

Fully Exploiting Vision Foundation Model's Profound Prior Knowledge for Generalizable RGB-Depth Driving Scene Parsing

Sicen Guo[✉], Tianyou Wen[✉], Chuang-Wei Liu[✉],
 Qijun Chen[✉] *Senior Member, IEEE*, Rui Fan[✉] *Senior Member, IEEE*

Abstract—Recent vision foundation models (VFM), typically based on Vision Transformer (ViT), have significantly advanced numerous computer vision tasks. Despite their success in tasks focused solely on RGB images, the potential of VFM in RGB-depth driving scene parsing remains largely under-explored. In this article, we take one step toward this emerging research area by investigating a feasible technique to fully exploit VFM for generalizable RGB-depth driving scene parsing. Specifically, we explore the inherent characteristics of RGB and depth data, thereby presenting a Heterogeneous Feature Integration Transformer (HFIT). This network enables the efficient extraction and integration of comprehensive heterogeneous features without re-training ViTs. Relative depth prediction results from VFM, used as inputs to the HFIT side adapter, overcome the limitations of the dependence on depth maps. Our proposed HFIT demonstrates superior performance compared to all other traditional single-modal and data-fusion scene parsing networks, pre-trained VFM, and ViT adapters on the Cityscapes and KITTI Semantics datasets. We believe this novel strategy paves the way for future innovations in VFM-based data-fusion techniques for driving scene parsing. Our source code is publicly available at <https://mias.group/HFIT>.

Index Terms—vision foundation models, Transformer, computer vision, scene parsing, data-fusion.

I. INTRODUCTION

The rise of vision foundation models (VFM) [1]–[4] has heralded a new era in computer vision. Traditional fundamental computer vision tasks, such as semantic segmentation [1], object detection [5], and depth estimation [3], [6], have all started turning to VFM for more compelling performance. This feasibility generally arises from the VFM's stronger capability in extracting profound, informative features through a range of related computer vision tasks [1]–[3]. Despite several attempts made in RGB-depth (often abbreviated as RGB-D) driving scene parsing [7]–[9], existing prior arts based on VFM are sketchy, and the achieved results remain

This research was supported by the National Natural Science Foundation of China under Grants 62473288 and 62233013, the Fundamental Research Funds for the Central Universities, and the Xiaomi Young Talents Program. (Corresponding author: Rui Fan)

The authors are with the College of Electronic and Information Engineering, Shanghai Institute of Intelligent Science and Technology, Shanghai Research Institute for Intelligent Autonomous Systems, Shanghai Key Laboratory of Intelligent Autonomous Systems, State Key Laboratory of Autonomous Intelligent Unmanned Systems, and Frontiers Science Center for Intelligent Autonomous Systems (Ministry of Education), Tongji University, Shanghai 201804, China (e-mail: {guosicen, tmi, cwliu, qjchen}@tongji.edu.cn; rui.fan@ieee.org)

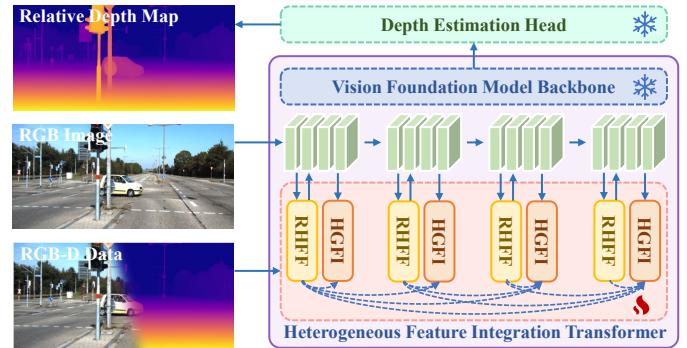


Fig. 1: An overview of our proposed HFIT. Through the interaction between recalibrated heterogeneous feature fusion (RHFF) modules and holistic gated feature integration (HGFI) modules, heterogeneous features are fully integrated with the profound prior features.

unsatisfactory. Therefore, this article primarily focuses on fully exploiting VFM's profound prior knowledge for RGB-D driving scene parsing.

The integration of RGB images and depth maps has been proven to significantly improve the performance of driving scene parsing [10]–[13]. However, a significant limitation of these fusion models is their reliance on depth maps, which restricts their applicability in scenarios where LiDARs are unavailable. Stereo cameras provide a practical alternative for obtaining depth information [14], [15]. However, even supervised methods fall short of achieving LiDAR-level accuracy, and the performance gap is even wider for unsupervised and self-supervised approaches [16]. Furthermore, integrating these heterogeneous characteristics can degrade scene parsing performance when the depth maps lack sufficient precision, *e.g.*, due to inaccuracies in camera-LiDAR calibration [17]–[19]. Surprisingly, VFM like Depth Anything V2 [20] have demonstrated exceptional effectiveness not only in extracting informative deep features but also in learning geometric information and performing relative depth estimation. Since these methods perceive only relative spatial relationships and cannot directly derive exact depth values, many researchers assume that their impact on scene parsing is minimal [21], [22]. However, we contend that relative depth maps are sufficient to enhance the performance of scene parsing. Utilizing VFM's

relative depth estimation output to assist data-fusion models deserves greater attention. Thus, this article marks the initial exploration in this uncharted territory, specifically focusing on the utilization of not only the VFM's deep features but also their relative depth outputs.

The main challenges for data-fusion models to have spatial understanding ability are in the following aspects. First, VFM's have limited capacity to understand depth maps as they are trained only on RGB images [2], [3], [20]. Consequently, directly inputting depth maps into VFM's results in poor performance. Most existing data-fusion models typically use dual encoders with separate trainable parameters to extract heterogeneous features from the RGB-D pairs [7], [10], [11]. Nevertheless, this strategy not only increases computational load but also overlooks the inherent characteristics between heterogeneous features. Therefore, developing an architecture that can leverage VFM's informative prior knowledge while enabling the extraction and integration of spatial features with minimum resource utilization, stands as an under-explored area requiring further investigation. Additionally, RGB images are abundant in global semantic cues [23], making them suitable for feature extraction via pre-trained Vision Transformers (ViTs), while depth images provide gradient-related semantics (local details) [24], which can be captured more efficiently through convolutional neural networks (CNNs). Thus, side adapters may be the optimal architecture choice. Freezing pre-trained ViTs and solely updating adapters not only leverages the profound prior knowledge of ViTs but also enables the model to extract spatial features with minimal resource usage [25], [26]. Although it may not achieve the performance of full fine-tuning strategies, its significantly lower number of trainable parameters and reduced resource consumption provide compelling advantages over full fine-tuning.

Additionally, strategies for fusing prior features from ViTs with heterogeneous features from the adapter are crucial for generalizable RGB-D driving scene parsing [27], [28]. Indiscriminate fusion strategies, such as direct overlaying interactions between modalities [12], not only introduce noise into the feature space but also exhibit the inadaptability of RGB-based models to mixed-modality inputs [29], [30]. Since traditional adapter inputs are typically RGB images, indiscriminate feature fusion [25], [31] is relatively harmless. However, if the adapter extracts heterogeneous information, indiscriminate fusion strategies will potentially drown out valid complementary information due to different semantic spaces [32], [33]. Therefore, it becomes essential to explore the inherent differences in heterogeneous feature characteristics, emphasizing their respective strengths and complementing each other based on reliability.

In addition to integrating heterogeneous features, the fusion of multi-level features with varying degrees of semantics has often been overlooked [11], [12], [23], [27]. As the network depth increases, semantic clues become richer, while fine-grained information, such as boundary details, tends to be lost [34]. SNE-RoadSeg [10] and UNet++ [35] employ densely connected skip connections to enhance comprehensive holistic feature integration. While this approach incorporates both fine-grained and coarse-grained details, it fails to address the

semantic gap between features extracted at different network depths. These limitations underscore the importance of assessing the significance of each feature vector within the feature map and aggregating the information accordingly.

