

# Multi-task Learning for Real-time Autonomous Driving Leveraging Task-adaptive Attention Generator

Wonhyeok Choi<sup>1\*</sup>, Mingyu Shin<sup>2\*</sup>, Hyukzae Lee<sup>3</sup>, Jaehoon Cho<sup>3</sup>, Jaehyeon Park<sup>3</sup>, and Sunghoon Im<sup>1,2†</sup>

**Abstract**—Real-time processing is crucial in autonomous driving systems due to the imperative of instantaneous decision-making and rapid response. In real-world scenarios, autonomous vehicles are continuously tasked with interpreting their surroundings, analyzing intricate sensor data, and making decisions within split seconds to ensure safety through numerous computer vision tasks. In this paper, we present a new real-time multi-task network adept at three vital autonomous driving tasks: monocular 3D object detection, semantic segmentation, and dense depth estimation. To counter the challenge of negative transfer — the prevalent issue in multi-task learning — we introduce a task-adaptive attention generator. This generator is designed to automatically discern interrelations across the three tasks and arrange the task-sharing pattern, all while leveraging the efficiency of the hard-parameter sharing approach. To the best of our knowledge, the proposed model is pioneering in its capability to concurrently handle multiple tasks, notably 3D object detection, while maintaining real-time processing speeds. Our rigorously optimized network, when tested on the Cityscapes-3D datasets, consistently outperforms various baseline models. Moreover, an in-depth ablation study substantiates the efficacy of the methodologies integrated into our framework.

**Index Terms**—Autonomous Driving, Real-time Multi-task Learning, Deep learning for Visual Perception.

## I. INTRODUCTION

With the rise of deep learning technologies, autonomous driving systems have also witnessed remarkable progress in alignment with this development because these perception tasks play a vital role in practical autonomous driving. Given the inherent nature of practical autonomous driving applications, which require swift and accurate responses in unpredictable environments, the algorithmic speed of deployed models (*e.g.*, latency, inference speed) is the crucial factor. Hence, training the model individually for each required task (*e.g.*, semantic segmentation or depth estimation) in this system is pragmatically inefficient for these scenarios.

In such scenarios, an applicable learning scheme is Multi-task learning [1], which trains multiple tasks within a single model. This method can improve the generalization ability by consolidating information from various related tasks into a




Fig. 1. Speed-performance trade-off curve for the Cityscapes-3D validation set (the top-right quadrant is optimal). The Y-axis represents the average relative performance compared to the single-task setup.

joint feature space. However, in the context of heterogeneous Multi-Task Learning (MTL), where tasks consisting of different methods are jointly trained, a challenge arises known as the “negative transfer problem” [2], [3]. This problem can lead to a decrease in overall performance, particularly when dealing with tasks that have low relevance to each other.

Recent works such as soft-parameter sharing [4], [5] and Nas-style MTL [6], [7], [8], [9], [10], [11] attempt to mitigate this problem by allowing the model to have separate sets of parameters while encouraging some level of parameter similarity or correlation across tasks. These approaches are particularly useful, as they can mitigate the risk of the negative transfer problem that can occur with hard-parameter sharing. However, most of these methods use additional network layers or parameters in proportion to the number of tasks, making the time/memory complexity almost equivalent to training multiple single-task models.

To design the real-time multi-task learning framework for autonomous driving, we first adopt the hard-parameter sharing scheme [12], [13], [14] (a.k.a. shared-bottom), where the layers of the model are explicitly shared across all tasks. These methods are suffered from negative transfer problems a lot, so we attempt to strike the balance between leveraging shared knowledge across tasks and allowing flexibility for task-adaptive learning while gaining computational efficiency. We consider two components to achieve this objective: 1. *What is the feasible multi-task architecture that can perform important tasks in autonomous driving?*, 2. *How to generate task-adaptive features while preserving the computational efficiency of hard-parameter sharing?*

<sup>1</sup>W Choi and S Im are with the Department of Electrical Engineering & Computer Science, DGIST, Daegu, Korea, {smu06117, sunghoonim}@dgist.ac.kr

<sup>2</sup>M Shin and S Im are with the Department of Interdisciplinary Studies of Artificial Intelligence, DGIST, Daegu, Korea, {alsrb4446, sunghoonim}@dgist.ac.kr

<sup>3</sup>H Lee, J Cho, and J Park are with the Autonomous Driving Center, Hyundai Motor Company, Seoul, Korea, {hyukzae.lee, jh.cho, hybrid}@hyundai.com

\* Equal-contribution, † Corresponding author

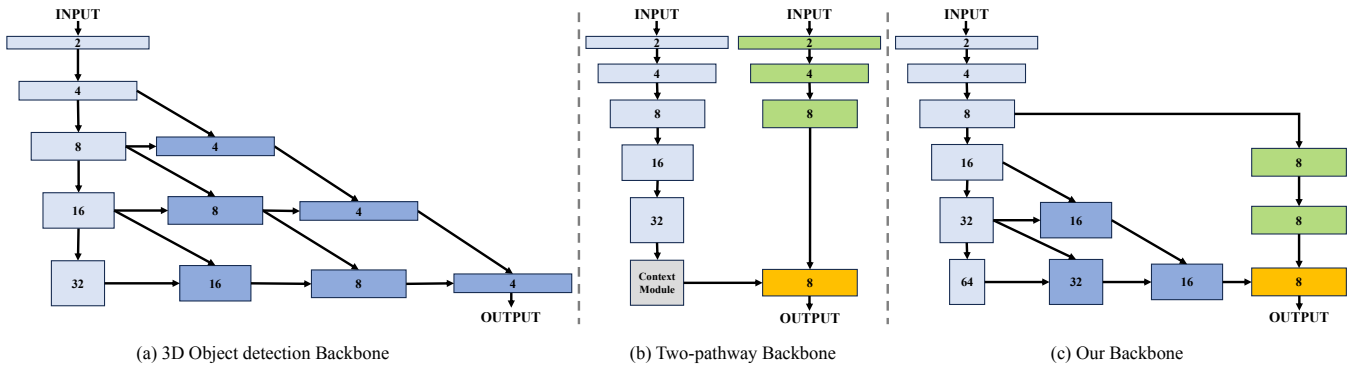


Fig. 2. Comparison between existing backbones and our backbone.

