# Efficient Multi-Task Scene Analysis with RGB-D Transformers

Söhnke Benedikt Fischedick, Daniel Seichter, Robin Schmidt, Leonard Rabes, and Horst-Michael Gross

Ilmenau University of Technology, Neuroinformatics and Cognitive Robotics Lab

98684 Ilmenau, Germany

soehnke.fischedick@tu-ilmenau.de, ORCID: 0000-0001-8447-0584

Abstract—Scene analysis is essential for enabling autonomous systems, such as mobile robots, to operate in real-world environments. However, obtaining a comprehensive understanding of the scene requires solving multiple tasks, such as panoptic segmentation, instance orientation estimation, and scene classification. Solving these tasks given limited computing and battery capabilities on mobile platforms is challenging. To address this challenge, we introduce an efficient multi-task scene analysis approach, called EMSAFormer, that uses an RGB-D Transformer-based encoder to simultaneously perform the aforementioned tasks. Our approach builds upon the previously published EMSANet. However, we show that the dual CNN-based encoder of EMSANet can be replaced with a single Transformer-based encoder. To achieve this, we investigate how information from both RGB and depth data can be effectively incorporated in a single encoder. To accelerate inference on robotic hardware, we provide a custom NVIDIA TensorRT extension enabling highly optimization for our EMSAFormer approach. Through extensive experiments on the commonly used indoor datasets NYUv2, SUNRGB-D, and ScanNet, we show that our approach achieves state-of-the-art performance while still enabling inference with up to 39.1 FPS on an NVIDIA Jetson AGX Orin 32 GB.

# I. INTRODUCTION

A broad understanding of the scene is crucial for mobile agents to operate autonomously in indoor environments. Traditional approaches often only provide semantic or instance information, which is not sufficient for many realworld applications. For example, in our research projects CO-HUMANICS and MORPHIA [1], a mobile robot should autonomously operate in indoor environments and should also be controlled by an operator from a remote location. To enable such a remote control to inexperienced operators, a more intuitive way for navigating is required. As shown in Fig. 1, the operator should be able to click onto a specific point or object visible in the camera image of the current surrounding, to which the robot should automatically navigate to. For example, the mobile robot should drive to a chair within a group of many chairs while respecting the chair's orientation to not block it. Furthermore, the operator should be able to select an entire room to which the robot should drive to. To achieve this, a broader scene understanding is required. The robot needs to combine information from multiple tasks, i.e., semantic and instance segmentation (panoptic segmentation),

This work has received funding from Carl-Zeiss-Stiftung to the project *Co-Presence of Humans and Interactive Companions for Seniors* (CO-HUMANICS).



Fig. 1. Application (bottom) of our proposed efficient multi-task scene analysis approach with an RGB-D Transformer encoder, called EMSAFormer, that simultaneously performs panoptic segmentation, instance orientation estimation, and scene classification (top). See Fig. 2 for prediction colors.

instance orientation estimation, and scene classification. Due to the limited computing and battery capabilities of a mobile robot, a multi-task approach should be preferred.

To gather information for all aforementioned tasks, we introduce an efficient multi-task scene analysis approach utilizing a Transformer-based encoder (EMSAFormer). It builds on top of our Efficient Multi-task Scene Analysis Network (EM-SANet) [2] that already realizes such a system. In [2], we have shown the effectiveness of using RGB and depth as complementary data to improve the performance of individual tasks. For processing the different modalities, a dual ResNet34 encoder is used. The approach enables real-time application (in our application scenario  $\geq 20$  FPS) on an NVIDIA Jetson AGX Xavier, while still reaching state-of-the-art performance. With the release of the NVIDIA Jetson AGX Orin 32 GB, computing capabilities have been increased, i.e., it is now possible to apply larger models in real time. Compared to NVIDIA Jetson AGX Xavier, inference throughput has almost doubled with the same power consumption. However, the authors of EMSANet have shown that a larger encoder, e.g., utilizing two ResNet101 backbones, only slightly improves performance. Motivated by the upcoming usage of Transformer-based architectures for computer vision, we address this limitation by replacing the dual encoder based on convolutional neural networks (CNNs) with a single Swin Transformer for processing RGB-D data. We conduct detailed experiments to show the effectiveness of using a Swin Transformer for processing RGB and depth data and how it compares to the traditional CNN-based encoders. Furthermore, we address differences and challenges when incorporating a Transformer-based encoder. Our experiments show that our EMSAFormer approach is able to outperform the state-of-the-art method EMSANet in most tasks. Finally, to enable inference in real time on the NVIDIA Jetson AGX Orin on our mobile robot, we provide a custom NVIDIA TensorRT extension that greatly accelerates inference. In summary, our main contributions include:

- A complementary RGB-D encoder approach that replaces the dual encoder in EMSANet with a novel single Swin Transformer encoder, while still effectively incorporating information from both RGB and depth data due to a modified Transformer design.
- A custom NVIDIA TensorRT extension for inference acceleration that enables using Swin Transformers as a general-purpose backbone for downstream tasks with import from Open Neural Network Exchange (ONNX) format and arbitrary input resolution.
- Detailed quantitative and qualitative experiments on the common indoor datasets NYUv2 [3], SUNRGB-D [4], and ScanNet [5] demonstrating the applicability and state-of-the-art performance of our approach.

Our code as well as the network weights are publicly available at: https://github.com/TUI-NICR/EMSAFormer.

# II. RELATED WORK

In the following, we summarize related work for scene analysis with focus on Transformer-based encoders and RGB-D processing. We give a brief introduction how the individual tasks can be accomplished in an encoder-decoder approach.

#### A. Transformer-based Encoders

In recent years, convolutional neural networks (CNNs) have been the dominant approach for various computer vision tasks, such as image classification, object detection, or semantic segmentation. Various backbones, such as ResNet [6], EfficientNet [7], or ConvNeXt [8] have been proposed and achieve state-of-the-art performance on a variety of benchmarks. Inspired by the success of Transformers for natural language processing (NLP) [9]-[11], the work of [12] introduces the Vision Transformer (ViT), which is able to reach similar performance to CNNs for image classification. While the proposed architecture achieves impressive results, it still requires extensive pretraining for comparable performance on the ImageNet benchmark [13] and does not realize a pyramid structure, making it less suitable as a general-purpose backbone [14], [15]. Motivated by the success of ViT, various Transformerbased architectures, such as Pyramid Vision Transformer [16], Swin Transformers [14], [17], and SegFormer [15] have been introduced, improving data efficiency and enabling the usage as general-purpose backbone due to the pyramid-like structure. These architectures achieve state-of-the-art performance on a variety of benchmarks; however, they have been proposed for processing RGB data and have not yet been extended to processing RGB-D data, which is the focus of this work.

#### B. RGB-D Encoders

Combining RGB and depth data can improve performance in various computer vision tasks as both provide complementary features [18]–[20]. RGB data provide information about semantic features, such as object color and texture, while depth data provide geometric information about the scene. Handling multiple modalities can be challenging as they contain different features with deviating statistics and characteristics. Thus, multiple approaches have been emerged.

