

# Efficient Multi-Task RGB-D Scene Analysis for Indoor Environments

Daniel Seichter, Söhnke Benedikt Fishedick, Mona Köhler, and Horst-Michael Groß

Ilmenau University of Technology, Neuroinformatics and Cognitive Robotics Lab

98684 Ilmenau, Germany

daniel.seichter@tu-ilmenau.de, ORCID: 0000-0002-3828-2926

**Abstract**—Semantic scene understanding is essential for mobile agents acting in various environments. Although semantic segmentation already provides a lot of information, details about individual objects as well as the general scene are missing but required for many real-world applications. However, solving multiple tasks separately is expensive and cannot be accomplished in real time given limited computing and battery capabilities on a mobile platform. In this paper, we propose an efficient multi-task approach for RGB-D scene analysis (EMSANet) that simultaneously performs semantic and instance segmentation (panoptic segmentation), instance orientation estimation, and scene classification. We show that all tasks can be accomplished using a single neural network in real time on a mobile platform without diminishing performance – by contrast, the individual tasks are able to benefit from each other. In order to evaluate our multi-task approach, we extend the annotations of the common RGB-D indoor datasets NYUv2 and SUNRGB-D for instance segmentation and orientation estimation. To the best of our knowledge, we are the first to provide results in such a comprehensive multi-task setting for indoor scene analysis on NYUv2 and SUNRGB-D.

**Index Terms**—Multi-task learning, orientation estimation, panoptic segmentation, scene classification, semantic segmentation, NYUv2, SUNRGB-D

## I. INTRODUCTION

In computer vision, semantic scene understanding is often equated with semantic segmentation as it enables gaining precise knowledge about the structure of a scene by assigning a semantic label to each pixel of an image. However, this kind of knowledge is not sufficient for the agents in our ongoing research projects MORPHIA and CO-HUMANICS that require operating autonomously in their environments. Imagine a mobile robot that is supposed to navigate to a semantic entity, e.g., a specific chair within a group of chairs in the living room, as shown in Fig. 1. Performing such a high-level task requires a much broader understanding of the scene. First, even with a semantic map of the environment [1], the robot still needs to know which part of its environment belongs to the living room. Subsequently, it needs to be able to distinguish individual instances of the same semantic class, and, finally, for approaching the chair from the right direction, its orientation is required.

In this paper, we present an approach called Efficient Multi-task Scene Analysis Network (EMSANet) for tackling all

This work has received funding from the German Federal Ministry of Education and Research (BMBF) to the project MORPHIA (grant agreement no. 16SV8426) and from Carl-Zeiss-Stiftung to the project CO-HUMANICS.

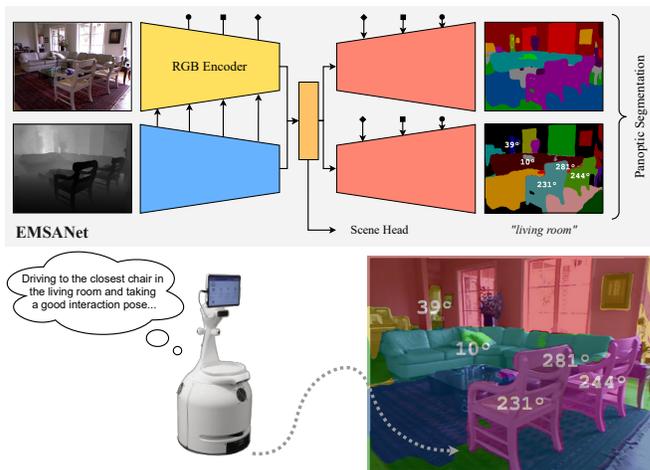


Fig. 1. Prediction of our proposed Efficient Multi-task Scene Analysis Network (EMSANet) that simultaneously performs panoptic segmentation, orientation estimation, and scene classification. With 24.5 FPS on an NVIDIA Jetson AGX Xavier it is well suited for mobile robotic applications. See Fig. 4 for semantic label colors. Color variations indicate individual instances.