**Multi-task architecture.** Initially, we select the fundamental tasks in autonomous driving as follows: *3D object detection*, *semantic segmentation*, and *dense depth estimation*. Inspired by real-time segmentation architectures [15], [16], [17] and real-time 3D object architecture [18], we propose a novel backbone that emerges the crucial architectural elements of each task to a hybrid structure. Our architectural design is based on two-pathway architecture, which is widely used in real-time segmentation frameworks [15], [17] that effectively and efficiently process the per-pixel classification/regression tasks. To enhance the performance of 3D object detection, where the model needs to estimate the objects with various sizes, we add the aggregation layer to the network that aggregates intermediate features from layers with different receptive fields. The comparison between existing architectures and ours is illustrated in Fig. 2.

**Task-adaptive attention generator.** According to some soft-parameter sharing MTL [4], [5] and Nas-style MTL [7], [18], conditional computation across tasks such as dynamic neural network [7] and attention-based method [19] can mitigate the negative transfer problem by separating the task-generic and task-specific information. To enhance the accuracy of our model while preserving the computational efficiency for real-time framework, we propose a Task-adaptive Attention Generator (TAG), which is the attention-based module that generates task-adaptive attention. Our TAG module can automatically pinpoint the useful feature location for each task along spatial/channel dimensions from the task-generic feature, and it highlights the task-related feature or suppresses the task-unrelated feature, respectively.

In the evaluation results on the Cityscapes 3D, which is the public autonomous driving dataset, our proposed model outperforms the existing baselines with notable margins. We conduct ablation studies of the proposed architecture and TAG module to validate the effectiveness of the proposed methods. Also, we report the extensive experiments that show how the TAG module effectively generates task-adaptive features by an attention-based mechanism. Our contributions are summarized as follows:

- To the best of our knowledge, we propose the first real-time autonomous driving framework capable of performing three crucial tasks: 3D object detection, semantic segmentation, and dense depth estimation.

- We propose a new architecture and attention-based module for these tasks that alleviate the negative transfer problem while maintaining the model inference cost.
- Our framework achieves the best performance and efficiency compared to other baselines consisting of various backbones.

## II. RELATED WORK

### A. Monocular 3D object detection

Monocular 3D object detection is broadly categorized into two main approaches. The first category [18], [20], [21], [22], [23], [24], [25], [26] relies solely on RGB images, annotations, and camera calibrations to predict object localizations around a self-driving vehicle. Most of these methods are based on the foundational work of CenterNet [27]. These methods primarily focus on accurately estimating the object depths and encompass various sub-tasks. The second category compensates for the absence of 3D information in monocular images by utilizing additional data like pre-trained models [28], [29], [30], [31], [32], [33] or CAD models [34], [35]. Some previous works employ pre-trained depth estimators [29], [30], [36], either by transforming the monocular settings to LiDAR/Stereo environments [28], [33] or dedicating more parameters to the depth estimation task. Others integrate depth maps from pre-trained depth completion models [31] or LiDAR-based detectors [28], [37] for model supervision. However, while many monocular 3D object detection models exhibit promising outcomes, their slow inference speeds often hinder real-time applications in autonomous driving. To the best of our knowledge, our model is the first real-time multi-task learning network encompassing monocular 3D object detection.

### B. Multi-Task Learning (MTL)

Hard parameter sharing [12], [13], [14] is the predominant approach to MTL. This scheme improves generalization performance and reduces the computational cost of the network by using shared hidden layers between all tasks. However, a significant downside is its susceptibility to the negative transfer problem [2], [3], which compromises performance when tasks with low relevance are improperly shared.

In contrast, soft-parameter sharing methods [4], [5], [19] address this problem by maintaining separate parameter sets

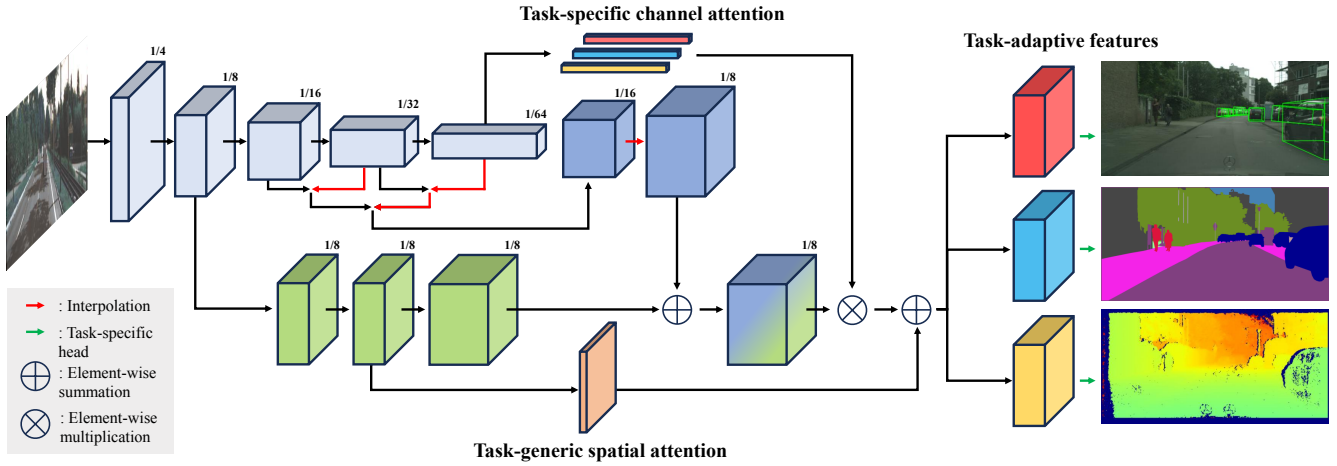


Fig. 3. The overall pipeline of the proposed network.