The majority of approaches [2], [18], [19], [21]–[26] handles both modalities in two separated encoders and fuse features using a dedicated fusion mechanism. The fusion is mostly done after each resolution stage [2], [18], [19], [26], [27] of the encoders, at the end of the encoders [28], or into a third encoder handling joint features [20], [23], [29], [30]. Moreover, most approaches additionally fuse features from the encoder to the decoder [2], [19], [26] similar to DeepLabV3+ [31]. While enabling independent processing of both modalities, this method implies a crucial design decision, i.e., it introduces the difficulty of deciding where the fusion should happen and how the features are fused. Additionally, depending on the used backbones, a dual-encoder design often leads to increased computational cost, making these approaches less suitable for embedded hardware.

Other approaches try to handle both modalities in a single encoder [32]–[36]. This is done by using specially tailored convolutions for incorporating depth data. However, such approaches often lack optimization and, thus, are less suitable for embedded hardware.

Motivated by the performance achieved by Swin Transformers [14], [17], the authors of OMNIVORE [37] propose a method for handling multiple modalities in a single encoder. However, as the approach mainly focuses on handling many different modalities and large backbones, it again lacks optimization and efficiency. Moreover, OMNIVORE requires heavy pretraining to achieve good performance.

In this paper, we follow the recent trend of using Transformer-based architectures. However, in contrast to the aforementioned approaches, we focus on modifying the Swin Transformer architecture to efficiently incorporate depth information while still relying on a single encoder.

# C. Task Decoders

The design of the decoder depends on the task to solve. Scene classification, i.e., assigning a scene label, such as living room or office, to the entire input, is similar to other classification tasks. It only requires a classification layer at the end of the encoder. By contrast, pixel-wise dense prediction tasks require more sophisticated decoders. As the encoder typically lowers the spatial resolution, dense-prediction decoders often incorporate multiple modules to gradually restore the resolution. Most approaches use a CNN-based decoder [26], [31], [38]–[40] of varying complexity. With the rise of Transformerbased architectures, SegFormer [15] proposes another more lightweight MLP-based decoder. Features from different stages are embedded using a fully-connected layer, upsampled to the same spatial resolution, concatenated, and passed through two additional fully-connected layers to encode the final prediction.

For panoptic segmentation, at least one dense-prediction decoder is required. Panoptic segmentation [41] combines semantic and instance segmentation and aims at assigning a semantic class to each pixel of the input image as well as a unique instance ID to each pixel belonging to a distinguishable instance. Panoptic segmentation can be done in a top-down, bottom-up, or end-to-end manner. Top-Down methods are typically based on existing approaches for instance segmentation and extend them with an additional decoder for semantic segmentation [42], [43]. While achieving state-of-the-art performance, these architectures typically feature sophisticated network architectures and require further logic to resolve overlapping instance masks. Bottom-up methods [2], [44], [45], on the other hand, are often based on encoder-decoder architectures for semantic segmentation and extend them by an additional decoder for instances. This additional decoder predicts individual instances by grouping corresponding pixels into clusters. As there are already efficient architectures for semantic segmentation [26], [40], bottom-up approaches are often more efficient than top-down approaches [2]. However, both approaches require an additional post-processing step for

combining instance and semantic segmentation. By contrast, end-to-end approaches, such as MaX-DeepLab [46], directly predict the panoptic segmentation without additional post-processing. While already achieving great performance, these methods are not established yet and also require complex architectures that currently do not focus on efficiency.

Instance orientation estimation can also be done in multiple ways. For extracting the orientation, patch-based methods, such as [47]–[49], can be used. Another way for estimating the orientation of an instance is to estimate the pose of its 3D bounding box as shown in [50]–[52]. In contrast to the aforementioned approaches, EMSANet [2] follows the bottom-up idea and incorporates orientation estimation into the instance decoder in a dense-prediction manner. This way, the computational overhead for orientation estimation is limited and averaging multiple predictions for an instance is enabled.

In this paper, we follow the bottom-up design of EMSANet for tackling the scene analysis tasks. However, to further speed up inference, we examine the lightweight MLP-based decoder of SegFormer.

# III. EFFICIENT MULTI-TASK SCENE ANALYSIS WITH RGB-D TRANSFORMERS

Our Transformer-based approach for scene analysis (EM-SAFormer) is shown in Fig. 2. It builds on top of the EM-SANet [2]. Both architectures share a similar encoder-decoder design and the same task encoding. The encoder extracts



Fig. 2. Architecture of our proposed efficient multi-task scene analysis approach with a single RGB-D Transformer encoder (EMSAFormer) that simultaneously performs panoptic segmentation, instance orientation estimation, and scene classification. For further details and explanations, see Sec. III. Semantic colors are chosen as in [2] and are the default colors for NYUv2 [3]. Panoptic is visualized by small color differences.

semantically rich features from the input and performs downsampling up to a factor of 1/32 of the input resolution to reduce computational effort. However, instead of using a fused dual encoder, our EMSAFormer features only a single encoder with a modified SwinV2 Transformer backbone to incorporate RGB and depth data. We address these modifications in Sec III-A. After the encoder, a context module (CM) similar to the pyramid pooling module in PSPNet [38] is attached. Although Transformers already enable a larger receptive field [15], we observe a performance boost due to an additional context module in our experiments (see results later in Fig. 4). Due to the large downsampling at the end of the encoder, the computational effort of the context module is almost negligible. While all tasks share the same encoder (often referred to as hard-parameter sharing [53]), we use three independent decoders, not sharing any network parameters, to handle the tasks for scene analysis. We introduce the task-specific decoders in Sec III-B. Similar to EMSANet, the entire network architecture is tailored to enable fast inference. However, as Transformer-based architectures are relatively new, inference optimization is crucial and currently rare. We address this key aspect with an additional NVIDIA TensorRT extension introduced in Sec III-C.