the aforementioned challenges in order to accomplish such a high-level task. Our approach performs scene classification, semantic and instance segmentation (panoptic segmentation), as well as instance orientation estimation. However, given limited computing and battery resources on a mobile platform, solving all these tasks separately is expensive and cannot be accomplished in real time. Therefore, we design our approach to solve all aforementioned tasks using a single efficient multi-task network. Our approach extends ESANet [2], an efficient approach for semantic segmentation, by adding additional heads for tackling panoptic segmentation, instance orientation estimation, and scene classification. ESANet processes both RGB and depth data as input. As shown in [2], especially for indoor environments, depth data provide complementary geometric information that help analyzing cluttered indoor scenes. In this paper, we show that this also holds true for panoptic segmentation, instance orientation estimation, and scene classification. Thus, our approach also relies on both RGB and depth data.

Training such a multi-task approach requires comprehensive datasets. However, to the best of our knowledge, there is no real-world RGB-D indoor dataset encompassing ground-truth annotations for all aforementioned tasks. Therefore, we enrich the existing datasets NYUv2 [3] and SUNRGB-D [4] with

additional annotations for instance segmentation and instance orientation estimation. With this data at hand, we first train single-task baselines and subsequently combine multiple tasks in several multi-task settings. We experimentally show that all tasks can be solved using a single neural network in real time without diminishing performance – by contrast, the individual tasks are able to boost each other. Our full multi-task approach reaches 24.5 FPS on the mobile platform NVIDIA Jetson AGX Xavier, while achieving state-of-the-art performance. Thus, it is well suited for real-world applications on mobile platforms. In summary, our main contributions are:

- an efficient RGB-D multi-task approach for panoptic segmentation, scene classification, and instance orientation estimation (EMSA<sub>Net</sub>) including a novel encoding for instance orientations
- enriched annotations for NYUv2 and SUNRGB-D
- detailed experiments regarding performance of each task in single- and multi-task settings as well as corresponding inference throughputs on an NVIDIA Jetson AGX Xavier.

Our code, the additional annotations for NYUv2 and SUNRGB-D as well as the trained models are publicly available at: <https://github.com/TUI-NICR/EMSA<sub>Net</sub>>.

## II. RELATED WORK

In the following, we briefly summarize related work for each task. Moreover, we give some insights on combining tasks in multi-task settings.

### A. Semantic Segmentation

Architectures for semantic segmentation typically follow an encoder-decoder design to accomplish dense pixel-wise predictions. Well-known approaches such as PSPNet [5] or the DeepLab series [6]–[8] achieve good results but cannot be executed in real time on mobile platforms due to their low downsampling of intermediate feature representation. Thus, another line of research emerged, focusing on low inference time while still keeping high performance. For example, ERFNet [9] introduces a more efficient block by spatially factorizing the expensive  $3 \times 3$  convolution into a  $3 \times 1$  and a  $1 \times 3$  convolution and, thus, reduces computational effort. By contrast, SwiftNet [10] simply uses a pretrained ResNet18 [11] as encoder with early and high downsampling, resulting in low inference time but still good performance as well.

While the aforementioned approaches only process RGB data, especially for indoor applications, others [12]–[16] also incorporate depth data as they provide complementary geometric information that help analyzing cluttered scenes. Most approaches use two encoders for processing RGB and depth data (RGB-D) first separately and fuse the resulting features later in the network. However, almost all RGB-D approaches use deep and complex network architectures and do not focus on fast inference. By contrast, our recently published ESA<sub>Net</sub> [2] combines the merits of efficient and RGB-D semantic segmentation. It utilizes a carefully designed architecture featuring a dual-branch RGB-D ResNet-based encoder with high downsampling and spatially factorized convolutions

enabling fast inference. Our experiments in [2] reveal that processing both RGB and depth data with shallow backbones is superior to utilizing only RGB data and a deeper backbone.