for each task, with looser constraints across tasks. Nas-style MTL methods have been proposed to adjust shared parameters using a dynamic neural network [7], [8], [9], [10], [11] and NAS [6], [38], [39]. Early works use Mixture-of-Experts [8], [9], [11] to dynamically adjust the sharing pattern across tasks by selectively choosing the experts from a fixed number of module layers. Additionally, some works focus on the skip-or-select strategy, opting for task-specific blocks from residual blocks [7] or designating a shared block for each layer [10], [38], [39]. Especially, the study [40] integrates prompting to encourage the model to learn the task-specific representation. While soft-parameter sharing and Nas-style MTL address negative transfer effectively by adapting their parameter-sharing methods, they fall short of offering the computational efficiency inherent in hard parameter sharing.

### C. Real-time semantic segmentation

Early works [41], [42], [43] have made advancements in improving the operations in network layers such as depth-wise convolution [43], and channel shuffling with group convolution [42]. The study [41] utilizes both low-resolution and high-resolution images to obtain a high-level semantic understanding. However, a challenge with these lightweight architectures is the information loss stemming from consistent downsampling. This hinders the effective retrieval of information during upsampling, resulting in a degradation of the semantic segmentation accuracy. To mitigate this, more recent works [15], [16], [17] have designed networks with two paths: one low-resolution path for extracting semantic information and another high-resolution path to capture spatial details for boundary delineation. Similarly, the work [44] employs a three-branch approach to separately extract details, context, and boundary information. It then utilizes boundary attention to guide the integration of details and context information, resulting in commendable real-time performance.

## III. METHOD

In this section, we describe the proposed method, which consists of two main components: the proposed enhanced

two-pathway architecture and the Task-adaptive Attention Generator (TAG). The overall pipeline of our network is illustrated in Fig. 3.

### A. Enhanced two-pathway Architecture

The detailed architecture of the proposed backbone is depicted in Tab. I. Drawing inspiration from previous works [15], [17], our architecture consists of two branches: a low-resolution branch for extracting semantic information and a high-resolution branch focused on spatial details. The high-resolution branch does not contain any downsampling operation to preserve the intricate spatial details. The branching point between the two resolutions is set to layer 3, where the feature map is at 1/8th the resolution of the input image.

Deeper layers within the low-resolution branch extract more semantic and global features. However, this does not confirm that the last layer is the best representation for all tasks. It is worth noting that skip connections are particularly effective for structured tasks like object detection [45], [46], [47]. Hence, moving away from the traditional context module, we introduce lightweight hierarchical aggregation layers. These layers pool features from the semantic branch, enhancing the 3D object detection performance. For the aggregation layers, we redesign the iterative deep aggregation approach [48] to progressively merge multi-stage features and retrieve a high-resolution feature.

As illustrated in Alg. 1, the aggregation layers aggregate the intermediate features of the model moving from shallow layers to deeper ones, as expressed by the equation:

$$f_{ij}(\mathbf{x}, \mathbf{y}) = W_{ij}^{\text{P}2}(\text{concat}(\mathcal{B}(W_{ij}^{\text{P}1}\mathbf{x}), \mathbf{y})), \quad (1)$$

where the depth of the aggregation layer  $N$  is set to 2. Intermediate features of the semantic branch are characterized by  $\mathbf{x}$  and  $\mathbf{y}$ . For any layer from the node  $i$  to  $j$ , denoted by  $f_{ij}$ , it is composed of two convolution layers with corresponding weights  $W_{ij}^{\text{P}1}, W_{ij}^{\text{P}2}$ . The process also incorporates a  $2 \times$  bilinear interpolation  $\mathcal{B}(\cdot)$ , followed by a single concatenation operation.

Lastly, the proposed backbone combines the aggregated feature  $\mathbf{x}_{\text{agg}}$  with the output of the high-resolution branch

TABLE I

DETAILED ARCHITECTURES OF OUR BACKBONE FOR CITYSCAPES3D.

Stage	Output	Backbone
stem	256 × 512	$\begin{bmatrix} 3 \times 3, 64, \text{stride } 2 \\ 3 \times 3, 64, \text{stride } 2 \end{bmatrix} \times 1$
layer1	256 × 512	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 2$
layer2	128 × 256	$3 \times 3, 128, \text{stride } 2$ $[3 \times 3, 128] \times 3$
layer3	64 × 128, 128 × 256	$\frac{3 \times 3, 256, \text{stride } 2}{[3 \times 3, 256] \times 3}$ $\begin{bmatrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{bmatrix} \times 2$
layer4	32 × 64, 128 × 256	$\frac{3 \times 3, 512, \text{stride } 2}{[3 \times 3, 512] \times 3}$ $\begin{bmatrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{bmatrix} \times 2$
layer5	16 × 32, 128 × 256	$\frac{3 \times 3, 1024, \text{stride } 2}{[3 \times 3, 1024] \times 3}$ identity
decode layer1	32 × 64, 128 × 256	$\begin{bmatrix} 3 \times 3, 512 \\ 3 \times 3, 512 \end{bmatrix} \times 1$ identity
decode layer2	64 × 128, 128 × 256	$\begin{bmatrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{bmatrix} \times 2$ identity
fusion	128 × 256	attention based fusion
head (MTL)	1024 × 2048	$\begin{bmatrix} 3 \times 3, 256 \\ 1 \times 1, C_t \end{bmatrix} \times 2 \times  \mathcal{T} $ $C_t$ : outplane of task $t$ $ \mathcal{T} $ : number of tasks