#### A. Encoder

The encoder of our EMSAFormer is derived from the SwinV2 Transformer [17] architecture. We build on top of the tiny model SwinV2-T as it is the only version that currently enables real-time inference on our target hardware. The next larger SwinV2-S and SwinV2-B increase inference time by a factor of 1.5 and 1.9, respectively, and, thus, are out of our scope. The architecture of SwinV2-T is shown Fig 3 (a). Each RGB input is processed in four stages. The first stage embeds the input using a  $4 \times 4$  convolution with stride 4 to 96 feature maps. As this convolution processes non-overlapping patches of  $4 \times 4$  pixels, this step is called patch embedding. Subsequent to the patch embedding, the first two SwinV2 blocks are attached within the same stage. Each SwinV2 block comprises

a multi-head self-attention module (MSA) and a subsequent 2-layer multilayer perceptron (MLP). Both the MSA and the MLP are followed by a layer normalization and further feature a skip connection as show in Fig. 3. The design of the SwinV2 block follows the original Transformer block introduced in [9]. However, in contrast to [9], [14], [15], the attention is computed using a scaled cosine function instead of the softmax function. Moreover, as computing the self-attention between all elements requires quadratic complexity, the authors adapt the MSA to a window multi-head self-attention (W-MSA) that divides the input in non-overlapping windows of size  $8 \times 8$ in order to reduce complexity. However, this approach lacks connections across windows and, thus, prevents the model from building context features. To overcome this limitation, the authors further introduce a shifted-window multi-head self-attention (SW-MSA). By alternating both modules, the network is able to exchange information between adjacent windows. For further details on the exact architecture, we refer to [14], [17]. The remaining stages follow the same design. However, the initial patch embedding is replaced with a patch merging operation, and the number of repeated SwinV2 blocks differs in stage 3. The patch merging aims at reducing complexity while creating hierarchical features as the network gets deeper similar to CNN-based architectures [6]-[8]. To achieve this, the spatial resolution gets downsampled by a factor of 2, while the number of feature maps is doubled.

To incorporate depth data, we examine the modification depicted in Fig 3 (bottom). The most straight-forward way is shown in Fig 3 (b) and integrates depth as additional modality to the patch embedding. The missing weights can be derived either by reusing the existing weights (D=R+G+B) or by performing an additional pretraining step. We refer to both modifications as *SwinV2-T* and *SwinV2-T-Pre*, respectively. Unfortunately, the ImageNet dataset [13] used for pretraining does not feature depth data. Therefore, we use a grayscale image instead for pretraining (D=gray). This way, the backbone already learns to handle four input channels during pretraining.



Fig. 3. Original SwinV2-T architecture [17] (top) and our modifications (bottom) to efficiently incorporate depth information in a single encoder backbone.

However, a major drawback of this approach is that both modalities are mixed right at the beginning of the network in the patch-embedding step. As already shown in [18], such an early mixing of both modalities introduces issues due to deviating statistics and characteristics and eliminates any benefits. Luckily, the Swin Transformer architecture enables to prevent mixing features up to the first patch merging. Similar to most other Transformer-based architectures [9], [11], [12], the attention is not computed across all channels but instead only on a subset of the channels. In Swin Transformers, each attention head only processes a subset of 32 channels of the input to a SwinV2 block (see MSA box in Fig 3). To take advantage of this property, only the patch embedding needs to be adapted. As shown in Fig 3 (c), we propose to split the  $4 \times 4$ convolution into two convolutions, the first embedding RGB to 64 channels and the second embedding depth to the remaining 32 channels. We refer to this modification as SwinV2-T-Multi. Note that the MLPs subsequent to the multi-head self-attention blocks (see Fig 3 red) still combine channels. However, there is a skip connection, which means that the network is able to learn whether to combine features or not in an adjustable way.

While this approach focuses on independent processing of depth, both modalities are embedded to the 96 channels of original SwinV2-T model. Embedding more information with the same number of channels, may lead to a bottleneck. To overcome this issue, we further propose to enlarge the width of the entire model and to use 128, i.e., 96+32, channels in the initial embedding. This way, the RGB embedding retains its original representation capabilities, and depth is embedded to additional 32 channels. Note, due to the subset property of the attention heads, depth is still processed in independent attention heads. We refer to this modification as SwinV2-T-128-Multi in Fig 3 (d). Note that this modification leads to a width similar to the larger SwinV2-B. However, the number of blocks and, thus, the depth remains the one of SwinV2-T. To further strengthen the independent processing of depth in the subset of channels, we further add an additional augmentation step to the pretraining pipeline that randomly masks out either the whole RGB or grayscale image. This way, the network is forced to learn to extract information from both images and cannot rely on the more meaningful RGB input solely. We refer to this modification as SwinV2-T-128-Multi-Aug.

In Sec IV-D, we examine the suitability of the aforementioned modifications. We further investigated modification to the patch merging to extend our design principle, i.e., processing depth in an independent subset of channels with an adjustable fusion mechanism in the MLPs, to the entire architecture. However, we could not observe any significant performance improvement when adapting subsequent stages.

# B. Decoders

The decoders for our multi-task architecture (see Fig 2) are designed to suit the specific needs of each task. To obtain the final prediction for scene classification, a single fully-connected layer is attached to the global-average-pooling branch of the context module. For panoptic segmentation, two

dense decoders are used to perform semantic and instance segmentation. Each decoder is followed by a task head that projects to the required number of channels for the corresponding tasks and, finally, restores the input resolution. For semantic segmentation, the task head projects to the number of semantic classes. For instance segmentation, we follow the bottom-up approach of Panoptic DeepLab [45] and EMSANet [2]. As shown in Fig 2, instances are represented by their center of mass encoded within a heatmap predicted by the first instance head. To assign pixels to instance centers, a second instance head further predicts offset vectors pointing towards a corresponding instance center. To ignore pixels belonging to stuff classes, e.g., wall or floor, a foreground mask derived from the semantic head is applied before assigning any pixel. As shown in [2], instance orientation estimation can also be handled in a dense manner with an additional head on top of the instance decoder. For each pixel the orientation is predicted as continuous angle around the axis perpendicular to the ground plane. To obtain the orientation of an instance, all predictions assigned to this instance are averaged.

For the dense decoders, we consider both the CNN-based decoder from EMSANet (see red block in instance branch in Fig. 2) as well as the SegFormer MLP decoder (see green block in semantic branch in Fig. 2). The original SegFormer MLP decoder in [15] uses four branches with equal number of channels, i.e., [128, 128, 128, 128] channels. We propose to follow the pyramid-like structure with increasing resolution of the EMSANet decoder and use [256, 128, 64, 64] channels. This way, both inference throughput and performance are increased. We refer to this decoder as modified SegFormer decoder. In our experiments, we examine and compare the performance of both decoder types, the EMSANet decoder and the modified SegFormer decoder.

# C. Optimization

For efficient and fast inference on embedded hardware, optimized inference frameworks, such as NVIDIA TensorRT, are crucial. Compared to CNN-based architectures, Transformerbased architectures are relatively new and, thus, currently lack the same kind of highly optimized inference engines. However, there is already ongoing effort to optimize inference throughput. FasterTransformer [54] provides a first extension to NVIDIA TensorRT with focus on Transformers. However, while already archiving a significant speedup, they currently only focus on optimizing architectures for image classification with inputs of fixed and square resolution. Moreover, the whole encoder is optimized as a single block, which makes it impossible to access intermediate features for skip connections or to incorporate any modification to the encoder architecture, such as a modified patch embedding or merging, a deviating number of channels, or another kind of normalization layer. To overcome these limitations, we propose a custom NVIDIA TensorRT extension. It is based on the existing FasterTransformer implementation but splits the encoder in smaller blocks to enable more flexibility in downstream tasks. We further added support for bottom-right padding to the CUDA kernels

if the input size is not a multiple of the window size used in the shifted-window attention. This is of great importance to enable inference with inputs of arbitrary input resolution. The modular design is complemented by the ability to import models from Open Neural Network Exchange (ONNX) format. The proposed extension enables us to examine the architecture modifications described in Sec III-A.