To tackle the limitations discussed above, we introduce **Heterogeneous Feature Integration Transformer (HFIT)** (see Fig. 1). HFIT investigates the utilization of VFM's for generalizable RGB-D driving scene parsing, thereby unleashing the power of profound prior knowledge embedded within these pre-trained VFM's. We begin by performing relative depth estimation [20]. Then, we develop a side adapter to capture robust and multi-scale spatial pyramid features from RGB-D pairs and adaptively integrate them with the prior features from VFM's. HFIT adapter starts with a duplex spatial prior extractor (DSPE), which aims to capture local spatial priors from the RGB-D data. Then, we develop a recalibrated heterogeneous feature fusion module, which supplements VFM's with multi-scale spatial geometric features. Finally, we design a holistic gated feature integration module, which first selectively integrates features from multiple levels and then fuses multi-scale features from VFM's and the adapter. Extensive experiments have been conducted on publicly available datasets [36], [37], and unequivocally demonstrate the superior performance of our proposed HFIT. Our contributions are as follows:

- To the best of our knowledge, this study is the first to investigate the adaptation of VFM's for RGB-D driving scene parsing. HFIT can directly extract and integrate heterogeneous features without re-training ViTs. By leveraging relative depth prediction results from VFM's as inputs to the HFIT adapter, our approach overcomes the limitations of the dependence on depth maps.
- We introduce a DSPE strategy to jointly capture local semantics from RGB-D data, enabling HFIT to effectively learn heterogeneous features.
- We design an RHFF module to collaboratively recalibrate spatial priors, enhancing the comprehensiveness of feature fusion.
- We propose an HGFI strategy to combine complementary strengths of multi-level features, equipping HFIT with enhanced fine-grained feature representation capabilities.

II. RELATED WORK

A. Conventional RGB-D Semantic Segmentation Networks

Considering the feature fusion stage of state-of-the-art (SoTA) RGB-D semantic segmentation networks, we can categorize them into three main types: early fusion, intermediate fusion, and late fusion [38]. Early fusion methods typically combine RGB and depth images at the input stage. While this approach is straightforward, it falls short in achieving a comprehensive understanding of the environment [38]. Intermediate fusion techniques [10] extract and fuse these heterogeneous features from RGB-D pairs using duplex encoders. Late fusion approaches [12], [39], [40] utilize two parallel encoders, with a primary focus on feature fusion within the decoder. Although these strategies have proven to be more effective than single-modal methods, they still face challenges such as insufficient utilization of depth features and high computational costs associated with multiple feature extractors [30].

B. Feature Extraction and Fusion Strategies

Feature extraction and fusion play essential roles in advancing various computer vision tasks [41]–[44]. Traditional fusion strategies perform effectively when features exhibit relevance and consistency. However, when dealing with multi-scale, multi-level, and multi-modality features, indiscriminate extraction and fusion can introduce redundancies and noise due to the semantic gap [17]. To address these limitations, the separation-and-aggregation gate (SAGate) [45] recalibrates and aggregates heterogeneous features to generate enhanced selective representations specifically for segmentation. Additionally, several studies [10], [35] have addressed this issue by introducing densely connected skip connections to integrate fine-grained and coarse-grained information. The cross feature module (CFM) [46] incorporates adaptive aggregation mechanisms to selectively retain useful feature maps and refine multi-level features while filtering out background noise. Nevertheless, these models often overlook the semantic gap between hierarchical features extracted at different network depths, which may result in incomplete fusion of high-level and low-level details. Hence, our primary focus in this article is on addressing this critical limitation.

C. Vision Foundation Models

Vision foundation models have emerged as a groundbreaking approach in artificial intelligence, demonstrating superior learning capabilities compared to traditional models. The segment anything model (SAM) [1] is a recently proposed VFM for image segmentation. Trained on a massive dataset consisting of over 1 billion masks from 11 million images, SAM is built to generalize across diverse segmentation tasks. By performing discriminative self-supervised learning at image and patch levels, DINOv2 [2] learns all-purpose visual features for various downstream tasks. Depth Anything V1 [3] inherits profound semantic priors from the pre-trained DINOv2 backbone via a simple feature alignment constraint, achieving superior performance in downstream tasks. Depth Anything V2 [20] produces finer and more robust depth predictions than the first version by mitigating the distributional shift and limited diversity of synthetic data. Larger models naturally excel at integrating vision features, making them well-suited for our RGB-D driving scene parsing tasks. However, they also come with increased memory requirements and higher training costs. To strike a balance between performance, resource efficiency, and experimental versatility, we select Large DINOv2, Large Depth Anything V1, and Large Depth Anything V2 as the backbones for our proposed HFIT.

III. METHODOLOGY

A. Overall Workflow

As shown in Fig. 2, our HFIT comprises two parts: the first is a plain ViT, which consists of a patch embedding layer followed by L encoder layers, and the second is our proposed HFIT adapter, which contains (1) a DSPE to capture local spatial features from the input RGB-D pairs, (2) RHFF modules to inject recalibrated spatial priors into the ViT, and

(3) HGFI modules to integrate holistic hierarchical features and to reconstruct multi-level and multi-scale features.

First, in the ViT branch, the RGB image $\mathbf{I}^R \in \mathbb{R}^{B \times 3 \times H \times W}$ is fed into the patch embedding layer to extract features with a resolution reduced to 1/16 of the original image. These patches are then flattened and projected into D -dimensional tokens, which, after being added with the position embedding, pass through L encoder layers. Simultaneously, in the adapter branch, the RGB image and its corresponding depth map $\mathbf{I}^D \in \mathbb{R}^{B \times 3 \times H \times W}$ pass through the DSPE to generate a feature pyramid $\{\mathbf{F}_{1, \frac{1}{8}}^S, \mathbf{F}_{1, \frac{1}{16}}^S, \mathbf{F}_{1, \frac{1}{32}}^S\}$. These pyramid features are then flattened and concatenated into $\mathbf{F}_1^S \in \mathbb{R}^{(\sum_{t=2}^4 \frac{HW}{S_t^2}) \times D}$ as input for subsequent feature interaction, where $S_t = 2^{t+1}$ ($t \in [2, 4] \cap \mathbb{Z}$) denotes the corresponding stride number. Specifically, given the number of interactions N (usually $N = 4$), we evenly split the Transformer encoders into N blocks, each containing L/N encoder layers. At the i -th stage, the spatial priors \mathbf{F}_i^S are first injected into the Transformer features $\mathbf{F}_i^V \in \mathbb{R}^{\frac{HW}{16^2} \times D}$ via the RHFF module, which adaptively emphasizes and diminishes complementary information to achieve comprehensive fusion. Then, the HGFI module combines multi-level features and reorganizes multi-scale features from the ViT's single-scale features at the end of each stage. After N bidirectional feature interactions, the heterogeneous features from two branches are aggregated and passed to the back-end decoders for dense prediction tasks.

B. Duplex Spatial Prior Extractor

Recent studies have shown that convolutions can help Transformers better capture local spatial semantics, underscoring their significance in computer vision tasks that necessitate clear object boundaries [47]. Building on this idea, we propose the DSPE module, which is designed to capture local spatial contexts from RGB-D pairs in conjunction with the patch embedding layer. As depicted in Fig. 2, the DSPE employs dual branches to extract spatial priors from RGB images and relative depth maps. By incorporating RGB images as complementary inputs, the DSPE extracts more comprehensive local spatial patterns than using relative depth maps alone. EfficientNet [48], known for its superior performance in extracting detailed visual features, outperforms other hierarchical CNN models such as the ResNet [49] series. Hence, we construct the DSPE using two identical EfficientNet blocks, enabling the efficient extraction of spatial priors from RGB-D pairs.