Therefore, our approach follows ESA<sub>Net</sub> and extends its architecture with additional heads tackling the remaining tasks.

### B. Panoptic Segmentation

Panoptic segmentation [17] was introduced to unify semantic segmentation (assigning a class label to each pixel) and instance segmentation (assigning a unique id to pixels of the same instance) in a single task. In panoptic segmentation, semantic classes for countable objects are regarded as thing classes and represent foreground. Background classes, such as wall or floor – known as stuff classes – do not require instances. Thus, all associated pixels have the same instance id. Approaches for panoptic segmentation can be categorized in top-down, bottom-up, and end-to-end approaches. Top-down approaches typically extend two-stage instance segmentation approaches such as Mask R-CNN [18] with an additional decoder for semantic segmentation [19], [20]. Although top-down approaches typically achieve superior performance, they have several major drawbacks: As instance segmentation approaches can output overlapping instance masks, further logic is required to resolve these issues in order to merge instance and semantic segmentation without contradictions. Moreover, they require complex training and inference pipelines, making them less suitable for mobile applications. On the other hand, bottom-up approaches extend encoder-decoder-based architectures for semantic segmentation and separate thing classes into instances by grouping their pixels into clusters [21]–[23]. As bottom-up approaches neither require region proposals, estimating multiple masks independently, nor further refinement steps, their pipelines for training and inference are much simpler compared to top-down approaches. However, until Panoptic DeepLab [23] bottom-up approaches could not compete with top-down approaches in terms of panoptic quality. Nevertheless, both top-down and bottom-up approaches require additional logic for merging instance and semantic segmentation. The recently proposed MaX-DeepLab [24] follows another approach based on a novel dual-path transformer architecture [25] and attempts to directly predict the panoptic segmentation using an end-to-end pipeline. However, research for this kind of approaches currently focuses on establishing new architectures and not on fast and efficient inference.

Unlike for semantic segmentation, there are only a few approaches targeting efficiency [26]–[30]. However, their target hardware is different as they only report inference times on high-end GPUs. Execution on mobile platforms, such as an NVIDIA Jetson AGX Xavier, is expected to be much slower.

Our approach follows the bottom-up idea as it is straightforward to be integrated into ESA<sub>Net</sub> and expected to enable faster inference on mobile platforms.

### C. Orientation Estimation

Orientation estimation is often done along with 3D bounding box detection [31]–[33] and deeply integrated into such

architectures. Adapting these detectors to also accomplish dense predictions would require fundamental changes and, thus, is not suitable for our application. Another field of research strongly related to orientation estimation is person perception [34]–[38]. Besides estimating a person’s orientation inherently using its skeleton [34], there are also approaches directly estimating the orientation from patches [35]–[38]. This can be performed using either classification or regression. However, as shown in [35], classification adds further discretization inaccuracy and does not account well for periodicity. Therefore, approaches such as [35], [36] rely on regression and estimate the angle through its sine and cosine parts, which is often called Baternion encoding [35]. The same authors also proposed to use the von Mises loss function [35] instead of L1 or MSE loss as it further improves accounting periodicity and avoiding discontinuities.

Our approach follows the latter idea and formulates orientation estimation as regression. However, instead of using a patch-based approach, we propose a novel way to accomplish dense orientation estimation.

#### D. Scene Classification

Scene classification, i.e., assigning a scene label such as kitchen or living room to an input image, is similar to other classification tasks such as the ImageNet-Challenge [39]. Thus, well known architectures [11], [40]–[42] can be used.