#### **IV. EXPERIMENTS**

We evaluate the performance of our proposed multi-task approach on the common indoor RGB-D datasets NYUv2 [3], SUNRGB-D [4], and ScanNet [5]. We start with a singletask setting on the smaller NYUv2 dataset to assess the performance of the SwinV2-based encoder and our proposed modifications across the individual tasks. The goal is to derive a suitable encoder that is capable of handling all tasks in a most efficient way. Moreover, we investigate the performance of our proposed encoder with both dense decoder types, the EMSANet decoder and the modified SegFormer decoder. With the results of these experiments at hand, we combine all tasks in a multi-task approach. The goal is to solve all four tasks, i.e., semantic segmentation, instance segmentation, instance orientation estimation, and scene classification, simultaneously using a single neural network. Finally, we demonstrate the suitability of our approach for the larger SUNRGB-D and ScanNet dataset and compare to other state-of-the-art approaches.

#### A. Implementation Details

Our implementation is built using PyTorch [55] and is based on the EMSANet implementation [2]. As commonly done in downstream tasks, we use weights pretrained on ImageNet [13] to initialize the encoders. However, as already stated in Sec. III-A, any modification except for SwinV2-T (D=R+G+B) requires an additional pretraining step. Pretraining Transformers from scratch is very time consuming and requires large batch sizes and, thus, heavy memory requirements. We used 8× NVIDIA A100 40 GB GPUs to accomplish these pretrainings. To enable other applications, we share the pretrained weights to the research community on GitHub. Note that pretraining larger models, such as SwinV2-S or SwinV2-B, implies even higher memory requirements and, thus, is out of our scope. Subsequent training of the EMSAFormer architecture requires less resources and can be done on any GPU with at least 25 GB of VRAM (all tasks simultaneously). We stick to the training pipeline of EMSANet [2] for data processing and augmentation. We use a fixed input resolution of  $640 \times 480$  pixels and a batch size of 8. Each network is trained for 500 epochs using SGD with a momentum of 0.9 and a small weight decay of 0.0001. The learning rate is varied in {0.00125, 0.0025, 0.005, 0.01, 0.02, 0.03, 0.04, 0.06} depending on the actual dataset and tasks. We further use a one-cycle learning rate scheduler, similar to the for Transformer commonly used cosine-annealing learning rate scheduler, to adjust the learning during training. For further details and other hyperparameters, we refer to our implementation available on GitHub.

# B. Datasets

Selecting datasets suitable for evaluating our multi-task approach is challenging as it requires RGB and depth data as well as annotations for the individual tasks. In the following, we give a brief overview over the RGB-D datasets used.

*NYUv2 [3]:* The NYUv2 dataset comprises 795 training samples and 654 test samples. It provides annotations for semantic and instance segmentation and scene classification. We use the semantic annotations with 40 classes. As the original dataset does not include annotations for instance orientation, we use the manually annotated ones from [2].

*SUNRGB-D* [4]: The SUNRGB-D dataset features 5,285 training and 5,050 test samples from multiple RGB-D cameras. The dataset provides annotations for the first 37 NYUv2 semantic classes and for scene classification. However, it lacks dense annotations for instance segmentation. We stick to the reconstructed instance annotations from 3D bounding boxes proposed in [2], which also provide orientation annotations.

ScanNet [5]: The ScanNet dataset comprises 1.89M training, 0.53M validation, and 0.21M test images. It provides annotations for semantic and instance segmentation as well as for scene classification. We use the semantic class mapping to the 40 NYUv2 classes. As the dataset is created from video sequences and, thus, contains many similar images, we follow the official recommendation [5] and use the subsample of 100 for the validation and test split. To reduce training time, we use a subsample of 50 and limit the number of samples to a random subset of 25% for each epoch. As the dataset lacks instance orientation annotations, we cannot train this task on ScanNet. However, given the size and the quality of the annotations, it is still an important dataset.

For creating panoptic annotations, we combine the dense annotations of the datasets and treat floor, wall, and ceiling as stuff classes. For scene classification, we use the unified class spectrum for the most relevant indoor classes presented in [2]. As the ScanNet dataset was not in the scope of [2], we created a similar scene class mapping for this dataset.

# C. Metrics

We follow the evaluation protocol of [2] and report the mean intersection over union (mIoU) for semantic segmentation, panoptic quality (PQ), segmentation quality (SQ) and recognition quality (RQ) for panoptic segmentation as well as the mean absolute angular error (MAAE) for instance orientation estimation. To enable experiments in a single-task setting, PQ is also reported for instance segmentation. Note that PQ tracks closely to the average precision (AP) [41] and, thus, also evaluates instance segmentation in a meaningful way. However, the instance decoder performs class-agnostic instance segmentation. Therefore, we use the ground-truth semantic segmentation for creating a foreground mask and for assigning semantic classes. The reported PQ is equal to perfect semantic segmentation. Similarly, for single-task instance orientation estimation, the MAAE is computed assuming perfect instances. For scene classification, the balanced accuracy (bAcc) is used to account for the imbalanced class distribution. As we aim at fast inference, we do not apply any evaluation tricks, such as test time augmentation. Furthermore, to enable fair comparison, we always upsample dense predictions to the full input resolution before determining any metric.

#### D. Single-task Performance

We start by evaluating the proposed SwinV2-Transformerbased encoder and its modifications in a single-task setting.

Semantic Segmentation (Sem): Fig 4 (a) visualizes the results for semantic segmentation and compares to the dualencoder approaches of EMSANet. It becomes obvious that SwinV2-T can also be used as backbone in a dual-encoder design leading to an mIoU similar to one with ResNet101 backbones except for processing depth solely (blue in Fig 4). This highlights that SwinV2 is tailored to processing RGB inputs. However, the dual-encoder design results in a significant drop in inference throughput, making such a design not suitable for our application scenario. Changing the decoder to the smaller modified SegFormer decoder leads to similar performance but cannot alleviate the gap in inference throughput. By contrast, relying on a single encoder greatly improves inference throughput. However, the results (red in Fig 4) also highlight the challenge of processing both modalities in a single encoder. The performance of SwinV2-T drops to an mIoU of ~48%. Additional pretraining on RGB-Gray inputs (SwinV2-T-Pre) can only halve the gap to the dual-encoder counterpart (green in Fig 4). Further splitting the patch embedding to emphasize processing both modalities independently, as done in SwinV2T-Multi, does not close the remaining gap in performance. This shows that the model is not capable of handling both modalities in the original embedding with 96 channels. The wider SwinV2-T-128, which uses 128 instead of 96 channels in the patch embedding, benefits much more from splitting the patch embedding (SwinV2-T-128-Multi). Adapting data augmentation during pretraining (SwinV2-T-128-Multi-Aug) to further strengthen independent processing of depth later in the downstream tasks results in another improvement. Our single encoder with SwinV2-T-128-Multi-Aug backbone leads to slightly better performance than the dual-encoder design with ResNet101 backbone at almost the same inference speed.

