}_g \in \mathbb{R}^{3 \times 3 \times C \times C}$ and a bias matrix $\mathbf{b}_g \in \mathbb{R}^C$. As all task features will go through this generic convolution, it will be optimized by the gradients of different tasks simultaneously, which can help extract common features among all tasks. During the training process, we stop the gradients of this path for further back-propagation. The gradient is back-propagated through the other two paths. We found such a simple operation can better ease the optimization process and well solve the gradient conflicts. Experimental results in Sec. 4.2 show that the simple task-sharing generic path can bring performance improvements on all tasks, demonstrating the effectiveness of the idea of explicitly building the relationship across all tasks from a global perspective.

Task-sharing low-rank expert path. We take inspiration from LoRA [21] and adopt low-rank convolution which is a low-rank format of a vanilla convolution. Each task-sharing expert network shares a similar structure consisting of a

3×3 convolution and a 1×1 convolution. The weights and bias of all task-sharing expert networks can be formulated as $\{\mathbf{W}_{lreb}^n, \mathbf{b}_{lreb}^n, \mathbf{W}_{lrea}^n, \mathbf{b}_{lrea}^n | n \in [1, \dots, N]\}$, where $\mathbf{W}_{lreb}^n \in \mathbb{R}^{3 \times 3 \times C \times r_n}$, $\mathbf{b}_{lreb}^n \in \mathbb{R}^{r_n}$, $\mathbf{W}_{lrea}^n \in \mathbb{R}^{1 \times 1 \times r_n \times C}$, and $\mathbf{b}_{lrea}^n \in \mathbb{R}^C$ ($r_n \ll C$). r_n denotes the rank for the n -th expert network. In our method, r_n for different expert networks is different, aiming to improve the diversity of parameters and features. For each task, the task-specific router network $f_{sr}^t(\cdot)$ learns gate values for these experts and activate the top- k experts according to the gate values. The output features of all the activated experts are summed up and then sent to a BatchNorm layer to generate task-specific features. The BatchNorm layer contains four parameters, including the accumulated channel-wise mean $\mu \in \mathbb{R}^C$, the accumulated channel-wise standard deviation $\sigma \in \mathbb{R}^C$, the scaling factor $\gamma \in \mathbb{R}^C$ and the $\beta \in \mathbb{R}^C$, respectively.

Task-specific low-rank expert path includes T specific expert networks and each processes one task feature. For each specific expert network, we utilize a similar structure with task-sharing expert path, which contains a 3×3 convolution with a weight matrix $\mathbf{W}_{seb}^t \in \mathbb{R}^{3 \times 3 \times C \times R}$ and a bias matrix $\mathbf{b}_{seb}^t \in \mathbb{R}^R$, followed by a 1×1 convolution with a weight matrix $\mathbf{W}_{sea}^t \in \mathbb{R}^{1 \times 1 \times R \times C}$ and a bias matrix $\mathbf{b}_{sea}^t \in \mathbb{R}^C$. R denotes the rank number ($R \ll C$). The task-specific expert path can enhance the distinctiveness of the task-specific features, which will be verified in the experiments.

Router network. As shown in Fig. 2, to generate task-specific features from the task-sharing low-rank expert path, we learn task-specific router networks to generate gate values for each expert and utilize them as the weight of the linear combination of feature output from different experts. The router network for each task is usually the simple linear layers followed by an average pooling layer and a prediction layer. Specifically, the router network $f_{sr}^t(\cdot)$ for the t -th task is designed as follows. Our router network takes the task-specific features $\mathbf{X}^t \in \mathbb{R}^{C \times H \times W}$ as input and feed them into two consecutive 1×1 convolutions, mapping the channel dimension from C to $C/4$, followed by a global pooling layer. The output is a global feature vector $\mathbf{X}_f \in \mathbb{R}^{\frac{C}{4}}$.

In addition, inspired by previous works [19, 20] showing that positional information is also important for modeling long-range spatial context, we introduce another parallel position-aware branch. Similarly, it consists of two linear layers. The first linear layer shrinks the feature along the spatial dimension, mapping the shape from $\mathbb{R}^{C \times HW}$ to $\mathbb{R}^{C \times 1}$, which will then be transformed to $\mathbb{R}^{\frac{C}{4}}$ via the second linear layer. The output feature vectors of these two branches are concatenated along the final dimension and then sent to the final prediction layer followed by a Softmax function to produce the gate values g_t for each expert.