$\mathbf{I}^R \in \mathbb{R}^{B \times 3 \times H \times W}$ and $\mathbf{I}^D \in \mathbb{R}^{B \times 3 \times H \times W}$ are independently fed into duplex EfficientNet [48] blocks, generating multi-scale feature pyramids $\mathcal{F}_1^R = \{\mathbf{F}_{1, \frac{1}{8}}^R, \mathbf{F}_{1, \frac{1}{16}}^R, \mathbf{F}_{1, \frac{1}{32}}^R\}$ and $\mathcal{F}_1^D = \{\mathbf{F}_{1, \frac{1}{8}}^D, \mathbf{F}_{1, \frac{1}{16}}^D, \mathbf{F}_{1, \frac{1}{32}}^D\}$, where $\mathbf{F}_{1, \frac{1}{2^l}}^{R,D} \in \mathbb{R}^{\frac{H}{S_l} \times \frac{W}{S_l} \times C_l}$ represents the features at the l -th stage, C_l and $S_l = 2^{l+1}$ ($l \in [2, 4] \cap \mathbb{Z}$) denote the corresponding channel and stride numbers, respectively. Subsequently, $\mathbf{F}_{1, \frac{1}{2^l}}^R$ and $\mathbf{F}_{1, \frac{1}{2^l}}^D$ are merged via element-wise summation, followed by 1×1 convolution layers to reduce the output to D channels, yielding heterogeneous features $\mathcal{F}_1^S = \{\mathbf{F}_{1, \frac{1}{8}}^S, \mathbf{F}_{1, \frac{1}{16}}^S, \mathbf{F}_{1, \frac{1}{32}}^S\}$, where $\mathbf{F}_{1, \frac{1}{2^l}}^S \in \mathbb{R}^{\frac{H}{S_l} \times \frac{W}{S_l} \times D}$. These three-scale features are then

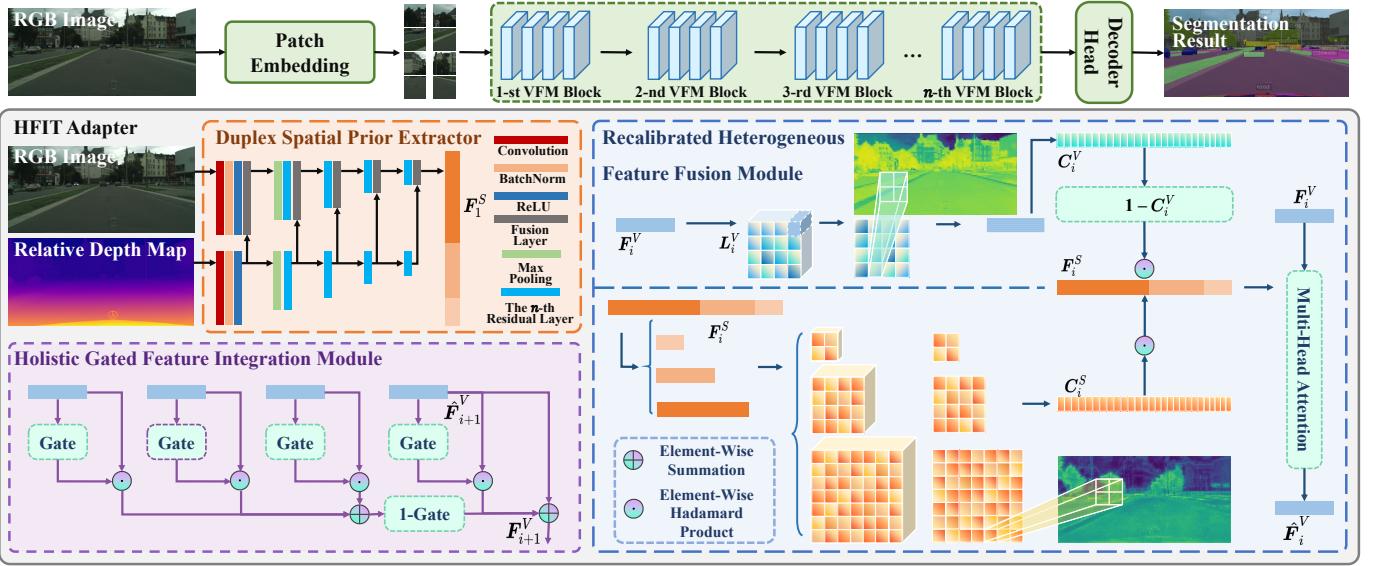


Fig. 2: An illustration of our proposed **HFIT**, consisting of (1) a plain ViT, (2) a duplex spatial prior extractor, (3) recalibrated heterogeneous feature fusion modules, and (4) holistic gated feature integration modules.

flattened and concatenated to form the heterogeneous spatial prior $\mathbf{F}_1^S \in \mathbb{R}^{(\sum_{t=2}^4 \frac{HW}{S_t^2}) \times D}$, which is utilized in subsequent heterogeneous feature fusion stages.

C. Recalibrated Heterogeneous Feature Fusion Module

The primary challenge in feature extraction for RGB-D scene parsing lies in the efficient fusion and integration of diverse heterogeneous features obtained from multiple data sources [24], [50], [51]. As discussed in Sect. II, previous studies [12], [27], [28] combined these heterogeneous features without adequately addressing their inherent differences. To tackle these limitations, RHFF selectively highlights key regions while diminishing the focus on less significant areas. Original heterogeneous spatial priors are adaptively recalibrated through dynamic weighting and seamlessly integrated into the ViT via a multi-scale deformable attention mechanism.

Before injecting heterogeneous features, it is crucial to distinguish and quantify the reliability of each feature. Convolutional layers are initially employed to generate a logit map $\mathbf{L}_i^V \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times C}$ for Transformer features through the following process:

$$\mathbf{L}_i^V = \text{ReLU} \left(\text{BN} \left(\text{Conv}_{1 \times 1} \left(\text{Reshape}(\mathbf{F}_i^V) \right) \right) \right), \quad (1)$$

where “BN(·)” denotes BatchNorm operation, “Reshape” operation reshapes the two-dimensional matrix $\mathbf{F}_i^V \in \mathbb{R}^{\frac{HW}{16^2} \times D}$ into a three-dimensional tensor $\mathbf{F}_i^V \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times D}$, and C denotes the segmentation classes. After modeling interdependencies across channel dimensions, the reliabilities are contrasted to compute the confidence map $\mathbf{C}_{i, \frac{1}{16}}^V \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times 1}$ using the following expression:

$$\mathbf{C}_{i, \frac{1}{16}}^V = \sigma \left(\text{Conv}_{k \times k, r} \left(-\mathbf{L}_i^V \odot \log \mathbf{L}_i^V \right) \right), \quad (2)$$

where \odot represents the element-wise Hadamard product between two metrics, $\sigma(\cdot)$ denotes the Sigmoid operation, and “Conv” denotes atrous convolutional layers with kernel size $k \times k, r$

k and dilation rate r . Each element in \mathbf{C}_i^V indicates the confidence level of a pixel relative to its segmentation class. A higher confidence level suggests the pixel significantly contributes to accurate segmentation, while a lower confidence level implies redundancy, which should be either ignored or supplemented with relevant information. To align the feature pyramid $\mathbf{F}_i^S \in \mathbb{R}^{(\sum_{t=2}^4 \frac{HW}{S_t^2}) \times D}$ across three resolutions, we downsample and upsample $\mathbf{C}_{i, \frac{1}{16}}^V$ to $\mathbf{C}_{i, \frac{1}{8}}^V \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times 1}$ as well as $\mathbf{C}_{i, \frac{1}{32}}^V \in \mathbb{R}^{\frac{H}{32} \times \frac{W}{32} \times 1}$. These features are then concatenated to yield the confidence map \mathbf{C}_i^V as follows:

$$\mathbf{C}_i^V = \left[\mathbf{C}_{i, \frac{1}{8}}^V; \mathbf{C}_{i, \frac{1}{16}}^V; \mathbf{C}_{i, \frac{1}{32}}^V \right] \in \mathbb{R}^{(\sum_{t=2}^4 \frac{HW}{S_t^2}) \times 1}. \quad (3)$$