$\mathbf{x}_{\text{detail}}$  to yield a task-generic feature  $\mathbf{h}$  as follows:

$$\mathbf{h} = \mathcal{B}(\mathbf{x}_{\text{agg}}) + g(\mathbf{x}_{\text{detail}}), \quad (2)$$

where  $g$  is the convolution layer that projects the high-resolution feature to match the channel dimension of  $\mathbf{x}_{\text{agg}}$ .

### B. Task-adaptive attention generator (TAG)

To mitigate the negative transfer problem prevalent in hard-parameter sharing, we propose the TAG module. This module is designed to automatically identify the location of task-adaptive features within the task-generic feature tensor  $\mathbf{h}$ , which takes the shape of  $(B, C, H/s, W/s)$ . The TAG module operates along the channel and spatial dimensions. Its primary function is to emphasize the task-relevant information while suppressing the unrelated details. For channel attention, TAG receives the last feature of the semantic branch  $\mathbf{x}_{\text{semantic}}$  to capture the model’s most semantic and global feature. On the other hand, the TAG module derives task-generic spatial attention from the output of the high-resolution branch  $\mathbf{x}_{\text{detail}}$  as follows:

$$\alpha_t = \sigma(S_{ch}^t(\mathbf{x}_{\text{semantic}})), \quad \beta = \sigma(S_{sp}(\mathbf{x}_{\text{detail}})), \quad (3)$$

where  $t \in \{0, 1, 2\}$  is the task index. The operation  $\sigma$  is the sigmoid function.  $S_{ch}^t$  is the channel attention layer consisting of two MLP layers following global average pooling. Meanwhile,  $S_{sp}$  is the task-generic spatial attention layer that includes single dilated convolution. Through Eq. 3, we obtain both task-specific channel attention and task-generic spatial attention, represented as  $\alpha_t$  and  $\beta$ , respectively.

The TAG module constructs the task-adaptive features  $\mathbf{h}_t$  by performing channel-wise multiplication between the task-specific channel attention map  $\alpha_t$  and the task-generic feature  $\mathbf{h}$ , followed by an addition of the spatial attention  $\beta$  as follows:

$$\mathbf{h}_t = \alpha_t \cdot \mathbf{h} + \beta. \quad (4)$$

### Algorithm 1 Forward procedure of aggregation layer

---

**Require:**  $L = [\mathbf{x}_1, \dots, \mathbf{x}_N]$   $\triangleright$  list contains  $N$  Multi-scale features  
**Ensure:**  $\mathbf{x}_{\text{agg}}$   $\triangleright$  aggregated feature  
 $\text{idx}_1, \text{idx}_2 \leftarrow 1, 0$   $\triangleright$  layer/scale index

```

while  $\text{idx}_1 \neq N$  do
   $\text{idx}_2 \leftarrow \text{idx}_1$ 
  while  $\text{idx}_2 \neq 0$  do
     $L[\text{idx}_2 - 1] \leftarrow f_{\text{idx}_1, \text{idx}_2}(L[\text{idx}_2 - 1], L[\text{idx}_2])$ 
     $\text{idx}_2 \leftarrow \text{idx}_2 - 1$   $\triangleright$  refer to Eq. 1
  end while
   $\text{idx}_1 \leftarrow \text{idx}_1 + 1$ 
end while
 $\mathbf{x}_{\text{agg}} \leftarrow L[0]$ 

```

---

### C. Multi-head Structure

3D object detection frameworks typically consist of multiple sub-tasks, including classification, center localization, and heading direction estimation. In our approach, we integrate heads for various roles: sub-tasks related to 3D object detection, semantic segmentation, and dense depth estimation to maintain a compact model structure. All these heads employ shallow CRC layers, which are composed of  $3 \times 3 \text{Conv-ReLU-}1 \times 1 \text{Conv}$ . For the sub-task associated with 3D detection, we basically adopt the approach of [23], but with an added functionality to deduce the pitch and roll, facilitating the inference of the 3-DoF heading direction.

### D. Loss function

For the 3D detection loss, our model is largely aligned with [23]. However, we introduce modifications for the pitch and roll tasks. Specifically, the heading loss for these tasks is informed by the method proposed in [49]. As for the semantic segmentation and dense depth estimation tasks, we employ cross-entropy and smooth L1 loss, respectively. In setting the loss weights for our three tasks, we follow the scheme presented in [40]. Here, the weights  $\lambda_{det}, \lambda_{seg}, \lambda_{dep}$  are set to 1, 100, and 1, respectively. Our final loss term is formulated as follows:

$$\mathcal{L}_{task} = \lambda_{det} \mathcal{L}_{det} + \lambda_{seg} \mathcal{L}_{seg} + \lambda_{dep} \mathcal{L}_{dep}. \quad (5)$$

## IV. EXPERIMENTS

### A. Dataset

1) *Cityscapes-3D*: Cityscapes-3D [50] is an extended version of the Cityscapes dataset [51] by incorporating additional 3D bounding box annotations. It retains the original Cityscapes configuration with 2,975 training images and 500 validation images, all finely annotated at a resolution of  $1024 \times 2048$ .