Re-parameterization during inference. We introduce linearity into the MLoRE module by removing all activation

functions, enabling the parameters of all paths to reparameterize to a simple 3×3 convolution for each task at inference. We first parameterize the task-sharing low-rank expert path and then parameterize the parameters of all paths. According to [12], the weight and bias matrices in the task-sharing expert path can be combined and formulated as :

$$\mathbf{W}_{lre}^t = \mathfrak{B}\left(\frac{\gamma}{\sigma}\right) \sum_{k \in \mathbb{K}_t} \mathbf{g}_k^t \mathbf{W}_{lreb}^k \mathbf{W}_{lrea}^k, \quad (3)$$

$$\mathbf{b}_{lre}^t = \frac{\gamma}{\sigma} \left(\sum_{k \in \mathbb{K}_t} \mathbf{g}_k^t (\mathbf{b}_{lreb}^k \mathbf{W}_{lrea}^k + \mathbf{b}_{lrea}^k) - \mu \right) + \beta, \quad (4)$$

where \mathfrak{B} denotes the broadcast operation and \mathbb{K}_t denotes the index set of the activated experts selected by the router network for task t . \mathbf{g}_k^t is the k -th gate value predicted by the router. At the inference time, the weight matrix and bias matrix of all these three paths can be re-parameterized as

$$\mathbf{W}_r^t = \mathbf{W}_g + \mathbf{W}_{sr}^t + \mathbf{W}_{seb}^t \mathbf{W}_{sea}^t, \quad (5)$$

$$\mathbf{b}_r^t = \mathbf{b}_g + \mathbf{b}_{sr}^t + \mathbf{b}_{seb}^t \mathbf{W}_{sea}^t + \mathbf{b}_{sea}^t. \quad (6)$$

\mathbf{W}_r^t and \mathbf{b}_r^t are the weight and bias of the re-parameterized convolution. Thus, Eqn. (2) can be reformulated as

$$\mathbf{S}^t = \mathbf{X}^t \circledast \mathbf{W}_r^t + \mathfrak{B}(\mathbf{b}_r^t), \quad (7)$$

where \circledast denotes the convolution operation and \mathbf{b}_r^t is with the same shape as \mathbf{X}^t via broadcast.

4. Experiments

4.1. Experimental Settings

Datasets. To demonstrate the effectiveness of our method, we evaluate the performance of our method on two popular multi-task datasets, including PASCAL-Context [7] and NYUD-v2 [39]. PASCAL-Context [7] contains high-quality annotations of several tasks, including semantic segmentation, human parsing, saliency detection, surface normals, and object boundary detection. There are 4,998 training images and 5,105 test images in this dataset. NYUDv2 [39] also provides high-quality multi-task annotations, including semantic segmentation, monocular depth estimation, surface normals, and object boundary detection. This dataset contains 795 training images and 654 test images.

Evaluation metrics. We introduce the evaluation metrics for the tasks mentioned above. Following previous multi-task works [48, 49], the mean intersection-over-union (mIoU) is used to evaluate semantic segmentation and human parsing. The root mean square error (RMSE) is used to evaluate the accuracy of monocular depth estimation. Saliency detection uses the maximum F-measure (maxF). Surface normal and object boundary detection adopt the mean error (mErr) and the optimal-dataset-scale F-measure

Table 2. Ablation study on different components in MLoRE on the PASCAL-Context dataset. Each row adds an extra setting to the above row. MoE: the standard mixture-of-experts structure; LoRE: task-sharing low-rank expert path; GC: task-sharing generic convolution path; SPE: task-specific expert path. \uparrow denotes higher is better. \downarrow denotes lower is better.

Settings	Semseg	Parsing	Sal.	Normal	Bound.	MTL	FLOPs	#Param
	mIoU \uparrow	mIoU \uparrow	maxF \uparrow	mErr \downarrow	odsF \uparrow	Δ_m \uparrow	(G)	(M)
Baseline	77.38	65.15	85.08	13.79	69.87	-3.41	391	115
+ MoE	78.56	66.78	85.18	13.57	73.91	-1.20	1834	676
Baseline								
+ LoRE	78.38	66.21	85.15	13.71	73.53	-1.71	568	213
+ GC	79.25	67.43	85.20	13.70	74.38	-0.88	568	243
+ SPE	79.26	67.82	85.31	13.65	74.69	-0.58	568	259

(odsF) as the evaluation metrics, respectively. In total, we evaluate the MTL gain Δ_m across all tasks, following [34].