In the adapter branch, the spatial prior $\mathbf{F}_i^S \in \mathbb{R}^{(\sum_{t=2}^4 \frac{HW}{S_t^2}) \times D}$ is first divided into three-scale feature maps $\mathbf{F}_{i, \frac{1}{2^t}}^S \in \mathbb{R}^{\frac{H}{S_t} \times \frac{W}{S_t} \times D}$. Then, we calculate the confidence map of the heterogeneous spatial prior \mathbf{C}_i^S as follows:

$$\mathbf{L}_{i, \frac{1}{2^t}}^S = \text{ReLU} \left(\text{BN} \left(\text{Conv}_{1 \times 1} \left(\text{Reshape}(\mathbf{F}_{i, \frac{1}{2^t}}^S) \right) \right) \right), \quad (4)$$

$$\mathbf{C}_{i, \frac{1}{2^t}}^S = \text{Reshape}' \left(\sigma \left(\text{Conv}_{k \times k, r} \left(-\mathbf{L}_{i, \frac{1}{2^t}}^S \odot \log \mathbf{L}_{i, \frac{1}{2^t}}^S \right) \right) \right), \quad (5)$$

$$\mathbf{C}_i^S = \left[\mathbf{C}_{i, \frac{1}{8}}^S; \mathbf{C}_{i, \frac{1}{16}}^S; \mathbf{C}_{i, \frac{1}{32}}^S \right] \in \mathbb{R}^{(\sum_{t=2}^4 \frac{HW}{S_t^2}) \times 1}, \quad (6)$$

where “Reshape’” expands the three-dimensional tensor into a two-dimensional matrix. In contrast to previous works like SAGate [45], CFM [46], and DDPM [52], which mainly focus on feature commonality, our RHFF module delves deeper into complementary and redundant features. Since a plain ViT may struggle to confidently segment certain regions, these areas

should be supplemented with heterogeneous spatial priors. Additionally, these priors need to be validated for reliability to avoid introducing redundant or incorrect information. Upon such inspirations, we use these two confidence maps to adaptively recalibrate the heterogeneous spatial priors as follows:

$$\mathbf{F}_i^S = (1 - \mathbf{C}_i^V) \odot \mathbf{C}_i^S \odot \mathbf{F}_i^S, \quad (7)$$

where $(1 - \mathbf{C}_i^V)$ selectively activates spatial priors that are deficient in Transformer features, and \mathbf{C}_i^S emphasizes high-reliability regions within spatial priors.

After adaptively recalibrating and purifying the heterogeneous spatial priors, we utilize a cross-attention mechanism to inject \mathbf{F}_i^S into the Transformer feature \mathbf{F}_i^V . The process can be described as:

$$\hat{\mathbf{F}}_i^V = \mathbf{F}_i^V + \gamma_i \text{MHA} \left(\text{LN} \left(\mathbf{F}_i^V \right), \text{LN} \left(\mathbf{F}_i^S \right) \right), \quad (8)$$

where “ $\text{LN}(\cdot)$ ” denotes LayerNorm operation [53], and “ $\text{MHA}(\cdot)$ ” is a multi-head attention operation [54]. In addition, we introduce a learnable vector $\gamma^i \in \mathbb{R}^D$ [31], [55] to ensure that the feature distribution of $\hat{\mathbf{F}}_i^V$ remains largely unchanged. This allows for effective utilization of the pre-trained weights of the ViT while maintaining stable feature integration.

D. Holistic Gated Feature Integration Module

Integrating low-level features to recover detailed information lost in high-level features is a natural approach [44], [56]. Unfortunately, directly combining them suffers from the semantic gap and will drown useful information [34]. To this end, we propose a holistic integration module to selectively extract features from multiple levels. Specifically, lower-level features, rich in fine-grained details, are used to enhance features at higher levels, while gate mechanisms regulate the flow of relevant information.

After injecting heterogeneous spatial priors into the plain ViT, we obtain the output feature $\hat{\mathbf{F}}_{i+1}^V$ by passing $\hat{\mathbf{F}}_i^V$ through the backbone layers of the i -th block. Then, we utilize the gate mechanism to control the feature flow as follows:

$$\mathbf{F}_{i+1}^V = (1 + \mathbf{G}_{i+1}^V) \odot \hat{\mathbf{F}}_{i+1}^V + (1 - \mathbf{G}_{i+1}^V) \odot \sum_{l=1}^i \mathbf{G}_l^V \odot \hat{\mathbf{F}}_l^V, \quad (9)$$

$$\mathbf{F}_i^S = (1 + \mathbf{G}_i^S) \odot \mathbf{F}_i^S + (1 - \mathbf{G}_i^S) \odot \sum_{l=1, l \neq i}^i \mathbf{G}_l^S \odot \mathbf{F}_l^S. \quad (10)$$

Features from lower levels are propagated to the i -th stage only when $\mathbf{G}_l^{V,S}$ has high values and $\mathbf{G}_i^{V,S}$ has low values. In addition to regulating the flow of useful information to appropriate stages, the gates can effectively prevent information redundancy by updating features only when the current information is insufficient or irrelevant. $\mathbf{G}_l^V \in \mathbb{R}^{\frac{H}{16^2} \times D}$ and $\mathbf{G}_l^S \in \mathbb{R}^{\left(\sum_{t=2}^4 \frac{H}{S_t^2}\right) \times D}$ are calculated as follows:

$$\mathbf{G}_l^V = \text{Reshape}' \left(\sigma \left(\text{Conv}_{1 \times 1} \left(\text{Reshape} \left(\hat{\mathbf{F}}_l^V \right) \right) \right) \right), \quad (11)$$

$$\mathbf{G}_l^S = \text{Reshape}' \left(\sigma \left(\text{DW} \left(\text{Reshape} \left(\mathbf{F}_{i, \frac{1}{2^t}}^S \right) \right) \right) \right), \quad (12)$$

$$\mathbf{G}_i^S = \left[\mathbf{G}_{i, \frac{1}{8}}^S; \mathbf{G}_{i, \frac{1}{16}}^S; \mathbf{G}_{i, \frac{1}{32}}^S \right], \quad (13)$$

where “ $\text{DW}(\cdot)$ ” represents a set of depth-wise convolutions with varying kernel sizes, designed to accommodate the diverse receptive field requirements of different feature groups.

Moreover, semantic representation often exhibits biases due to modality differences. To alleviate this problem, we employ a module comprising a cross-attention layer and a feed-forward network (FFN) to unify spatial priors with Transformer features. This process can be formulated as:

$$\mathbf{F}_{i+1}^S = \hat{\mathbf{F}}_i^S + \text{FFN} \left(\text{LN} \left(\hat{\mathbf{F}}_i^S \right) \right), \quad (14)$$

$$\hat{\mathbf{F}}_i^S = \mathbf{F}_i^S + \text{MHA} \left(\text{LN} \left(\mathbf{F}_i^S \right), \text{LN} \left(\mathbf{F}_{i+1}^V \right) \right). \quad (15)$$

Moreover, unlike traditional Transformer architecture [57], [58], which employs the attention mechanism solely on single-scale features, our HGFI strategy incorporates multi-resolution features at scales of 1/8, 1/16, and 1/32. The generated feature \mathbf{F}_{i+1}^S is then passed to the next RHFF module.