In contrast to previous 3D object detection datasets, such as nuScenes [52], Waymo [53], and KITTI [54], Cityscapes-3D has been designed specifically for the purpose of monocular 3D object detection. This distinction arises because Cityscapes-3D relies on ground-truth depth labels

TABLE II

EVALUATION ON CITYSCAPES-3D VALIDATION SET (BEST, SECOND-BEST,  $\geq 30$ FPS, POSITIVE/NEGATIVE PERFORMANCE DEVIATION FROM STL).

Method	Backbone	3D Det.	Semantic seg.	Dense depth	$\Delta_{\mathcal{T}}$ (%) $\uparrow$	FPS $\uparrow$	GFLOPs $\downarrow$	Params (M) $\downarrow$
		DS $\uparrow$	mIoU $\uparrow$	RMSE $\downarrow$				
STL	DLA34	54.55	-	-	0.00	17.4	246	20.1
		-	66.38	-	0.00	17.8	222	19.9
		-	-	5.50	0.00	17.8	223	19.9
MTL	DLA34	38.05 (-30.25%)	63.79 (-3.90%)	6.15 (-11.82%)	-15.32	17.3	285	20.4
	ResNet50	51.92 (-4.82%)	66.40 (+0.03%)	5.87 (-6.73%)	-3.84	12.4	373	44.5
	ICNet	39.67 (-27.28%)	61.18 (-7.83%)	5.68 (-3.27%)	-12.79	21.7	226	44.4
	DDRNet-23-slim	32.36 (-40.68%)	60.06 (-9.52%)	6.26 (-13.82%)	-21.34	50.9	75	7.5
	DDRNet-23	49.19 (-9.83%)	62.80 (-5.39%)	5.93 (-7.82%)	-7.68	30.1	216	23.9
	DDRNet-39	50.83 (-6.82%)	<b>67.17 (+1.19%)</b>	5.59 (-1.64%)	-2.42	20.3	345	35.5
	PIDNet-S	40.12 (-26.45%)	50.57 (-23.82%)	6.82 (-24.00%)	-24.76	42.9	84	9.4
	PIDNet-M	50.78 (-6.91%)	58.31 (-12.16%)	6.27 (-14.00%)	-11.02	24.9	241	32.2
	PIDNet-L	53.48 (-1.96%)	58.81 (-11.40%)	6.33 (-15.09%)	-9.48	20.4	339	40.3
MTL	Ours (w/o TAG)	<u>57.46 (+5.33%)</u>	65.90 (-0.72%)	<u>5.49 (+0.18%)</u>	+1.60	30.8	218	32.8
	Ours	<b>61.21 (+12.21%)</b>	<u>66.50 (+0.18%)</u>	<b>5.36 (+2.55%)</b>	<b>+4.98</b>	30.2	219	33.9

from stereo cameras, rather than LiDAR. This often introduces disparities between images and depth maps. Moreover, Cityscapes-3D provides labels for both pitch and roll angles of vehicles, whereas many previous datasets primarily concentrate on yaw. Recognizing these angles is crucial in instances where roads are not flat, a common occurrence.

### B. Baselines

We compare our model to various backbones with multi-head structures to validate the effectiveness of the proposed architecture and TAG module. We select five well-known baselines as comparison sets: ResNet50 [55], DLA34 [48], ICNet [56], DDRNet [17], and PIDNet [44].

### C. Evaluation Metric

For the task of 3D object detection, we employ the Detection Score, which is the official evaluation index provided by Cityscapes-3D. In the context of semantic segmentation, the mean Intersection over Union (mIoU) serves as our metric, while for dense depth estimation, we use the Root Mean Square Error (RMSE) for evaluation. To understand the in/decremental shift in overall task performance from a single-task setup, we adopt the averaged relative performance  $\Delta_{\mathcal{T}}$  in [7] as follows:

$$\Delta_{\mathcal{T}} = \frac{100}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (-1)^{l_t} \frac{(\mathcal{M}_t - \mathcal{M}_t^{single})}{\mathcal{M}_t^{single}}, \quad (6)$$

where  $\mathcal{M}_t$  and  $\mathcal{M}_t^{single}$  are the metric of task  $t$  from each method and the single task baseline, respectively. The constant  $l_t$  is 1 if a lower value represents better for the metric  $\mathcal{M}_t$  and 0 otherwise.

### D. Implementation details

Like all previous baselines [17], [44], [48], [55], we initialize our model by pretraining on ImageNet [57]. All models undergo training with an input resolution set at  $224 \times 224$ , a batch size of 256, spanning a total of 100 epochs. We commence with a learning rate of 0.1, and reduce it by a factor of 10 during the 30th, 60th, and 90th epochs.

All networks are trained using SGD with a weight decay of  $1 \times 10^{-4}$  and Nesterov momentum set to 0.9.

For Cityscapes-3D training, the same multi-head structure and configuration are maintained across all baselines to ensure a fair comparison, differing only in the backbone (*e.g.*,  $\mathcal{L}_{task}$ , loss scale, and learning rate). Training parameters include an input resolution at the original  $1024 \times 2048$  dimensions, a batch size of 2, and no image augmentation. For all our baseline models, the Adam optimizer is employed with an initial learning rate of  $2 \times 10^{-5}$ . Moreover, we employ a polynomial learning rate policy, with a power of 0.9, for a progressive reduction in the learning rate.

### E. Evaluation results on Cityscapes-3D

To compare the MTL baselines and ours to single-task setups, we select the DLA34 [48] which is widely used in 3D object detection, as the benchmark for single-task setups. It is important to note that among all tasks, 3D object detection is the most computationally intensive. To investigate the trade-off between performance and computational costs, we analyze the GFLOPs, number of parameters, and inference speed. We measure the inference speed for all baselines on an NVIDIA Titan RTX GPU, with a batch size of 1.