Training settings. We utilize ViT-large [13] as the backbone. The channel number of the decoder is set to 384. For the ablation study, the ViT-base network is set as the backbone. Following the previous work [49], the proposed MTL network is trained for 40,000 iterations with a batch size of 4 on both the two datasets. The optimizer and the loss functions for different tasks follow the previous work [49].

4.2. Ablation Study

In this subsection, we conduct extensive experiments to demonstrate the effectiveness of different components and find the best settings of different hyper-parameters. All the ablation experiments are conducted based on the ViT-based backbone if not specified otherwise. The baseline is built upon the ViT-base backbone with 12 layers, where the backbone features from the 3-*rd*, 6-*th*, 9-*th*, and 12-*th* layers are utilized as the multi-scale features, each of which is followed by a linear layer to project the channel dimension to the output channels for each task.

Effectiveness of different components. We first conduct experiments to verify the effectiveness of different components of the MLoRE module. The quantitative results are shown in Tab. 2. We first examine the performance of the baseline with the standard MoE and the parameter size and FLOPs of their model. The expert networks in the standard MoE (15 experts) are similar to ours, each of which consists of a 3×3 convolution and a 1×1 convolution with ReLU among them. When adding the MoE to the baseline, we observe the MTL gain can be improved, but the parameters and FLOPs also expand about 5 times and 4 times, which is a heavy burden for the whole network. When adding the low-rank expert networks (LoRE) to the baseline, the performance is also improved, but the parameters and FLOPs are only 1/3 and 1/3 of the MoE-based model. When the low-rank property is applied to the expert net-




Figure 3. Ablation study on the number of experts N and the number of activated experts K . In the right figure, we also present the parameter change of the MLoRE module with the increase in the number of experts.

Table 3. Ablation on the rank setting in the task-sharing low-rank expert path of MLoRE on PASCAL-Context dataset.

Min/Max Rank	Semseg mIoU \uparrow	Parsing mIoU \uparrow	Saliency maxF \uparrow	Normal mErr \downarrow	Boundary odsF \uparrow	MTL Δ_m \uparrow
16/16	78.84	68.01	85.32	13.69	74.39	-0.77
16/128	79.26	67.82	85.30	13.65	74.69	-0.58
128/128	78.79	66.98	85.35	13.67	74.29	-1.07

works, the parameter size is reduced several times. When further introducing linearity in the expert networks by removing all activation functions, the computations are saved by re-parameterizing all experts to a single convolution.

Furthermore, we also emphasize the importance of explicitly building global task correlations in the MoE and introduce the task-sharing generic path to achieve this goal. It can be seen that adding the task-sharing generic path to LoRE can further improve the performance and outperform the baseline with MoE on most metrics, which proves the effectiveness of modeling the global relationships among all tasks. In addition, adding a task-specific low-rank expert for each task also improves the performance, which demonstrates that the specialized expert can enhance the discrimination of task-specific features. We empirically set the rank of the task-specific low-rank expert to 64 in this paper.

Number of task-sharing low-rank experts and top- k selection. We ablate the number of low-rank experts in the MLoRE module and top- k selection of the experts by the task-specific router networks. We first fix one and ablate another parameter to study their impact on multi-task performance. As shown in Fig. 3, when increasing the number of experts, the MTL gain of the model is significantly improved and allows us to achieve the best performance when the number of experts is 15. Further increasing the number of experts, we do not observe obvious performance gain. Besides, when fixing the number of experts to 15, we ablate the ratio of the activated experts. We observe that activating 60% experts for each task is the best choice in our experiments. When selecting all expert networks, the performance decreases by a large margin, which reflects the importance of sparsity for feature discrimination.




Figure 4. (a) The relations between tasks and low-rank experts. (b) The ratio of an expert activated by different numbers of tasks in the MLoRE module without the task-sharing generic path. We can see that without the task-sharing generic path, there is only a few experts can be activated by all five tasks. Horizontal coordinates represent the ranks of different experts.