Instance Segmentation (Ins): For instance segmentation, a similar trend is emerged. However, the results in Fig 4 (b) show two new aspects. First, there is a gap of ~6% in PQ between CNN-based and Transformer-based encoders independently of the modality or the encoder design. As discussed later in Sec. IV-E, this gap highlights optimization issues in the Transformer-based encoder due to a small number of training samples along with a more challenging task. Second, the modified SegFormer MLP decoder leads consistently to even worse results. Therefore, we stick to the EMSANet decoder for instance segmentation for the remaining experiments.

*Orientation Estimation (Or):* The results in Tab. I (top) show that the gap experienced for instance segmentation is not present for instance orientation estimation. This could be due to the fact that this task is easier to accomplish in general and of more similar complexity to semantic segmentation.



Fig. 4. Results on NYUv2 test split when performing semantic segmentation (top) and instance segmentation (bottom) in a single-task setting with various encoder configurations over inference throughput (NIVIDA Jetson AGX Orin 32 GB, Jetpack 5.1.1, TensorRT 8.5.2, Float16, 50 W). See Sec. IV-C for metrics.

#### TABLE I

Results on NYUv2 test split when performing instance orientation estimation (top) and scene classification (bottom) in a single-task setting with various encoder configurations. See Sec. IV-C for metric abbreviations.

Ori	ientation Estimation			RGB-	D
MA	AE↓	RGB	Depth	Fused Dual	Single
	ResNet34-NBt1d	22.24	18.36	17.91	_
5	ResNet50	23.09	18.81	18.41	_
	ResNet101	22.06	18.02	17.50	—
	SwinV2-T	24.08	19.13	19.02	19.52
50	SwinV2-T-Pre	_	_	_	18.91
In	SwinV2-T-128	_	_	_	18.89
0	SwinV2-T-128-Multi-Aug	_	—	—	17.85
<b>.</b>	ana Classification			DCD	<u> </u>
Sc	ene Classification			RGB-	D
Sco bA	ene Classification cc <sup>↑</sup>	RGB	Depth	RGB- Fused Dual	D Single
Sco bA	ene Classification cc <sup>↑</sup> ResNet34-NBt1d	<b>RGB</b> 74.40	<b>Depth</b> 67.26	RGB- Fused Dual 72.40	D Single —
Sco bA	ene Classification cc <sup>↑</sup> ResNet34-NBt1d ResNet50	<b>RGB</b> 74.40 74.19	<b>Depth</b> 67.26 69.92	<b>RGB</b> - <b>Fused Dual</b> 72.40 74.91	D Single 
Sco bA	ene Classification cc <sup>↑</sup> ResNet34-NBt1d ResNet50 ResNet101	<b>RGB</b> 74.40 74.19 74.95	<b>Depth</b> 67.26 69.92 70.53	RGB- Fused Dual 72.40 74.91 <b>75.86</b>	D Single — — —
Sce bA	ene Classification cc <sup>↑</sup> ResNet34-NBt1d ResNet50 ResNet101 SwinV2-T	<b>RGB</b> 74.40 74.19 74.95 76.84	<b>Depth</b> 67.26 69.92 70.53 66.72	RGB- Fused Dual 72.40 74.91 75.86 73.39	D Single — — — 76.52
Sci bA [7]	ene Classification cc <sup>↑</sup> ResNet34-NBt1d ResNet50 ResNet101 SwinV2-T SwinV2-T-Pre	<b>RGB</b> 74.40 74.95 76.84 —	<b>Depth</b> 67.26 69.92 70.53 66.72	RGB- Fused Dual 72.40 74.91 75.86 73.39	D Single — — 76.52 77.32
Sco bA	ene Classification cc <sup>↑</sup> ResNet34-NBt1d ResNet50 ResNet101 SwinV2-T SwinV2-T-Pre SwinV2-T-128	<b>RGB</b> 74.40 74.19 74.95 76.84	<b>Depth</b> 67.26 69.92 70.53 66.72	RGB- Fused Dual 72.40 74.91 75.86 73.39	D Single — — 76.52 77.32 78.21

Our SwinV2-T-128-Multi-Aug encoder achieves comparable performance to a dual encoder with ResNet101 backbone, while outperforming all other dual-encoder approaches.

Scene Classification (Sce): The results in Tab. I (bottom) again highlight the strength of Transformer-based architectures for image classification. The performance of the Transformer-based single encoder model processing RGB solely already outperforms all CNN-based approaches with ResNet backbone. Furthermore, each single RGB-D encoder model outperforms all dual-encoder approaches. Our proposed SwinV2-T-128-Multi-Aug achieves the best result.

The results of this set of experiments show that our proposed SwinV2-T-128-Multi-Aug backbone is capable of handling all tasks. Issues for instance segmentation are addressed below.

#### E. Multi-task Performance

Learning multiple tasks using a single neural network is challenging as the tasks may influence each other. Thus, balancing the losses to each other and selecting the best epoch are crucial. We put less focus on orientation estimation as the results are already close to annotation quality [2]. However, as we focus on the Transformer-based encoder in this publication, we refer to our implementation for the actual task balancing. The best epoch is chosen based on the task most relevant for our application, i.e., the PQ for panoptic segmentation. However, to get a better impression on the performance of the individual tasks when trained simultaneously in a multi-task setting and independent of selecting a specific checkpoint, we also report the best result for each metric within the same run.

Tab II shows the multi-task results for NYUv2. It becomes obvious that all tasks can be solved using a single neural network. The results for the individual tasks are close to single-task performance. The PQ for instance segmentation increases noticeably and closes the gap to CNN-based dual encoders partially. This indicates that the encoder features learned in the multi-task setting are more beneficial for instance segmentation. Surprisingly, exchanging the modified SegFormer decoder (denoted by Sem(SegFormer) in Tab. II) with the EMSANet decoder (denoted by Sem in Tab. II) for semantic segmentation improves almost the entire multi-task performance. This suggests that – at least for the smaller NYUv2 dataset – two identical dense decoders lead to better encoder features. Except for instance segmentation, the multitask results are close to EMSANet with ResNet101 backbone.

The results for SUNRGB-D and ScanNet in Tab. III show a different picture. Performing semantic segmentation with the modified SegFormer decoder (denoted by Sem(SegFormer) in Tab. III) consistently leads to better multi-task results. The gap to the EMSANet decoder further gets larger as the dataset size increases, i.e., for ScanNet. We deduce that the NYUv2 dataset is too small to fit the parameters of our model with modified SegFormer decoder. The proposed EMSAFormer configuration (see Fig 2) outperforms both EMSANet approaches for semantic and panoptic segmentation and scene classification.