IV. EXPERIMENTS

A. Datasets, Implementation Details, and Evaluation Metrics

We utilize two public datasets to evaluate the performance of our proposed methods: Cityscapes [36] and KITTI Semantics [37] datasets. Our model was trained for 20,000 iterations on an NVIDIA RTX 3090 GPU. Images were randomly cropped to a size of 448×448 pixels during the training process. Standard data augmentation techniques, such as random color adjustments, photometric distortion, rescaling, and flipping, were also used. We employed five widely used metrics to evaluate the scene parsing performance: mean F₁-score (mFsc), mean intersection over union (mIoU), average accuracy (aAcc), mean precision (mPre), and mean recall (mRec). More details are given in the supplementary material.

B. Comparisons with State-of-the-Art Methods

As illustrated in Figs. 3 and 4, HFIT consistently outperforms other scene parsing methods on the KITTI Semantics and Cityscapes datasets. As depicted in Figs. 3 and 4, our HFIT achieves segmentation results with clear boundaries, effectively separating adjacent objects such as vehicles, pedestrians, and buildings. Additionally, HFIT demonstrates significant improvements in segmenting small and distant categories, including streetlights, traffic signs, and fences. In dense, visually complex regions like intersections, HFIT ensures robust semantic consistency, accurately segmenting objects such as sidewalks, vegetation, and poles, while minimizing overlap and misclassification.

Tables I and II highlight the best results in bold, with \uparrow indicating that higher values correspond to improved performance. To ensure a thorough evaluation, we conduct additional experiments by averaging the results from three independent runs for each metric. These results indicate that our proposed HFIT outperforms all other traditional RGB-D methods and single-modal methods, with an increase in mIoU by 17.45–45.54% on the Cityscapes dataset and an increase in mIoU

TABLE I: Quantitative comparisons with SoTA scene parsing methods on the Cityscapes dataset.

Backbone	Type	Methods	mFsc (%) ↑	mIoU (%) ↑	aAcc (%) ↑	mPre (%) ↑	mRec (%) ↑	error range (%) ↓
Traditional	Single-Modal	OCRNet [59]	79.78	67.29	94.99	79.73	77.91	1.23
		KNet [60]	74.43	62.25	93.50	81.58	71.74	0.59
		EMANet [61]	72.10	61.52	93.91	78.85	69.69	0.67
	Data-Fusion	SNE-RoadSeg [10]	54.04	42.52	89.34	67.01	49.78	2.32
		OFFNet [11]	59.22	49.16	64.33	71.21	53.41	1.54
		MFNet [12]	51.50	39.20	89.39	62.69	46.56	0.81
		FuseNet [13]	53.28	40.40	89.79	63.63	49.57	2.03
VFM	Single-Modal w/o Adapter	DINOv2 [2]	88.11	80.23	96.12	89.29	88.91	0.52
		Depth Anything V1 [3]	89.48	81.11	96.30	90.26	88.66	0.29
		Depth Anything V2 [20]	89.00	80.31	96.03	89.51	88.27	0.41
	Single-Modal w/ Adapter	ViT-Adapter [25]	89.69	82.11	96.53	90.91	89.31	0.35
		ViT-CoMer [31]	88.84	80.02	96.15	89.21	88.01	0.61
	Data-Fusion w/ Adapter (ours)	DINOv2 [2]	91.32	84.53	97.15	91.42	91.19	0.42
		Depth Anything V1 [3]	91.64	84.74	97.09	92.27	90.26	0.64
		Depth Anything V2 [20]	91.06	84.59	96.94	91.87	90.49	0.52

TABLE II: Quantitative comparisons with SoTA scene parsing methods on the KITTI Semantics dataset.

Backbone	Type	Methods	mFsc (%) ↑	mIoU (%) ↑	aAcc (%) ↑	mPre (%) ↑	mRec (%) ↑	error range (%) ↓
Traditional	Single-Modal	OCRNet [59]	79.77	62.24	93.26	81.40	68.70	0.54
		KNet [60]	74.56	56.47	91.22	84.74	62.79	0.72
		EMANet [61]	67.39	57.03	90.42	80.28	63.71	0.67
	Data-Fusion	SNE-RoadSeg [10]	67.74	41.28	89.92	70.39	50.81	1.03
		OFFNet [11]	55.12	39.24	91.65	67.47	46.26	1.23
		MFNet [12]	58.55	37.97	88.65	69.70	42.98	0.64
		FuseNet [13]	50.91	34.12	84.38	50.94	42.68	0.75
VFM	Single-Modal w/o Adapters	DINOv2 [2]	86.82	78.63	95.49	87.84	86.39	0.34
		Depth Anything V1 [3]	86.62	77.85	95.63	89.41	84.66	0.22
		Depth Anything V2 [20]	87.05	79.30	95.77	89.61	86.04	0.45
	Single-Modal w/ Adapters	ViT-Adapter [25]	87.28	79.28	95.91	88.27	87.23	0.23
		ViT-CoMer [31]	85.13	76.54	95.47	88.93	83.77	0.44
	Data-Fusion w/ Adapters (ours)	DINOv2 [2]	87.58	79.91	96.22	89.41	86.88	0.52
		Depth Anything V1 [3]	88.07	80.24	95.82	89.37	87.27	0.51
		Depth Anything V2 [20]	87.23	79.38	95.88	89.13	86.25	0.63

by up to 46.12% on the KITTI Semantics dataset. Compared to VFM-based algorithms, HFIT demonstrates greater performance, delivering a 2.63-4.72% increase in mIoU on the Cityscapes dataset and a 0.94-2.39% improvement on the KITTI dataset. These results underscore HFIT's superior visual representation capabilities and reliability in real-world driving scenarios. They also validate the hypothesis proposed in Sect. I, demonstrating that incorporating heterogeneous spatial information significantly enhances the scene parsing performance of VFs.

Compared to traditional RGB-D methods, single-modal VFs demonstrate significant advantages, largely due to their profound prior knowledge, which enables them to achieve competitive results in constrained settings. Additionally, we select ViT-Adapter [25] and ViT-CoMer [31] as the baselines

of single-modal VFs with adapters. ViT-Adapter enhances ViT's ability to capture multi-scale information by adding spatial prior, injector, and extractor modules. ViT-CoMer efficiently handles multi-scale interactions, balancing complex scene comprehension with memory efficiency. Notably, ViT-Adapter [25] outperforms standard VFs, attributed to its dynamic fine-tuning mechanism, which allows for better refinement of feature representations. Nevertheless, ViT-CoMer's [31] performance is underwhelming. This may be due to its bidirectional summation feature fusion mechanism, which fails to bridge the semantic gap between two distinct feature types.

Additionally, we also validate the compatibility of HFIT with various cutting-edge VFs. Among them, Depth Anything V1 [3] achieves the best performance, with the highest mIoU of 84.74%. However, an interesting observation emerged

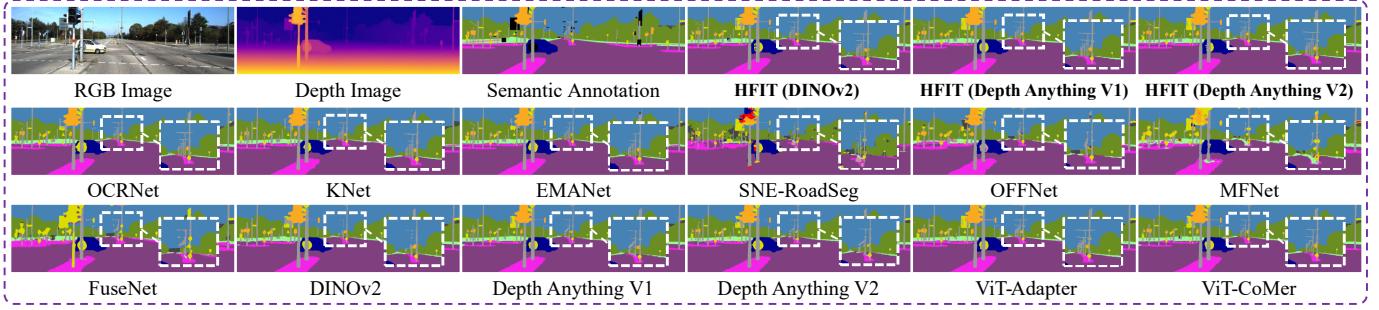


Fig. 3: Qualitative comparisons with SoTA scene parsing approaches on the KITTI Semantics [37] dataset.