Rank number setting. The expert networks utilize the low-rank format of a vanilla 3×3 convolution with 640 output channels. The rank r of different expert networks also plays an important role. We study different settings, including 1) all experts with the same rank number 16, 2) all experts with the rank number 128, 3) all experts with the rank numbers from 16 to 128 with an increased step of 8. As shown in Tab. 3, it can be seen that selecting different rank numbers for expert networks achieves the best MTL gain. We analyze the expert networks with different rank numbers that can bring more feature diversity than the same rank numbers, which is more useful for assembling task-specific features, and utilize this as our settings.

The relation visualizations between tasks and the low-rank experts are shown in Fig. 4(a). We count the relations from the second MLoRE module at the last stage and calculate the activation ratio of every expert selected by different tasks on the entire dataset. It can be seen that experts with different ranks tend to learn different subsets of tasks. Specifically, for this module, experts with lower ranks tend to learn shared knowledge for 3-4 related tasks, while experts with higher ranks tend to specialize in 1-2 tasks. Moreover, we also show the ratio of different experts to be activated by different numbers of tasks when the task-sharing generic path is not added to the MLoRE module in Fig. 4(b). It can be seen that in a fully dynamic manner, all these experts are seldom or even never activated by all the tasks in one sample. This proves the fact that almost no expert can learn global relationships across all tasks when utilizing the MoE in the decoder directly. This phenomenon strongly supports the necessity of the task-sharing generic path for explicitly modeling global task relationships.

Task-specific router network. The router network is important for generating task-specific gates, which decide how to activate the experts and assemble their features. We ablate several settings of the router network and the results are shown in Tab. 4. Adding the position-aware branch to the basic router network can improve the MTL gain by +0.17. The position-aware branch can obtain more context information, which benefits for the router network. Moreover, when replacing the router’s input from the learnable pa-

Table 4. Ablation on the setting of the router network. **basic**: the basic router network. **pos.**: the positional-aware router network. **w/o sample-dep**: the input is the learnable parameters.

Router Variant	Semseg mIoU \uparrow	Parsing mIoU \uparrow	Saliency maxF \uparrow	Normal mErr \downarrow	Boundary odsF \uparrow	MTL $\Delta_m \uparrow$
basic	79.15	67.40	85.21	13.58	74.34	-0.75
basic+pos.	79.26	67.82	85.31	13.65	74.69	-0.58
only pos.	79.10	67.76	85.11	13.71	74.51	-0.82
basic	79.15	67.40	85.21	13.58	74.34	-0.75
w/o sample-dep	78.86	67.38	85.41	13.65	74.25	-0.91

rameters to the sample features, the MTL gain increases by +0.16, which demonstrates that dynamic information from samples is vital for gates.

4.3. Comparisons with Other Methods

The quantitative comparisons with previous state-of-the-art (SOTA) methods are shown in Tab. 5 and Tab. 6. It can be seen that our method clearly outperforms previous methods in terms of all metrics on both the PASCAL-Context dataset and the NYUDv2 dataset. In particular, on the PASCAL-Context dataset, for semantic segmentation, human parsing, and boundary detection, the performance of our method outperforms the previous best method by +0.52 mIoU, +1.10 mIoU and +1.92 odsF, respectively.

Previous methods, M³ViT [14], Mod-Squad [10], and TaskExpert [50] all utilize the MoE technique in their network. However, our method shows performance superior to theirs, which demonstrates the effectiveness of the MLoRE module. Compared with the decoder-focused method, TaskExpert, the performance of semantic segmentation, human parsing, and object boundary is significantly improved by +0.77 mIoU, +1.10 mIoU and +2.12 odsF, but using fewer parameters and FLOPs.

Furthermore, we also present an intuitive visualization comparison among different methods in Fig. 5. Our method can generate better visualization results than the previous SoTA methods, especially for semantic segmentation, human parsing, and object boundary detection. More visualization comparisons are in the supplemental material.

Table 5. Quantitative comparison of different methods on PASCAL-Context dataset. * denotes the reproduced performance of methods based on the ViT-large backbone in [50].