#### E. Multi-task Learning

Multi-task learning refers to learning multiple tasks simultaneously in a single neural network. As these tasks commonly share at least some network parameters, inference is faster compared to using an independent network for each task. Moreover, in [43] it is shown, that dense prediction tasks may benefit from another when trained together. Especially early network layers are known to learn common features and, thus, can be shared among multiple tasks – in literature this is referred to as hard-parameter sharing [44]. Some approaches [45], [46] also exchange information in the task-specific heads, which is called soft-parameter sharing. However, when utilizing soft-parameter-shared task heads, these tasks cannot be decoupled anymore. This means that the whole network needs to be applied during inference, even though only a single task may be of interest. Therefore, our approach uses a hard-parameter shared RGB-D encoder and independent task-specific heads, not sharing any network parameters or information. We show that semantic and instance segmentation, instance orientation estimation as well as scene classification benefit from such a multi-task setting.

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

Our Efficient Multi-task Scene Analysis Network (EM-SANet) extends the encoder-decoder-based ESANet [2] for efficient RGB-D semantic segmentation. As shown in Fig. 2 (top), ESANet features two identical encoders, one for processing RGB images and one for depth images. For efficiency reasons, both encoders are based on a ResNet34 [11]

backbone. For even faster inference and improved accuracy, the  $3\times 3$  convolutions are spatially factorized, resulting in the NonBottleneck1D block (NBt1D) [9] (see Fig. 2 violet). At each resolution stage of the encoders, an attention-based mechanism is used to fuse depth features into the RGB branch enhancing its representation with additional geometric information. After the last fusion, a context module similar to the Pyramid Pooling Module in PSPNet [5] is attached. It incorporates context information at multiple scales using several branches with different pooling sizes (see [2] for details). The decoder is comprised of three decoder modules (see Fig. 2 light red), each is refining and upsampling the intermediate feature maps to gradually restore input resolution. This is done by a  $3\times 3$  convolution followed by three NonBottleneck1D blocks and a final learned upsampling by a factor of two. The learned upsampling (see Fig. 2 dark green) is initialized to mimic bilinear upsampling first. However, as its weights are not fixed, the network is able to learn to combine adjacent features in a more useful manner during training. Additional encoder-decoder skip connections further help to restore spatial details that were lost during downsampling in the encoders. Following the last decoder module, a  $3\times 3$  convolution maps the features to semantic classes. Finally, two additional learned upsamplings restore the input resolution. The entire network is trained end-to-end with additional side outputs and multi-scale supervision as depicted in Fig. 2.

ESANet builds a strong and efficient baseline for semantic segmentation. However, the architecture is specifically tailored for semantic segmentation. In order to improve its generalization capability for the remaining tasks, we further add a slight dropout with rate of 0.1 to all NonBottleneck1D blocks. Furthermore, we change the initialization in all RGB-D fusion modules to He-initialization [47] and force zero-initialization [48] in all NonBottleneck1D blocks. Finally, to incorporate the loss of other tasks more effectively, we do not reduce the accumulated loss by the sum of the applied semantic class weights but only by the number of all pixels across all outputs of the network.

Next, we present the extension to a multi-task network.

#### A. Panoptic Segmentation

For panoptic segmentation, a second decoder for instance segmentation is required. As shown in Fig. 2 (middle), our instance decoder follows the same architecture as the semantic decoder except the task-specific heads. The instance encoding follows the implementation of Panoptic DeepLab [23]: Instances are represented by their center of mass encoded as small 2D Gaussian within a heatmap, similar to keypoint estimation in other domains [34]. With an additional head, the instance decoder also predicts offset vectors for each pixel pointing towards a corresponding instance center in x and y direction. As instances are only required for pixels belonging to thing classes – i.e., all classes except wall, floor, and ceiling – a corresponding foreground mask is derived from the semantic segmentation. All thing pixels are then grouped into class-agnostic instances by combining both instance centers