Results detailed in Tab. II highlight the prowess of the proposed architecture, especially in the detection task. For semantic segmentation, the absence of the TAG module results in a performance drop of  $-0.72\%$ p and  $-1.91\%$ p when compared to the single-task configurations of DLA34 and DDRNet-39, respectively. However, incorporating the TAG module elevates performance across all metrics, incurring only a slight rise in inference time by 0.56 ms. This increase is minimal when compared to the variant devoid of the TAG module. Intriguingly, regardless of the presence of the TAG module, our method emerges as the sole configuration showcasing a surge in the averaged relative performance  $\Delta_{\mathcal{T}}$  relative to the single-task setup. This amounts to a notable performance boost of  $+7.4\%$ p against the second-best model, DDRNet-39, while also achieving a 32.8% reduction in inference time.

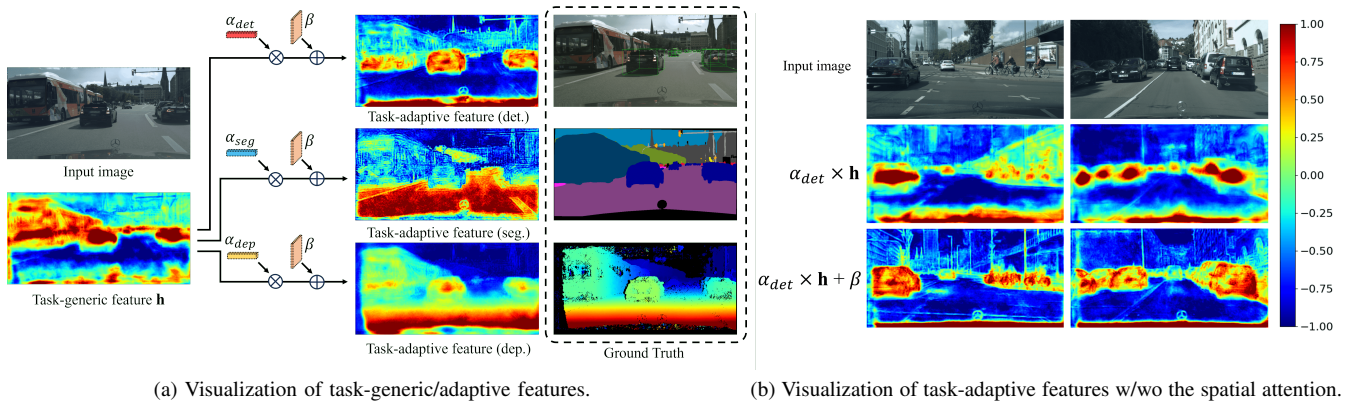


Fig. 4. Feature heatmap visualization with respect to attention in TAG module.

## F. Additional experiments

1) *Visualizing the impact of TAG modules:* To demonstrate the efficacy of channel attention within the TAG, we visualize task-adaptive features  $\mathbf{h}_t$  produced from channel/spatial attention. These features are derived by using task-specific channel attention  $\alpha_t$  and task-generic spatial attention  $\beta$ , with the task-generic feature  $\mathbf{h}$ . Following this, we apply global channel pooling to condense the resultant features into a single channel for visualization purposes. As shown in Fig. 4-(a), it is evident that the TAG module can generate distinctive task-adaptive features. For example,  $\alpha_{det}$  mainly emphasizes prominent objects, highlighted in red on the feature map, and obstacles pertinent to the 3D detection task. Conversely,  $\alpha_{seg}$  focuses on the semantic regions, with a particular emphasis on the ‘road’ class in the given example. Meanwhile,  $\alpha_{dep}$  adjusts the features to closely match the patterns observed in the depth ground truth.

Alongside our examination of the TAG module, we delve into the influence of task-generic spatial attention  $\beta$  by contrasting the visualization of the feature map both with and without the incorporation of the attention. As shown in Fig. 4-(b), the spatial attention drawing from detailed features of a high-resolution branch distinctly highlights spatial details. This is particularly prominent in object boundaries and smaller elements, notably in traffic participants such as cars and cyclists. Additionally, the intricate designs of poles and trees are faithfully retained. This underscores the effectiveness of spatial attention in directing the model to emphasize spatial data.

2) *Ablation studies on architectural elements:* To delve deeper into the individual contributions of specific architectural elements, we conduct comprehensive ablation studies. Our design encapsulates four core components: the dual-branch structure, aggregation layer, semantic channel attention, and spatial attention. In the experiment presented in Tab. III, we either remove or replace one of these components for each trial (e.g., aggregation layer  $\rightarrow$  Pyramid pooling module (PPM) [58]). These experiments adopt the same configurations as those in Tab. II. The results consistently reveal that models lacking or altering one of our architectural components tend to underperform compared to our original

TABLE III

ABLATION STUDIES ON ARCHITECTURAL ELEMENTS (BR: HIGH-RES BRANCH, AGG: AGGREGATION LAYER,  $\{\alpha_t, \beta\}$ : CHANNEL/SPATIAL ATTN).

br	agg	$\alpha_t$	$\beta$	3D Det.	Seg.	Dep.	FPS	$\Delta_{\mathcal{T}}$ (%)
				DS	mIoU	RMSE		
DLA34 (STL)				54.55	66.38	5.50	17.4~17.8	0.00
✓	✓	✓	✓	49.13	58.99	6.21	37.1	-11.33
✓	✓	✓	✓	49.10	60.09	5.74	35.0	-7.94
✓	✓	✓	✓	55.12	61.02	5.56	30.3	-2.71
✓	✓	✓	✓	58.21	63.71	5.62	30.7	+0.17
✓	✓	✓	✓	<b>61.21</b>	<b>66.50</b>	<b>5.36</b>	30.2	<b>+4.98</b>

proposal. The proposed architecture with the TAG module achieves the best performance among the existing baselines. This attests to the value and effectiveness of each subcomponent in enhancing performance within the MTL framework.