Method	Publication	Backbone	Semseg mIoU \uparrow	Parsing mIoU \uparrow	Saliency maxF \uparrow	Normal mErr \downarrow	Boundary odsF \uparrow	FLOPs (G)	#Param (M)
PAD-Net [46]	CVPR'18	HRNet18	53.60	59.60	65.80	15.30	72.50	124	81
MTI-Net [45]	ECCV'20	HRNet18	61.70	60.18	84.78	14.23	70.80	161	128
ATRC [4]	ICCV'21	HRNet18	67.67	62.93	82.29	14.24	72.42	216	96
PAD-Net* [46]	CVPR'18	ViT-large	78.01	67.12	79.21	14.37	72.60	773	330
MTI-Net* [45]	ECCV'20	ViT-large	78.31	67.40	84.75	14.67	73.00	774	851
ATRC* [4]	ICCV'21	ViT-large	77.11	66.84	81.20	14.23	72.10	871	340
InvPT [48]	ECCV'22	ViT-large	79.03	67.61	84.81	14.15	73.00	669	423
TaskPrompter [49]	ICLR'23	ViT-large	80.89	68.89	84.83	13.72	73.50	497	401
TaskExpert [50]	ICCV'23	ViT-large	80.64	69.42	84.87	13.56	73.30	622	420
Ours	-	ViT-large	81.41	70.52	84.90	13.51	75.42	571	407

Table 6. Quantitative comparison of different methods on the NYUD-v2 dataset. Our method performs the best on all four tasks.

Method	Backbone	Semseg mIoU \uparrow	Depth RMSE \downarrow	Normal mErr \downarrow	Boundary odsF \uparrow
PAD-Net [46]	HRNet18	36.61	0.6246	20.88	76.38
MTI-Net [45]	HRNet48	45.97	0.5365	20.27	77.86
ATRC [4]	HRNet48	46.33	0.5363	20.18	77.94
InvPT [48]	ViT-large	53.56	0.5183	19.04	78.10
TaskPrompter [49]	ViT-large	55.30	0.5152	18.47	78.20
TaskExpert [50]	ViT-large	55.35	0.5157	18.54	78.40
Ours	ViT-large	55.96	0.5076	18.33	78.43

4.4. Efficient MTL Models

We also apply our MLoRE module to the ViT-small backbone to check the performance of the efficient models. Specifically, the channel number of the decoder decreases from 384 to 192. As shown in Tab. 7, using about 35% GFLOPs of TaskExpert, our method can achieve a highly competitive result. In particular, the performance of the semantic segmentation and object boundary is improved by 0.6% mIoU and 1.01% odsF while the metrics of other tasks are close to TaskExpert. Furthermore, the parameters are fewer than TaskExpert by 11M.

4.5. Conclusions

We present a novel decoder-focused multi-task learning method, called MLoRE. We delve deep into the standard mixture-of-experts (MoE) technique and improve it for dense multi-task learning from two aspects. First, to address the neglected global relationship modeling in the MoE, we add a simple generic convolution path to MoE, enabling different task features to be able to share this path. Furthermore, we apply the low-rank format of a vanilla convolution to different expert networks to free the MoE from high computational cost and a large number of parameters when increasing the number of experts. Experiments




Figure 5. Qualitative comparison among different methods, including InvPT [48], TaskPrompter [49], and ours. Best viewed with zoom-in. It can be seen that our method achieves better visual results than other methods on all five tasks thanks to the proposed MLoRE module.

Table 7. Quantitative comparison of the MoE-based efficient models on PASCAL-Context dataset.

Method	Semseg mIoU \uparrow	Parsing mIoU \uparrow	Sal. maxF \uparrow	Nor. mErr \downarrow	Bound. odsF \uparrow	FLOPs (G)	#Param (M)
M ³ ViT	72.80	62.10	66.30	14.50	71.70	420	42
Mod-Squad	74.10	62.70	66.90	13.70	72.00	420	52
TaskExpert	75.04	62.68	84.68	14.22	68.80	204	55
ours	75.64	62.65	84.70	14.43	69.81	72	44

demonstrate that the proposed method clearly outperforms previous state-of-the-art methods on all metrics.

Acknowledgments. This research was supported by NSFC (NO. 62276145), the Fundamental Research Funds for the Central Universities (Nankai University, 070-63223049), CAST through Young Elite Scientist Sponsorship Program (No. YESS20210377). Computations were supported by the Supercomputing Center of Nankai University (NKSC).