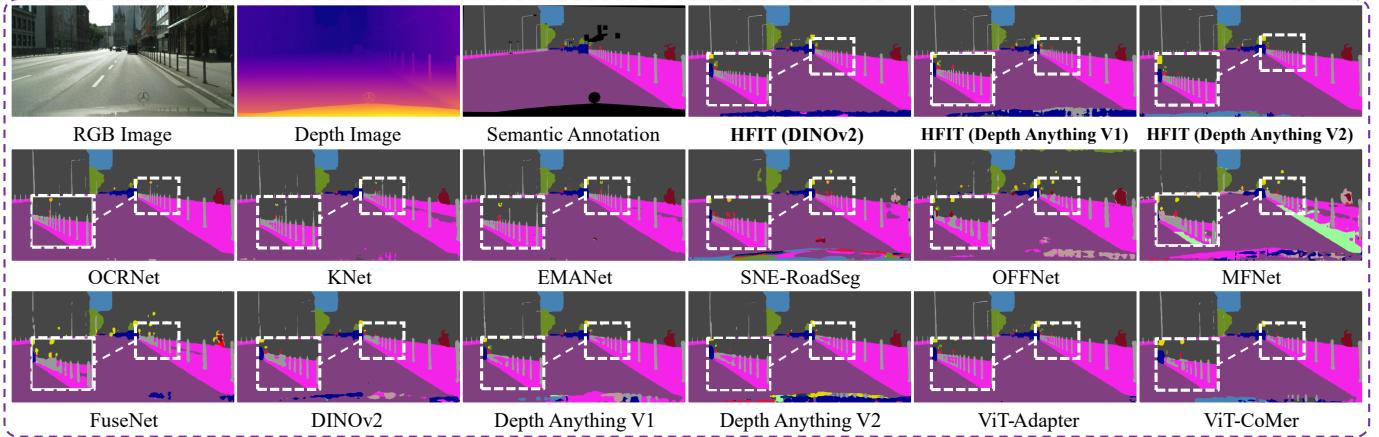


Fig. 4: Qualitative comparisons with SoTA scene parsing approaches on the Cityscapes [36] dataset.

TABLE III: Ablation Study on DSPE with different inputs.

Input Type	mFsc (%) ↑	mIoU (%) ↑	aAcc (%) ↑
RGB	90.39	83.07	96.76
Depth	90.53	83.85	96.74
Relative Depth	90.42	83.82	96.78
RGB+Depth	91.02	84.10	96.99
RGB+Relative Depth	91.11	84.22	96.97

during our experiments. Despite the evaluation matrices having comparable values on the Cityscapes [36] and KITTI Semantics [37] datasets, the performance improvements achieved by HFIT compared to VFM-based methods are relatively less significant on the KITTI Semantics dataset. This performance gap can be attributed to KITTI's smaller dataset size and limited diversity, which restricts the ability of large models like HFIT to fully utilize their representational capacity. As a result, the model is more prone to overfitting and exhibits reduced generalization on the KITTI dataset.

C. Ablation Studies

As shown in Table III, we first evaluate HFIT's performance using different input strategies. Single-modal configurations, using either RGB or depth images alone, achieve mIoU scores of 83.07% and 83.85%, respectively. Inputting both RGB and depth images achieves superior performance, with an mIoU of

TABLE IV: Ablation Study on DSPE blocks.

Input Type	Params (M)	mFsc (%) ↑	mIoU (%) ↑	aAcc (%) ↑
ResNet-18	23.34	90.55	83.31	96.87
ResNet-34	43.56	90.78	83.68	96.90
ResNet-50	50.95	90.81	83.76	96.90
Swin-T	56.51	90.76	83.65	96.86
Swin-S	99.15	90.84	83.78	96.88
Swin-B	175.45	90.69	83.56	96.90
EfficientNet-B0	6.35	90.53	83.32	96.81
EfficientNet-B1	9.21	90.69	83.59	96.91
EfficientNet-B7	79.09	90.85	83.82	96.89

TABLE V: Ablation Study on RHFF Modules.

RGB Weight	Depth Weight	mFsc (%) ↑	mIoU (%) ↑	aAcc (%) ↑
✓		90.24	82.83	96.79
	✓	90.31	82.98	96.77
	✓	90.43	83.14	96.83
✓	✓	90.96	84.01	96.97

84.10%, validating the importance of modality complementation. Notably, the performance difference between using scaled metric depth obtained from RAFT-Stereo [62] and relative

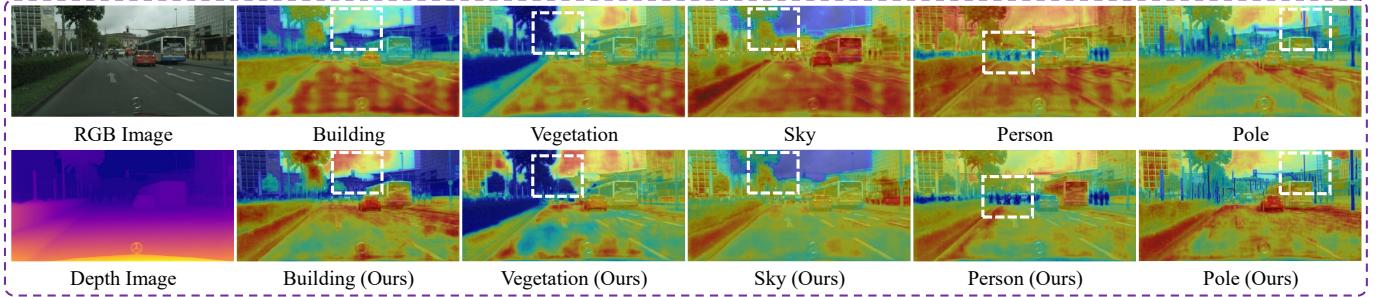


Fig. 5: Probability maps for different classes generated by VFM and HFIT, where blue indicates high prediction confidence and red represents low prediction confidence.

TABLE VI: Ablation Study on HGFI Modules.

HGFI-ViT	HGFI-Adapter	mFsc (%) ↑	mIoU (%) ↑	aAcc (%) ↑
✓		89.91	82.36	96.73
	✓	90.44	83.17	96.80
	✓	90.30	82.96	96.77
✓	✓	91.04	84.11	96.94

depth derived from the VFM is minimal, within a margin of 0.12%. This observation further supports the hypothesis presented in Sect. I.

To investigate the most effective structure for the DSPE, we substitute the basic blocks with alternatives, as detailed in Table IV. Compared to earlier models like the ResNet [49] series and the more recent Swin Transformer [57] series, EfficientNet [48] demonstrates superior performance, achieving the highest mIoU of 83.82%. This suggests that the lightweight design of EfficientNet is well-suited for effectively extracting spatial priors.