## V. CONCLUSION

In this paper, we present a new multi-task learning approach tailored for real-time autonomous driving, encompassing monocular 3D object detection, semantic segmentation, and dense depth estimation. Our uniquely designed network structure addresses the negative transfer problem inherent in heterogeneous multi-task learning, ensuring computational efficiency. Our proposed architecture, complemented by an attention-based module, capitalizes on shared knowledge across tasks, fostering task-adaptive learning. Drawing from domains of 3D object detection, semantic segmentation, and dense depth estimation, our architecture is optimized to adeptly handle per-pixel classification/regression tasks, accommodating objects of varied sizes. The TAG mechanism effectively accentuates task-relevant features, suppressing unrelated data. Our model’s superior performance over baseline models highlights its pivotal role in advancing safe and efficient autonomous vehicles. We believe that this study signifies a substantial progression in multi-task learning tailored for real-time autonomous driving.

## ACKNOWLEDGEMENT

This research was supported in part by Autonomous Driving Center, Hyundai Motor Company.

## REFERENCES

- [1] S. Ruder, "An overview of multi-task learning in deep neural networks," *arXiv preprint arXiv:1706.05098*, 2017.
- [2] Z. Kang, K. Grauman, and F. Sha, "Learning with whom to share in multi-task feature learning," in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 521–528.
- [3] T. Standley, A. Zamir, D. Chen, L. Guibas, J. Malik, and S. Savarese, "Which tasks should be learned together in multi-task learning?" in *International Conference on Machine Learning*. PMLR, 2020, pp. 9120–9132.
- [4] I. Misra, A. Shrivastava, A. Gupta, and M. Hebert, "Cross-stitch networks for multi-task learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3994–4003.
- [5] S. Ruder, J. Bingel, I. Augenstein, and A. Søgaard, "Sluice networks: Learning what to share between loosely related tasks," *arXiv preprint arXiv:1705.08142*, vol. 2, 2017.
- [6] Y. Gao, H. Bai, Z. Jie, J. Ma, K. Jia, and W. Liu, "Mtl-nas: Task-agnostic neural architecture search towards general-purpose multi-task learning," in *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 2020, pp. 11 543–11 552.
- [7] X. Sun, R. Panda, R. Feris, and K. Saenko, "Adashare: Learning what to share for efficient deep multi-task learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 8728–8740, 2020.
- [8] K. Maziarz, E. Kokkopolou, A. Gesmundo, L. Sbaiz, G. Bartok, and J. Berent, "Flexible multi-task networks by learning parameter allocation," *arXiv preprint arXiv:1910.04915*, 2019.
- [9] J. Ma, Z. Zhao, J. Chen, A. Li, L. Hong, and E. H. Chi, "Snr: Sub-network routing for flexible parameter sharing in multi-task learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 216–223.
- [10] P. Guo, C.-Y. Lee, and D. Ulbricht, "Learning to branch for multi-task learning," in *International conference on machine learning*. PMLR, 2020, pp. 3854–3863.
- [11] C. Fernando, D. Banarse, C. Blundell, Y. Zwols, D. Ha, A. A. Rusu, A. Pritzel, and D. Wierstra, "Pathnet: Evolution channels gradient descent in super neural networks," *arXiv preprint arXiv:1701.08734*, 2017.
- [12] B. Jou and S.-F. Chang, "Deep cross residual learning for multitask visual recognition," in *Proceedings of the 24th ACM international conference on Multimedia*, 2016, pp. 998–1007.
- [13] J. Huang, R. S. Feris, Q. Chen, and S. Yan, "Cross-domain image retrieval with a dual attribute-aware ranking network," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1062–1070.
- [14] H. Bilen and A. Vedaldi, "Integrated perception with recurrent multi-task neural networks," *Advances in neural information processing systems*, vol. 29, 2016.
- [15] C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, and N. Sang, "Bisenet: Bilateral segmentation network for real-time semantic segmentation," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 325–341.
- [16] C. Yu, C. Gao, J. Wang, G. Yu, C. Shen, and N. Sang, "Bisenet v2: Bilateral network with guided aggregation for real-time semantic segmentation," *International Journal of Computer Vision*, vol. 129, pp. 3051–3068, 2021.
- [17] Y. Hong, H. Pan, W. Sun, and Y. Jia, "Deep dual-resolution networks for real-time and accurate semantic segmentation of road scenes," *arXiv preprint arXiv:2101.06085*, 2021.
- [18] X. Liu, N. Xue, and T. Wu, "Learning auxiliary monocular contexts helps monocular 3d object detection," in *AAAI*, vol. 36, no. 2, 2022, pp. 1810–1818.
- [19] S. Liu, E. Johns, and A. J. Davison, "End-to-end multi-task learning with attention," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 1871–1880.
- [20] G. Brazil and X. Liu, "M3d-rpn: Monocular 3d region proposal network for object detection," in *ICCV*, 2019, pp. 9287–9296.
- [21] A. Kumar, G. Brazil, E. Corona, A. Parchami, and X. Liu, "Deviant: Depth equivariant network for monocular 3d object detection," in *ECCV*. Springer, 2022, pp. 664–683.
- [22] Y. Lu, X. Ma, L. Yang, T. Zhang, Y. Liu, Q. Chu, J. Yan, and W. Ouyang, "Geometry uncertainty projection network for monocular 3d object detection," in *ICCV*, 2021, pp. 3111–3121.
- [23] X. Ma, Y. Zhang, D. Xu, D. Zhou, S. Yi, H. Li, and W. Ouyang, "Delving into localization errors for monocular 3d object detection," in *CVPR*, 2021, pp. 4721–4730.