References

- [1] Armen Aghajanyan, Luke Zettlemoyer, and Sonal Gupta. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. *arXiv preprint arXiv:2012.13255*, 2020. 3
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4009–4018, 2021. 1
- [3] David Bruggemann, Menelaos Kanakis, Stamatios Georgoulis, and Luc Van Gool. Automated search for resource-efficient branched multi-task networks. *arXiv preprint arXiv:2008.10292*, 2020. 3
- [4] David Bruggemann, Menelaos Kanakis, Anton Obukhov, Stamatios Georgoulis, and Luc Van Gool. Exploring relational context for multi-task dense prediction. In *Int. Conf. Comput. Vis.*, pages 15869–15878, 2021. 1, 3, 8
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017. 1
- [6] Tianlong Chen, Xuxi Chen, Xianzhi Du, Abdullah Rashwan, Fan Yang, Huizhong Chen, Zhangyang Wang, and Yeqing Li. Adamv-moe: Adaptive multi-task vision mixture-of-experts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17346–17357, 2023. 1, 3
- [7] Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, and Alan Yuille. Detect what you can: Detecting and representing objects using holistic models and body parts. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1971–1978, 2014. 5
- [8] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International conference on machine learning*, pages 794–803. PMLR, 2018. 3
- [9] Zhao Chen, Jiquan Ngiam, Yanping Huang, Thang Luong, Henrik Kretzschmar, Yuning Chai, and Dragomir Anguelov. Just pick a sign: Optimizing deep multitask models with gradient sign dropout. *Advances in Neural Information Processing Systems*, 33:2039–2050, 2020. 3
- [10] Zitian Chen, Yikang Shen, Mingyu Ding, Zhenfang Chen, Hengshuang Zhao, Erik G Learned-Miller, and Chuang Gan. Mod-squad: Designing mixtures of experts as modular multi-task learners. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11828–11837, 2023. 1, 2, 3, 7
- [11] Aidan Clark, Diego De Las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake Hechtman, Trevor Cai, Sebastian Borgeaud, et al. Unified scaling laws for routed language models. In *Int. Conf. Mach. Learn.*, pages 4057–4086. PMLR, 2022. 2
- [12] Xiaohan Ding, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Diverse branch block: Building a convolution as an inception-like unit. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10886–10895, 2021. 5
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Int. Conf. Learn. Represent.*, 2020. 6
- [14] Zhiwen Fan, Rishov Sarkar, Ziyu Jiang, Tianlong Chen, Kai Zou, Yu Cheng, Cong Hao, Zhangyang Wang, et al. M³vit: Mixture-of-experts vision transformer for efficient multi-task learning with model-accelerator co-design. In *Adv. Neural Inform. Process. Syst.*, volume 35, pages 28441–28457, 2022. 1, 3, 7
- [15] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1):5232–5270, 2022. 2
- [16] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3205–3214, 2019. 1, 3
- [17] Michelle Guo, Albert Haque, De-An Huang, Serena Yeung, and Li Fei-Fei. Dynamic task prioritization for multitask learning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 270–287, 2018. 3
- [18] Pengsheng Guo, Chen-Yu Lee, and Daniel Ulbricht. Learning to branch for multi-task learning. In *Int. Conf. Mach. Learn.*, pages 3854–3863. PMLR, 2020. 3
- [19] Qibin Hou, Li Zhang, Ming-Ming Cheng, and Jiashi Feng. Strip pooling: Rethinking spatial pooling for scene parsing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4003–4012, 2020. 5
- [20] Qibin Hou, Daquan Zhou, and Jiashi Feng. Coordinate attention for efficient mobile network design. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 13713–13722, 2021. 5
- [21] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 3, 4
- [22] Yerlan Idelbayev and Miguel A Carreira-Perpinán. Low-rank compression of neural nets: Learning the rank of each layer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8049–8059, 2020. 3
- [23] Keishi Ishihara, Anssi Kanervisto, Jun Miura, and Ville Hautamaki. Multi-task learning with attention for end-to-end autonomous driving. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2902–2911, 2021. 1
- [24] Robert A Jacobs and Michael I Jordan. Learning piecewise control strategies in a modular neural network architecture. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(2):337–345, 1993. 3
- [25] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural Comput.*, 3(1):79–87, 1991. 3
- [26] Shibo Jie and Zhi-Hong Deng. Fact: Factor-tuning for lightweight adaptation on vision transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 1060–1068, 2023. 3