Theoretically, higher confidence and greater complementarity in features should result in improved fusion. To validate this hypothesis, we experimentally evaluate the contributions of the RGB weights ($1 - C_i^V$) and depth weights C_i^S in the RHFF module. As expected, removing either component leads to a significant decline in HFIT's performance (see Table V). Moreover, we compare the probability maps produced by VFM and HFIT. As shown in Fig. 5, the probability maps generated by HFIT demonstrate more reliable and confident predictions for each category, resulting in clearer and sharper probability distributions. This improvement is particularly evident for categories such as buildings, riders, and cars, especially when compared to predictions at neighboring pixels.

We also investigate the impact of holistic gated feature integration on both Transformer and CNN features. As presented in Table VI, the baseline setup, which relies solely on the cross-attention mechanism for feature fusion, achieves an mIoU that is 1.75% lower than the full HGFI strategy. As expected, the holistic gated feature integration of ViT features \hat{F}_{i+1}^V and spatial priors F_i^S from the side adapter significantly enhances HFIT's performance.

V. CONCLUSION AND FUTURE WORK

This article discussed a new challenge in computer vision: fully exploiting profound prior knowledge of VFM for RGB-D driving scene parsing. Our contributions can be summarized in four key aspects: the successful development of an RGB-D scene parsing structure and the introduction of three novel components in the side adapter—DSPE, RHFF, and HGFI. Experimental results demonstrated the significant advantages of HFIT over existing single-modal and data-fusion algorithms, emphasizing its effectiveness and potential for autonomous driving applications.

Looking ahead, advancements in multi-modal VFM for autonomous driving present new opportunities for robust scene understanding and decision-making. Incorporating emerging techniques like LiDAR-LLM [63] could enable HFIT to simultaneously process LiDAR and depth-based cues, enhancing 3D spatial comprehension in driving environments. Additionally, integrating large language models could improve HFIT's contextual understanding by interpreting ambiguous or complex visual inputs through language-informed perspectives. Furthermore, drawing inspiration from DriveLM's [64] graph-based reasoning for sequential decision-making, a promising future direction for HFIT lies in its evolution from a static model to a dynamic system capable of localization, prediction, and action planning. Such advancements would not only increase HFIT's adaptability and responsiveness to real-time challenges but also pave the way for smarter and more interactive autonomous driving systems. Further discussions are provided in the supplementary materials.

REFERENCES

- [1] A. Kirillov *et al.*, “Segment Anything,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 4015–4026.
- [2] M. Oquab *et al.*, “DINOv2: Learning robust visual features without supervision,” *Computing Research Repository (CoRR)*, vol. abs/2304.07193, 2023. [Online]. Available: <https://arxiv.org/abs/2304.07193>
- [3] L. Yang *et al.*, “Depth Anything: Unleashing the power of large-scale unlabeled data,” in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 10 371–10 381.
- [4] Y. Feng *et al.*, “ViPOcc: leveraging visual priors from vision foundation models for single-view 3D occupancy prediction,” in *Proceedings of the AAAI conference on Artificial Intelligence (AAAI)*, 2025, in press.
- [5] G. Han and S.-N. Lim, “Few-shot object detection with foundation models,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 28 608–28 618.

- [6] C.-W. Liu *et al.*, “Playing to vision foundation model’s strengths in stereo matching,” *IEEE Transactions on Intelligent Vehicles*, 2024, DOI: 10.1109/TIV.2024.3467287.
- [7] M. Shvets *et al.*, “Joint depth prediction and semantic segmentation with multi-view SAM,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2024, pp. 1328–1338.
- [8] Z. Chen *et al.*, “Bridging the domain gap: Self-supervised 3D scene understanding with foundation models,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 36, 2024.
- [9] S. Guo *et al.*, “LIX: Implicitly infusing spatial geometric prior knowledge into visual semantic segmentation for autonomous driving,” *IEEE Transactions on Image Processing*, vol. 34, pp. 7250–7263, 2025.
- [10] R. Fan *et al.*, “SNE-RoadSeg: Incorporating surface normal information into semantic segmentation for accurate freespace detection,” in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 340–356.
- [11] C. Min *et al.*, “ORFD: A dataset and benchmark for OFF-Road freespace detection,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2532–2538.
- [12] Q. Ha *et al.*, “MFNet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 5108–5115.
- [13] Hazirbas *et al.*, “FuseNet: Incorporating depth into semantic segmentation via fusion-based cnn architecture,” in *Proceedings of the Asian Conference on Computer Vision (ACCV)*. Springer, 2017, pp. 213–228.
- [14] R. Fan *et al.*, “Pothole detection based on disparity transformation and road surface modeling,” *IEEE Transactions on Image Processing*, vol. 29, pp. 897–908, 2020.
- [15] R. Fan *et al.*, “Graph attention layer evolves semantic segmentation for road pothole detection: A benchmark and algorithms,” *IEEE Transactions on Image Processing*, vol. 30, pp. 8144–8154, 2021.
- [16] C.-W. Liu *et al.*, “These maps are made by propagation: Adapting deep stereo networks to road scenarios with decisive disparity diffusion,” *IEEE Transactions on Image Processing*, vol. 34, pp. 1516–1528, 2025.
- [17] Z. Wu *et al.*, “S³M-Net: Joint learning of semantic segmentation and stereo matching for autonomous driving,” *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 2, pp. 3940–3951, 2024.
- [18] Z. Huang *et al.*, “Online, target-free LiDAR-camera extrinsic calibration via cross-modal mask matching,” *IEEE Transactions on Intelligent Vehicles*, vol. 10, no. 5, pp. 3531–3542, 2025.
- [19] Z. Huang *et al.*, “Environment-driven online LiDAR-camera extrinsic calibration,” *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 24 164–24 176, 2025.
- [20] Yang, Lihe *et al.*, “Depth Anything V2,” *Computing Research Repository (CoRR)*, vol. abs/2406.09414, 2024. [Online]. Available: <https://arxiv.org/abs/2406.09414>
- [21] L. Hoyer *et al.*, “Three ways to improve semantic segmentation with self-supervised depth estimation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 11 130–11 140.
- [22] S. Saha *et al.*, “Learning to relate depth and semantics for unsupervised domain adaptation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 8197–8207.
- [23] D. Seichter *et al.*, “Efficient RGB-D semantic segmentation for indoor scene analysis,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 525–13 531.
- [24] J. Wang *et al.*, “Learning common and specific features for RGB-D semantic segmentation with deconvolutional networks,” in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2016, pp. 664–679.
- [25] Z. Chen *et al.*, “Vision Transformer adapter for dense predictions,” in *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
- [26] H. Chen *et al.*, “Conv-Adapter: Exploring parameter efficient transfer learning for convnets,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 1551–1561.
- [27] H. Zhou *et al.*, “CANet: Co-attention network for RGB-D semantic segmentation,” *Pattern Recognition*, vol. 124, p. 108468, 2022.
- [28] X. Zhou *et al.*, “FANet: Feature aggregation network for RGBD saliency detection,” *Signal Processing: Image Communication*, vol. 102, p. 116591, 2022.
- [29] J. Li *et al.*, “RoadFormer: Duplex Transformer for RGB-Normal semantic road scene parsing,” *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 7, pp. 5163–5172, 2024.
- [30] J. Huang *et al.*, “RoadFormer+: Delivering RGB-X scene parsing through scale-aware information decoupling and advanced heterogeneous feature fusion,” *IEEE Transactions on Intelligent Vehicles*, 2024, DOI: 10.1109/TIV.2024.3448251.
- [31] C. Xia *et al.*, “ViT-CoMer: Vision Transformer with convolutional multi-scale feature interaction for dense predictions,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 5493–5502.
- [32] Z. Wu *et al.*, “Robust RGB-D fusion for saliency detection,” in *2022 International Conference on 3D Vision (3DV)*. IEEE, 2022, pp. 403–413.
- [33] P. Sun *et al.*, “Deep RGB-D saliency detection with depth-sensitive attention and automatic multi-modal fusion,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 1407–1417.
- [34] X. Li *et al.*, “Gated fully fusion for semantic segmentation,” in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, vol. 34, no. 07, 2020, pp. 11 418–11 425.
- [35] Z. Zhou *et al.*, “UNet++: Redesigning skip connections to exploit multi-scale features in image segmentation,” *IEEE Transactions on Medical Imaging*, vol. 39, no. 6, pp. 1856–1867, 2019.
- [36] M. Cordts *et al.*, “The CityScapes dataset for semantic urban scene understanding,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 3213–3223.
- [37] M. Menze and A. Geiger, “Object scene flow for autonomous vehicles,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 3061–3070.
- [38] Y. Zhang *et al.*, “Deep multimodal fusion for semantic image segmentation: A survey,” *Image and Vision Computing*, vol. 105, p. 104042, 2021.
- [39] A. Valada *et al.*, “AdapNet: Adaptive semantic segmentation in adverse environmental conditions,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4644–4651.
- [40] Y. Cheng *et al.*, “Locality-sensitive deconvolution networks with gated fusion for RGB-D indoor semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 3029–3037.
- [41] E. Hassan *et al.*, “Learning feature fusion in deep learning-based object detector,” *Journal of Engineering*, vol. 2020, no. 1, p. 7286187, 2020.
- [42] W. Zhou *et al.*, “ECFFNet: Effective and consistent feature fusion network for RGB-T salient object detection,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 3, pp. 1224–1235, 2021.
- [43] H. Jiang *et al.*, “Salient object detection: A discriminative regional feature integration approach,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 2083–2090.
- [44] Z. Zhang *et al.*, “ExFuse: Enhancing feature fusion for semantic segmentation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 269–284.
- [45] X. Chen *et al.*, “Bi-directional cross-modality feature propagation with separation-and-aggregation gate for RGB-D semantic segmentation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 561–577.
- [46] J. Wei *et al.*, “F³Net: fusion, feedback and focus for salient object detection,” in *Proceedings of the AAAI conference on Artificial Intelligence (AAAI)*, vol. 34, no. 07, 2020, pp. 12 321–12 328.
- [47] K. Yuan *et al.*, “Incorporating convolution designs into visual Transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 579–588.
- [48] M. Tan and Q. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in *International Conference on Machine Learning (ICML)*. PMLR, 2019, pp. 6105–6114.
- [49] K. He *et al.*, “Deep residual learning for image recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [50] L. Sun *et al.*, “Real-time fusion network for RGB-D semantic segmentation incorporating unexpected obstacle detection for road-driving images,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5558–5565, 2020.
- [51] G. Xu *et al.*, “THCANet: Two-layer hop cascaded asymptotic network for robot-driving road-scene semantic segmentation in RGB-D images,” *Digital Signal Processing*, vol. 136, p. 104011, 2023.
- [52] Y. Pang *et al.*, “Hierarchical dynamic filtering network for RGB-D salient object detection,” in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 235–252.