Tab II further reports the inference throughput for all approaches and task settings. Note that the values also apply for SUNRGD-D and ScanNet as they feature less or the same number of semantic classes. Even with a lower power profile (measured power consumption of 30 W), our

TABLE II

Results on NYUv2 test split when training our multi-task EMSAFormer and in comparison to EMSANet [2]. See Sec. IV-C for the reported metrics. Panoptic results are obtained after merging semantic and instance prediction. Legend: italic: metric used for determining the best checkpoint, gray: best result within the same run, FPS<sub>x</sub>: frames per second on an NVIDIA Jetson AGX Orin 32 GB (Jetpack 5.1.1, TensorRT 8.5.2, Float16) at measured power consumption x.

	Backbone	Task(s)	Semantic Decoder mIoU↑	Inst Dec PQ↑	ance oder MAAE↓	Scene Head bAcc ↑	mIoU↑	Pano (afte PQ↑	optic Ro er merş RQ↑	esults ging) SQ↑	MAEE ↓	Infer Throu FPS <sup>↑</sup> <sub>50W</sub>	ence ghput FPS ↑ 30W
EMSAFormer (ours)	SwinV2-T-128-Multi-Aug	Semantic Segmentation (Sem) Instance Segmentation (Ins) Orientation Estimation (Or) Scene Classification (Sce)	50.53 — — —	56.44 	 17.85 	  78.66	 	 	 		 	47.1 48.4 47.7 58.9	30.5 33.5 32.9 40.7
		Sem(SegFormer) + Sce + Ins + Or	<b>50.23</b> 50.51	<b>58.75</b> 59.25	<b>20.95</b> 20.95	<b>77.70</b> 80.02	<b>51.34</b> 51.34	<b>43.41</b> 43.41	<b>52.53</b> 52.53	<b>81.75</b> 81.79	<b>18.94</b> 18.94	39.1	27.3
		Sem + Sce + Ins + Or	<b>51.06</b> 51.26	<b>59.06</b> 59.27	<b>20.01</b> 18.09	<b>78.80</b> 78.80	<b>51.76</b> 51.76	<b>43.28</b> 43.28	<b>52.48</b> 52.48	<b>81.43</b> 81.51	18.26 18.09	36.5	25.6
EMSANet	2x ResNet101	Sem + Sce + Ins + Or	<b>50.83</b> 51.01	<b>62.64</b> 62.81	<b>17.87</b> 17.82	<b>77.41</b> 78.43	<b>50.67</b> 51.23	<b>45.12</b> 45.12	<b>54.02</b> 54.02	<b>82.49</b> 82.99	<b>15.33</b> 14.73	42.9	30.1
	2x ResNet34-NBt1D	Sem + Sce + Ins + Or	<b>50.97</b> 51.15	<b>61.33</b> 61.53	<b>18.37</b> 18.37	<b>76.46</b> 78.18	<b>50.61</b> 51.31	<b>43.59</b> 43.59	<b>52.23</b> 52.27	<b>82.48</b> 82.70	<b>16.39</b> 15.76	70.5	49.9

#### TABLE III

Results on SUNRGB-D test split and ScanNet validation split when training our multi-task EMSAFormer and in comparison to EMSANet [2]. See Sec. IV-C for details on the reported metrics. Legend: italic: metric used for determining the best checkpoint, gray: best result within the same run.

		Model	Semantic Decoder mIoU↑	Inst Dec PO↑	ance oder MAAE↓	Scene Head bAcc↑	mIoU ↑	Pano (afte PO↑	ptic Re r merg RO↑	sults ing) SO↑	MAEE ↓
SUNRGB-D	EMSAFormer	SwinV2-T-128-Multi-Aug	<b>48.52</b> 48.67	<b>61.14</b> 61.60	<b>16.99</b> 16.99	<b>62.01</b> 64.96	<b>45.12</b> 45.27	<b>50.08</b> 50.08	<b>59.08</b> 59.08	<b>84.68</b> 84.83	<b>15.32</b> 14.93
		SwinV2-T-128-Multi-Aug (Sem(SegFormer))	<b>48.61</b> 48.82	<b>61.20</b> 61.78	<b>15.91</b> 15.91	<b>61.97</b> 64.50	<b>45.79</b> 45.94	<b>51.70</b> 51.70	<b>60.12</b> 60.15	<b>84.65</b> 84.65	<b>14.00</b> 13.90
	EMSANet	2x ResNet101	<b>47.99</b> 47.99	<b>62.07</b> 62.96	<b>15.17</b> 15.17	<b>59.40</b> 61.21	<b>43.22</b> 44.19	<b>51.06</b> 51.75	<b>58.88</b> 59.74	<b>85.53</b> 85.64	<b>13.34</b> 13.00
		2x ResNet34-NBt1D	<b>48.39</b> 48.39	<b>60.62</b> 61.48	<b>16.28</b> 14.83	<b>61.76</b> 62.66	<b>45.53</b> 45.66	<b>49.88</b> 50.53	<b>57.79</b> 58.66	<b>84.91</b> 85.20	<b>14.23</b> 14.15
ScanNet	EMSAFormer	SwinV2-T-128-Multi-Aug	<b>63.78</b> 63.78	<b>66.69</b> 66.71		<b>48.82</b> 49.70	<b>61.93</b> 61.93	<b>49.70</b> 49.70	<b>59.15</b> 59.15	<b>83.31</b> 83.36	_
		SwinV2-T-128-Multi-Aug (Sem(SegFormer))	<b>64.75</b> 64.75	<b>67.71</b> 67.84	_	<b>49.69</b> 49.73	<b>62.66</b> 62.66	<b>51.18</b> 51.18	<b>61.01</b> 61.01	<b>83.20</b> 83.38	_
	EMSANet	2x ResNet101	<b>63.63</b> 64.11	<b>66.36</b> 66.64	_	<b>44.63</b> 46.32	<b>61.92</b> 61.96	<b>50.35</b> 50.35	<b>59.82</b> 59.82	<b>83.51</b> 83.80	_
		2x ResNet34-NBt1D	<b>61.25</b> 61.25	<b>65.57</b> 65.57	_	<b>45.47</b> 46.35	<b>58.32</b> 58.32	<b>47.76</b> 47.76	<b>56.85</b> 56.85	<b>83.39</b> 83.47	_

proposed multi-task EMSAFormer approach meets our realtime requirements of at least 20 FPS.

Fig. 5 presents qualitative results. For all indoor datasets, our approach is able to analyze the scenes thoroughly. The obtained predictions are well suited for enabling a mobile robot to operate autonomously in indoor environments.