- [27] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018. 3
- [28] Zhihao Li, Toshiyuki Motoyoshi, Kazuma Sasaki, Tetsuya Ogata, and Shigeki Sugano. Rethinking self-driving: Multi-task knowledge for better generalization and accident explanation ability. *arXiv preprint arXiv:1809.11100*, 2018. 1
- [29] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1925–1934, 2017. 1
- [30] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1871–1880, 2019. 3
- [31] Yen-Cheng Liu, Chih-Yao Ma, Junjiao Tian, Zijian He, and Zsolt Kira. Polyhistor: Parameter-efficient multi-task adaptation for dense vision tasks. *Advances in Neural Information Processing Systems*, 35:36889–36901, 2022. 3
- [32] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3431–3440, 2015. 1
- [33] Yongxi Lu, Abhishek Kumar, Shuangfei Zhai, Yu Cheng, Tara Javidi, and Rogerio Feris. Fully-adaptive feature sharing in multi-task networks with applications in person attribute classification. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5334–5343, 2017. 1, 3
- [34] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive single-tasking of multiple tasks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1851–1860, 2019. 6
- [35] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3994–4003, 2016. 1, 3
- [36] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3
- [37] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Int. Conf. Comput. Vis.*, pages 12179–12188, 2021. 1
- [38] Sebastian Ruder, Joachim Bingel, Isabelle Augenstein, and Anders Søgaard. Latent multi-task architecture learning. In *AAAI Conf. Artif. Intell.*, volume 33, pages 4822–4829, 2019. 3
- [39] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *Eur. Conf. Comput. Vis.*, pages 746–760. Springer, 2012. 5
- [40] Chi Su, Fan Yang, Shiliang Zhang, Qi Tian, Larry S Davis, and Wen Gao. Multi-task learning with low rank attribute embedding for person re-identification. In *Proceedings of the IEEE international conference on computer vision*, pages 3739–3747, 2015. 3
- [41] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. VI-adapter: Parameter-efficient transfer learning for vision-and-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5227–5237, 2022. 3
- [42] Madeleine Udell, Corinne Horn, Reza Zadeh, Stephen Boyd, et al. Generalized low rank models. *Foundations and Trends® in Machine Learning*, 9(1):1–118, 2016. 3
- [43] Simon Vandenhende, Stamatios Georgoulis, Bert De Brabandere, and Luc Van Gool. Branched multi-task networks: deciding what layers to share. In *Brit. Mach. Vis. Conf.*, 2019. 3
- [44] 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 Trans. Pattern Anal. Mach. Intell.*, 44(7):3614–3633, 2021. 3
- [45] Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Mti-net: Multi-scale task interaction networks for multi-task learning. In *Eur. Conf. Comput. Vis.*, pages 527–543. Springer, 2020. 1, 3, 8
- [46] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 675–684, 2018. 1, 3, 8
- [47] Yongxin Yang and Timothy M Hospedales. Trace norm regularised deep multi-task learning. *arXiv preprint arXiv:1606.04038*, 2016. 3
- [48] Hanrong Ye and Dan Xu. Inverted pyramid multi-task transformer for dense scene understanding. In *Eur. Conf. Comput. Vis.*, pages 514–530. Springer, 2022. 1, 2, 3, 5, 8
- [49] Hanrong Ye and Dan Xu. Taskprompter: Spatial-channel multi-task prompting for dense scene understanding. In *Int. Conf. Learn. Represent.*, 2022. 2, 3, 5, 6, 8
- [50] Hanrong Ye and Dan Xu. Taskexpert: Dynamically assembling multi-task representations with memorial mixture-of-experts. In *Int. Conf. Comput. Vis.*, 2023. 2, 3, 7, 8
- [51] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Eur. Conf. Comput. Vis.*, pages 325–341, 2018. 1
- [52] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2636–2645, 2020. 1
- [53] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Zequn Jie, Xiang Li, and Jian Yang. Joint task-recursive learning for semantic segmentation and depth estimation. In *Eur. Conf. Comput. Vis.*, pages 235–251, 2018. 3
- [54] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Yan Yan, Nicu Sebe, and Jian Yang. Pattern-affinitive propagation across depth, surface normal and semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4106–4115, 2019. 3
- [55] Xiangyun Zhao, Haoxiang Li, Xiaohui Shen, Xiaodan Liang, and Ying Wu. A modulation module for multi-task learning

with applications in image retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 401–416, 2018. 3

- [56] Ling Zhou, Zhen Cui, Chunyan Xu, Zhenyu Zhang, Chaoqun Wang, Tong Zhang, and Jian Yang. Pattern-structure diffusion for multi-task learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4514–4523, 2020. 3