- [53] J. L. Ba *et al.*, “Layer normalization,” *Computing Research Repository (CoRR)*, vol. abs/1607.06450, 2016. [Online]. Available: <https://arxiv.org/abs/1607.06450>
- [54] B. Cheng *et al.*, “Masked-attention mask Transformer for universal image segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 1290–1299.
- [55] S. Huang *et al.*, “FaPN: Feature-aligned pyramid network for dense image prediction,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 864–873.
- [56] H. Fu and G. Qiu, “Integrating low-level and semantic features for object consistent segmentation,” *Neurocomputing*, vol. 119, pp. 74–81, 2013.
- [57] Z. Liu *et al.*, “Swin Transformer: hierarchical vision Transformer using shifted windows,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 10012–10022.
- [58] E. Xie *et al.*, “SegFormer: Simple and efficient design for semantic segmentation with Transformers,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 34, pp. 12077–12090, 2021.
- [59] Y. Yuan *et al.*, “Object-contextual representations for semantic segmentation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 173–190.
- [60] W. Zhang *et al.*, “K-Net: Towards unified image segmentation,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 34, pp. 10326–10338, 2021.
- [61] X. Li *et al.*, “Expectation-maximization attention networks for semantic segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 9167–9176.
- [62] L. Lipson *et al.*, “RAFT-Stereo: Multilevel recurrent field transforms for stereo matching,” in *2021 International Conference on 3D Vision (3DV)*, 2021, pp. 218–227.
- [63] S. Yang *et al.*, “Lidar-LLM: Exploring the potential of large language models for 3D lidar understanding,” *Computing Research Repository (CoRR)*, vol. abs/2312.14074, 2023. [Online]. Available: <https://arxiv.org/abs/2312.14074>
- [64] C. Sima *et al.*, “DriveLM: Driving with graph visual question answering,” *Computing Research Repository (CoRR)*, vol. abs/2312.14150, 2023. [Online]. Available: <https://arxiv.org/abs/2312.14150>



Sicen Guo (Graduate Student Member, IEEE) received the B.E. degree in automation from Tongji University, Shanghai, China, in 2022. She is currently pursuing her Ph.D. degree with the MIAS Group at Tongji University. Her research interests include stereo matching and semantic segmentation.



Tianyou Wen (Undergraduate Student Member, IEEE) is pursuing his B.E. degree in Electrical Engineering at Tongji University. His interest lies in semantic segmentation. He is also dedicated to using advanced spatial perception technologies to address real-world challenges.



Chuang-Wei Liu (Graduate Student Member, IEEE) received his B.E. degree in automation from Tongji University in 2020. He is currently pursuing his Ph.D. degree, supervised by Prof. Rui Fan, with the MIAS Group at Tongji University. His research interest lies in computer stereo vision, especially for unsupervised and long-term learning.



Cijun Chen (Senior Member, IEEE) received the B.S. degree in automation from Huazhong University of Science and Technology, Wuhan, China, in 1987, the M.S. degree in information and control engineering from Xi'an Jiaotong University, Xi'an, China, in 1990, and the Ph.D. degree in control theory and control engineering from Tongji University, Shanghai, China, in 1999. He is currently a Full Professor with the College of Electronics and Information Engineering at Tongji University. His research interest lies in robotics.



Rui Fan (Senior Member, IEEE) received the B.Eng. degree in Automation from the Harbin Institute of Technology in 2015 and the Ph.D. degree in Electrical and Electronic Engineering from the University of Bristol in 2018. He worked as a Research Associate at the Hong Kong University of Science and Technology from 2018 to 2020 and a Postdoctoral Scholar-Employee at the University of California San Diego between 2020 and 2021. He began his faculty career as a Full Research Professor in the College of Electronics & Information Engineering at Tongji University in 2021. He was promoted to Full Professor in 2022 and attained tenure in 2024, both in the same college and at the Shanghai Research Institute for Intelligent Autonomous Systems. His research interests include computer vision, deep learning, and robotics, with a specific focus on humanoid visual perception under the two-streams hypothesis. Prof. Fan served as an associate editor for ICRA'23/25 and IROS'23/24, an area chair for ICIP'24, and a senior program committee member for AAAI'23/24/25/26. He organized several impactful workshops and special sessions in conjunction with WACV'21, ICIP'21/22/23, ICCV'21/25, and ECCV'22. He was honored by being included in the Stanford University List of Top 2% Scientists Worldwide between 2022 and 2025, recognized on the Forbes China List of 100 Outstanding Overseas Returnees in 2023, acknowledged as one of Xiaomi Young Talents in 2023, and awarded the Shanghai Science & Technology 35 Under 35 honor in 2024 as its youngest recipient.