#### F. Comparison to State of the Art

Comparing our proposed EMSAFormer to other approaches is challenging, as they mainly focus on single-task semantic segmentation and rarely account for efficiency. Moreover, most approaches use test-time augmentation, which is not applicable on a mobile robot with limited computational resources. The only approach that fits our multi-task setting is EMSANet [2], which we already compared to above. However, Tab. IV shows additional comparisons to RGB-D approaches on all three common indoor datasets for semantic segmentation. The results on NYUv2 and SUNRGB-D reveal that our approach outperforms other CNN-based approaches. For NYUv2, the results are close to the OMNIVORE approach with the larger Swin-B backbone. We also report results for the official Scan-Net benchmark (hidden test split). Our approach outperforms EMSANet with a dual ResNet101 encoder on this split as well.



Fig. 5. Qualitative results as RGB image overlayed with predicted panoptic segmentation, predicted scene class, and estimated orientations if available.

#### V. CONCLUSION AND FUTURE WORK

We have presented a Transformer-based RGB-D approach for multi-task scene analysis, called EMSAFormer, that simultaneously performs panoptic segmentation, instance orientation estimation, and scene classification. We have shown that the CNN-based dual encoder of EMSANet [2] can be replaced with a single Trasformer-based encoder. Our extensive experiments on the three common indoor datasets NYUv2, SUNRGB-D, and ScanNet highlight the strong performance of our proposed EMSAFormer. We further have revealed limitations of Transformer-based approaches in a single-task setting on smaller datasets, such as NYUv2. However, we have demonstrated that these issues can be addressed using a multi-task approach. Due to the proposed NVIDIA TensorRT extension, our EMSAFormer approaches can be applied in real time with 39.1 FPS on an NVIDIA Jetson AGX Orin 32 GB, demonstrating its suitability for deployment to mobile robots. In future work, we intend to explore the benefits of additional training on large-scale RGB-D datasets such as Hypersim [56].

TABLE IV

Comparison to other state-of-the-art approaches on NYUv2 test split, SUNRGB-D test split, and ScanNet test (benchmark) split. \* indicates additional test-time augmentation.

		e	
		Backbone	mIoU $^{\uparrow}$
NYUv2	OMNIVORE [37]	Swin-T	47.9
		Swin-B	51.1
	ShapeConv [36]	ResNet50	47.3
		ResNext101	50.2
	SA-Gate [25]	2xResNet50	50.4
	EMSAFormer (ours)	SwinV2-T-128-Multi-Aug	51.26
0	ShapeConv [36]	ResNet101	47.6
3	AC-Net [20]	3x ResNet50	48.1
G	2.5D Conv [33]	ResNet-101	48.2
SUNR	ESANet [26]	2x ResNet50	48.31
	EMSAFormer (ours)	SwinV2-T-128-Multi-Aug (Sem(SegFormer))	48.82
ScanNet	FuseNet (from [23])	2x VGG16	53.5
	SSMA [23]	2x mod. ResNet50	57.7*
	EMSANet	2x ResNet101	54.0
	EMSAFormer (ours)	SwinV2-T-128-Multi-Aug (Sem(SegFormer))	56.4

#### REFERENCES

- [1] T. Wengefeld, B. Schuetz, G. Girdziunaite, A. Scheidig, and H.-M. Gross, "The MORPHIA Project: First Results of a Long-Term User Study in an Elderly Care Scenario from Robotic Point of View," in Proc. of ISR, 2022.
- [2] D. Seichter, S. Fischedick, M. Köhler, and H.-M. Gross, "Efficient Multi-Task RGB-D Scene Analysis for Indoor Environments," in Proc. of IJCNN, 2022, pp. 1-10.
- N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor Segmentation [3] and Support Inference from RGBD Images," in Proc. of ECCV, 2012.
- S. Song, S. P. Lichtenberg, and J. Xiao, "SUN RGB-D: A RGB-D Scene [4] Understanding Benchmark Suite," in Proc. of CVPR, 2015, pp. 567-576.
- A. Dai, A. X. Chang, M. Savva et al., "ScanNet: Richly-annotated 3D [5] Reconstructions of Indoor Scenes," in Proc. of CVPR, 2017.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. of CVPR, 2016, pp. 770-778.
- M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in Proc. of ICML, 2019, pp. 6105-6114.
- Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A ConvNet for the 2020s," in Proc. of CVPR, 2022, pp. 770-778.
- [9] A. Vaswani, N. Shazeer, N. Parmar et al., "Attention is all you need," Proc. of NeurIPS, vol. 30, 2017.
- [10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," arXiv preprint arXiv:1810.04805, 2018.
- [11] T. Brown, B. Mann, N. Ryder et al., "Language models are few-shot learners," Proc. of NeurIPS, vol. 33, 2020.
- [12] A. Dosovitskiy, L. Beyer, A. Kolesnikov et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," in Proc. of ICLR, 2021.
- [13] O. Russakovsky, W. Dong, R. Socher et al., "ImageNet Large Scale Visual Recognition Challenge," in *IJCV*, 2015, pp. 211–252. [14] Z. Liu, Y. Lin, Y. Cao *et al.*, "Swin Transformer: Hierarchical Vision
- Transformer using Shifted Windows," in Proc. of ICCV, 2021.
- [15] E. Xie, W. Wang et al., "SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers," in Proc. of NeurIPS, 2021.
- W. Wang, E. Xie, X. Li et al., "Pyramid Vision Transformer: A Versatile [16] Backbone for Dense Prediction Without Convolutions," in Proc. of ICCV, 2021, pp. 568-578.
- Z. Liu, H. Hu, Y. Lin et al., "Swin Transformer V2: Scaling Up Capacity [17] and Resolution," in Proc. of CVPR, 2022.
- [18] C. Hazirbas, L. Ma, C. Domokos, and D. Cremers, "FuseNet: Incorporating Depth into Semantic Segmentation via Fusion-based CNN Architecture," in Proc. of ACCV, 2016, pp. 213-228.
- [19] J. Jiang, L. Zheng, F. Luo, and Z. Zhang, "RedNet: Residual Encoder-Decoder Network for indoor RGB-D Semantic Segmentation," arXiv preprint arXiv:1806.01054, 2018.
- [20] X. Hu, K. Yang, L. Fei, and K. Wang, "ACNet: Attention Based Network to Exploit Complementary Features for RGBD Semantic Segmentation," Proc. of ICIP, 2019.
- [21] S. Lee, S. J. Park, and K. S. Hong, "RDFNet: RGB-D Multi-Level Residual Feature Fusion for Indoor Semantic Segmentation," Proc. of ICCV, pp. 4990-4999, 2017.
- [22] Y. Xing, J. Wang, X. Chen, and G. Zeng, "Coupling Two-Stream RGB-D Semantic Segmentation Network by Idempotent Mappings," in Proc. of ICIP, 2019, pp. 1850-1854.
- [23] A. Valada, R. Mohan, and W. Burgard, "Self-supervised Model Adaptation For Multimodal Semantic Segmentation," IJCV, 2019.
- [24] F. Fooladgar and S. Kasaei, "Multi-Modal Attention-based Fusion Model for Semantic Segmentation of RGB-Depth Images," arXiv preprint arXiv:1912.11691, pp. 1-12, 2019.
- [25] X. Chen, K.-Y. Lin, J. Wang et al., "Bi-directional Cross-Modality Feature Propagation with Separation-and-Aggregation Gate for RGB-D Semantic Segmentation," in Proc. of ECCV, 2020, pp. 561-577.
- [26] D. Seichter, M. Köhler, B. Lewandowski, T. Wengefeld, and H.-M. Gross, "Efficient RGB-D Semantic Segmentation for Indoor Scene Analysis," in Proc. of ICRA, 2021, pp. 13525-13531.
- [27] S.-W. Hung, S.-Y. Lo, and H.-M. Hang, "Incorporating Luminance, Depth and Color Information by a Fusion-Based Network for Semantic Segmentation," in Proc. of ICIP, 2019, pp. 2374-2378.
- [28] Z. Li, Y. Gan, X. Liang, Y. Yu, H. Cheng, and L. Lin, "LSTM-CF: Unifying Context Modeling and Fusion with LSTMs for RGB-D Scene Labeling," in Proc. of ECCV. Springer, 2016, pp. 541-557.