- [24] A. Mousavian, D. Anguelov, J. Flynn, and J. Kosecka, "3d bounding box estimation using deep learning and geometry," in *CVPR*, 2017, pp. 7074–7082.
- [25] Z. Wu, Y. Wu, J. Pu, X. Li, and X. Wang, "Attention-based depth distillation with 3d-aware positional encoding for monocular 3d object detection," *arXiv preprint arXiv:2211.16779*, 2022.
- [26] Y. Zhang, J. Lu, and J. Zhou, "Objects are different: Flexible monocular 3d object detection," in *CVPR*, 2021, pp. 3289–3298.
- [27] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as points," *arXiv preprint arXiv:1904.07850*, 2019.
- [28] Y.-N. Chen, H. Dai, and Y. Ding, "Pseudo-stereo for monocular 3d object detection in autonomous driving," in *CVPR*, 2022, pp. 887–897.
- [29] M. Ding, Y. Huo, H. Yi, Z. Wang, J. Shi, Z. Lu, and P. Luo, "Learning depth-guided convolutions for monocular 3d object detection," in *CVPRW*, 2020, pp. 1000–1001.
- [30] K.-C. Huang, T.-H. Wu, H.-T. Su, and W. H. Hsu, "Monodr: Monocular 3d object detection with depth-aware transformer," in *CVPR*, 2022, pp. 4012–4021.
- [31] L. Peng, X. Wu, Z. Yang, H. Liu, and D. Cai, "Did-m3d: Decoupling instance depth for monocular 3d object detection," in *ECCV*. Springer, 2022, pp. 71–88.
- [32] C. Reading, A. Harakeh, J. Chae, and S. L. Waslander, "Categorical depth distribution network for monocular 3d object detection," in *CVPR*, 2021, pp. 8555–8564.
- [33] Y. Wang, W.-L. Chao, D. Garg, B. Hariharan, M. Campbell, and K. Q. Weinberger, "Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving," in *CVPR*, 2019, pp. 8445–8453.
- [34] F. Manhardt, W. Kehl, and A. Gaidon, "Roi-10d: Monocular lifting of 2d detection to 6d pose and metric shape," in *CVPR*, 2019, pp. 2069–2078.
- [35] Y. Xiang, W. Choi, Y. Lin, and S. Savarese, "Data-driven 3d voxel patterns for object category recognition," in *CVPR*, 2015, pp. 1903–1911.
- [36] R. Zhang, H. Qiu, T. Wang, X. Xu, Z. Guo, Y. Qiao, P. Gao, and H. Li, "Monodetr: Depth-aware transformer for monocular 3d object detection," *arXiv preprint arXiv:2203.13310*, 2022.
- [37] Y. Hong, H. Dai, and Y. Ding, "Cross-modality knowledge distillation network for monocular 3d object detection," in *ECCV*. Springer, 2022, pp. 87–104.
- [38] D. S. Raychaudhuri, Y. Suh, S. Schuler, X. Yu, M. Faraki, A. K. Roy-Chowdhury, and M. Chandraker, "Controllable dynamic multi-task architectures," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10955–10964.
- [39] W. Choi and S. Im, "Dynamic neural network for multi-task learning searching across diverse network topologies," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 3779–3788.
- [40] H. Ye and D. Xu, "Taskprompter: Spatial-channel multi-task prompting for dense scene understanding," in *The Eleventh International Conference on Learning Representations*, 2022.
- [41] M. Orsic, I. Kreso, P. Bevandic, and S. Segvic, "In defense of pre-trained imagenet architectures for real-time semantic segmentation of road-driving images," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 607–12 616.
- [42] M. Gamal, M. Siam, and M. Abdel-Razek, "Shuffleleg: Real-time semantic segmentation network," *arXiv preprint arXiv:1803.03816*, 2018.
- [43] H. Li, P. Xiong, H. Fan, and J. Sun, "Dfanet: Deep feature aggregation for real-time semantic segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 9522–9531.
- [44] J. Xu, Z. Xiong, and S. P. Bhattacharyya, "Pidnet: A real-time semantic segmentation network inspired by pid controllers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 19 529–19 539.
- [45] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik, "Hypercolumns for object segmentation and fine-grained localization," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 447–456.
- [46] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [47] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the*

- IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [48] F. Yu, D. Wang, E. Shelhamer, and T. Darrell, “Deep layer aggregation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2403–2412.
- [49] T. Wang, X. Zhu, J. Pang, and D. Lin, “Fcos3d: Fully convolutional one-stage monocular 3d object detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 913–922.
- [50] N. Gähler, N. Jourdan, M. Cordts, U. Franke, and J. Denzler, “Cityscapes 3d: Dataset and benchmark for 9 dof vehicle detection,” *arXiv preprint arXiv:2006.07864*, 2020.
- [51] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The cityscapes dataset for semantic urban scene understanding,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223.
- [52] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nusenes: A multimodal dataset for autonomous driving,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [53] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, *et al.*, “Scalability in perception for autonomous driving: Waymo open dataset,” in *CVPR*, 2020, pp. 2446–2454.
- [54] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in *CVPR*, 2012.
- [55] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [56] H. Zhao, X. Qi, X. Shen, J. Shi, and J. Jia, “Icnet for real-time semantic segmentation on high-resolution images,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 405–420.
- [57] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, *et al.*, “Imagenet large scale visual recognition challenge,” *International journal of computer vision*, vol. 115, pp. 211–252, 2015.
- [58] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.