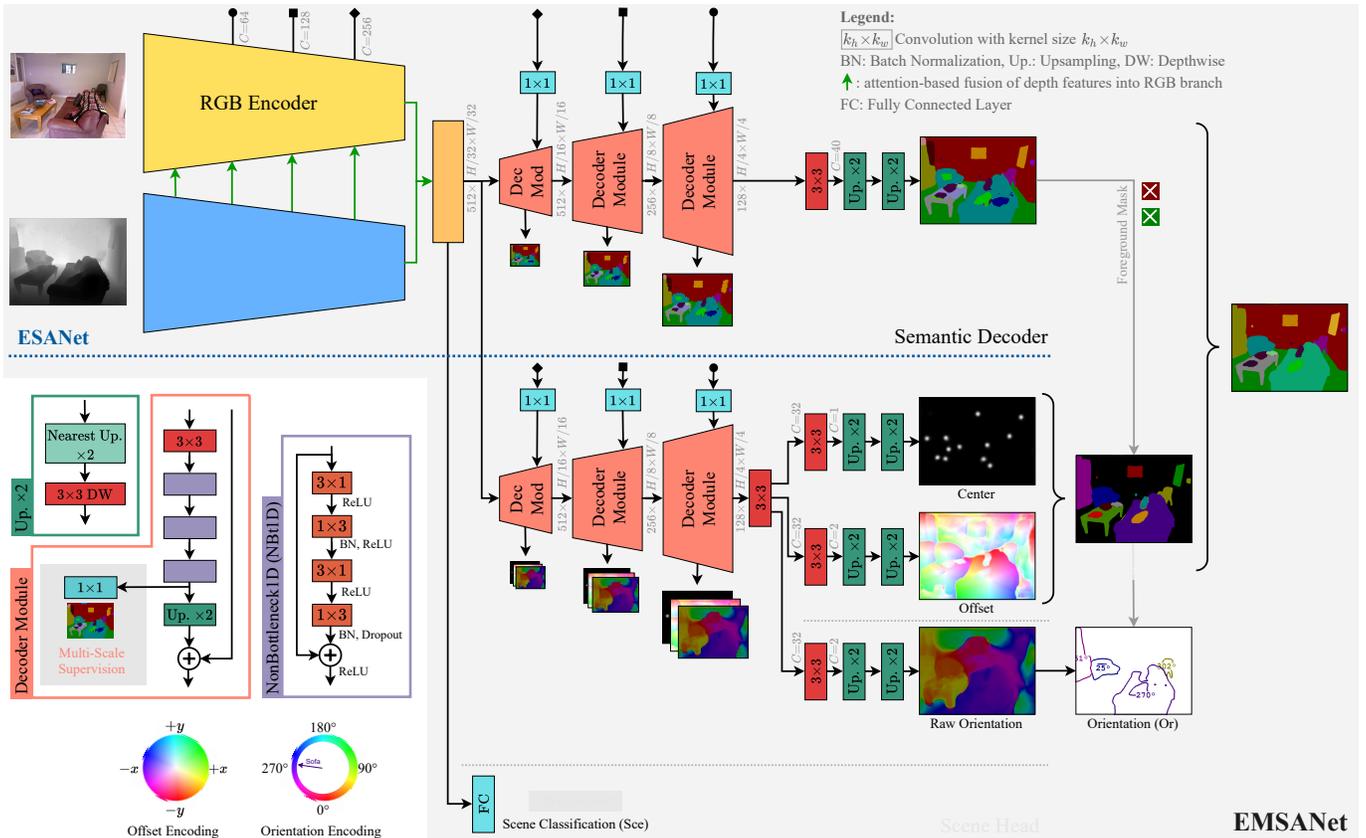


Fig. 2. Architecture of our Efficient Multi-task Scene Analysis Network (EMSANet) extending ESANet [2] for semantic segmentation (top) with an additional decoder for instance segmentation and instance orientation estimation as well as a head for scene classification. See Fig. 4 for semantic label colors.

and offset predictions. The semantic class for each instance is derived by a majority vote from the semantic segmentation.