- [29] Q. Ha, K. Watanabe, T. Karasawa, Y. Ushiku, and T. Harada, "MFNet: Towards Real-Time Semantic Segmentation for Autonomous Vehicles with Multi-Spectral Scenes," in Proc. of IROS, 2017, pp. 5108-5115.
- Y. Li, J. Zhang, Y. Cheng, K. Huang, and T. Tan, "Semantics-guided [30] Multi-Level RGB-D Feature Fusion for Indoor Semantic Segmentation," in Proc. of ICIP, 2017, pp. 1262-1266.
- [31] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation," in Proc. of ECCV, 2018, pp. 801-818.
- W. Wang and U. Neumann, "Depth-Aware CNN for RGB-D Segmen-[32] tation," in Proc. of ECCV, 2018, pp. 144-161.
- Y. Xing, J. Wang, X. Chen, and G. Zeng, "2.5D Convolution for RGB-D [33] Semantic Segmentation," in Proc. of ICIP, 2019, pp. 1410-1414.
- [34] Y. Xing, J. Wang, and G. Zeng, "Malleable 2.5D Convolution: Learning Receptive Fields along the Depth-axis for RGB-D Scene Parsing," in Proc. of ECCV, 2020, pp. 1-17.
- L.-Z. Chen, Z. Lin, Z. Wang, Y.-L. Yang, and M.-M. Cheng, "Spatial [35] Information Guided Convolution for Real-Time RGBD Semantic Segmentation," arXiv preprint arXiv:2004.04534v1, pp. 1-11, 2020.
- [36] M. Cao, H. Leng, D. Lischinski, D. Cohen-Or, C. Tu, and Y. Li, "ShapeConv: Shape-aware Convolutional Layer for Indoor RGB-D Semantic Segmentation," in Proc. of ICCV, 2021.
- [37] R. Girdhar, M. Singh, N. Ravi et al., "Omnivore: A Single Model for Many Visual Modalities," in Proc. of CVPR, 2022.
- [38] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid Scene Parsing Network," in Proc. of CVPR, 2017, pp. 2881-2890.
- [39] L.-C. Chen, G. Papandreou, F. Schroff et al., "Rethinking Atrous Convolution for Semantic Image Segmentation," arXiv preprint arXiv:1706.05587, 2017.
- M. Oršić, I. Krešo, P. Bevandić, and S. Šegvić, "In Defense of Pre-[40] trained ImageNet Architectures for Real-time Semantic Segmentation of Road-driving Images," in Proc. of CVPR, 2019, pp. 12607-12616.
- [41] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár, "Panoptic Segmentation," in Proc. of CVPR, 2019, pp. 9404-9413.
- [42] A. Kirillov, R. Girshick, K. He, and P. Dollár, "Panoptic Feature Pyramid Networks," in Proc. of CVPR, 2019, pp. 6399-6408.
- Y. Xiong, R. Liao, H. Zhao et al., "UPSNet: A Unified Panoptic [43] Segmentation Network," in Proc. of CVPR, 2019, pp. 8818-8826.
- T.-J. Yang, M. D. Collins, Y. Zhu et al., "DeeperLab: Single-shot Image [44] Parser," arXiv preprint arXiv:1902.05093, 2019.
- [45] B. Cheng, M. D. Collins, Y. Zhu et al., "Panoptic-DeepLab: A Simple, Strong, and Fast Baseline for Bottom-Up Panoptic Segmentation," in Proc. of CVPR, 2020, pp. 12475-12485.
- H. Wang, Y. Zhu, H. Adam, A. Yuille, and L.-C. Chen, "MaX-DeepLab: [46] End-to-End Panoptic Segmentation with Mask Transformers," in Proc. of CVPR, 2021, pp. 5463-5474.
- L. Beyer, A. Hermans, and B. Leibe, "Biternion Nets: Continuous Head [47] Pose Regression from Discrete Training Labels," in Proc. of GCPR. Springer, 2015, pp. 157-168.
- [48] B. Lewandowski, D. Seichter, T. Wengefeld et al., "Deep orientation: Fast and Robust Upper Body orientation Estimation for Mobile Robotic Applications," in Proc. of IROS, 2019, pp. 441-448
- [49] D. Seichter, B. Lewandowski, D. Höchemer, T. Wengefeld, and H.-M. Gross, "Multi-Task Deep Learning for Depth-based Person Perception in Mobile Robotics," in Proc. of IROS. IEEE, 2020, pp. 10497-10504.
- [50] A. Mousavian, D. Anguelov, J. Flynn, and J. Kosecka, "3D Bounding Box Estimation Using Deep Learning and Geometry," in Proc. of CVPR, 2017, pp. 7074-7082.
- [51] L. Liu, J. Lu, C. Xu, Q. Tian, and J. Zhou, "Deep Fitting Degree Scoring Network for Monocular 3D Object Detection," in Proc. of CVPR, 2019, pp. 1057-1066.
- [52] Y. Zhang, J. Lu, and J. Zhou, "Objects are different: Flexible monocular 3d object detection," in Proc. of CVPR, 2021, pp. 3289-3298
- [53] S. Vandenhende, S. Georgoulis, W. Van Gansbeke et al., "Multi-Task Learning for Dense Prediction Tasks: A Survey," TPAMI, 2021.
- [54] NVIDIA, "FasterTransformer," github.com/NVIDIA/FasterTransformer, Retrieved at 04.02.2023, 2022.
- [55] A. Paszke, S. Gross, F. Massa et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in Proc. of NeurIPS. Curran Associates, Inc., 2019, pp. 8024-8035.
- M. Roberts, J. Ramapuram, A. Ranjan et al., "Hypersim: A Photorealis-[56] tic Synthetic Dataset for Holistic Indoor Scene Understanding," in Proc. of ICCV, 2021.