Similar to Panoptic DeepLab, we use MSE loss for center prediction and L1 loss for offset prediction. However, unlike Panoptic DeepLab, we mask predicted centers using the ground-truth instance mask instead of the thing-class mask to account for missing instance annotations in the ground truth. We also adopt their postprocessing, including thresholding and keypoint non-maximum suppression using max-pooling for the centers and the final merging of instance and semantic segmentation putting more focus on instances. However, we faced some problems when applying their training regime. Panoptic DeepLab uses linear outputs for estimating both centers and absolute offsets. This results in losses being unbounded and quite imbalanced, and, thus, requires a carefully tuned initialization and loss weights such as 200:0.1 for center:offsets as used in their implementation. Moreover, absolute offset vectors do not generalize to varying input resolutions. To address these issues, we use sigmoid activation for instance centers and a tanh activation for encoding relative instance offsets. Thus, the outputs are bounded within  $[0, 1]$  and  $[-1, 1]$ , respectively. We observed great improvements in terms of stability during optimization and performance for instance segmentation.

### B. Instance Orientation Estimation

Our approach further predicts the orientation for instances of thing classes relevant for our indoor scenario, i.e., cabinet,

bed, chair, sofa, bookshelf, shelves, dresser, refrigerator, tv, person, nightstand, and toilet. The orientation is crucial when approaching objects from the right direction (e.g., chairs or persons) or to restrict waiting positions (e.g., the robot should not wait in the sightline to a TV or in front of a freestanding chair or cabinet). To accomplish this, as shown in Fig. 2 (middle), our instance decoder also predicts orientations as continuous angles around the axis perpendicular to the ground (see bottom-left legend for orientation encoding in Fig. 2). Instead of relying on a patch-based orientation estimation, we follow our dense prediction design and propose to predict the orientation for all pixels of an instance. This way, the instance awareness of our instance decoder can further be strengthened. Moreover, to determine an instance’s orientation, we are able to average multiple predictions approximating an ensemble effect. We use the biternion encoding [35] and the von Mises loss function [35] to account for the periodicity of angles and to avoid discontinuities in the loss.

### C. Scene Classification

For scene classification, as shown in Fig. 2 (bottom), we simply apply a fully-connected layer on top of the context module. However, as scene classification requires global context, we connect the fully-connected layer directly to the global average pooled branch of the context module. Due to the noisy nature of scene classes, we further utilize label smoothing during training.

## IV. DATASETS

Training our proposed multi-task approach is challenging as it requires comprehensive data providing annotations for all tasks. Moreover, our approach relies on both RGB and depth images as input. Based on these requirements and our application scenario, below, we examine common RGB-D datasets for their suitability. Furthermore, we describe how we enriched these datasets to enable training our multi-task approach. Additional annotations are publicly available.

*NYUv2*: The NYUv2 dataset [3] provides dense annotations for both semantic and instance segmentation. For semantics, we use the common 40 classes setting. However, this may lead to very small instances getting assigned to misleading classes, e.g. door knobs are assigned to cabinet, dresser, or nightstand. To avoid such bad assignments, we restrict instances to have at least an area of 0.25% of the image area. For enabling panoptic segmentation, we declare wall, floor, and ceiling to be background (stuff classes) and consider the remaining classes as thing classes. In addition to these dense annotations, NYUv2 also provides ground-truth labels for scene classification. However, annotations for instance orientations are missing so far. Therefore, we manually annotated the orientation for instances of the semantic classes mentioned in Sec. III-B. Due to perspective distortions, exact annotation of the orientation as egocentric angle around the axis perpendicular to the ground is not possible in the pure RGB image, which is why we annotated them directly in a point cloud.

*SUNRGB-D*: The SUNRGB-D dataset [4] combines multiple indoor RGB-D datasets, including NYUv2, and enriches them with additional annotations, making SUNRGB-D to be one of the most important datasets for real-world applications. The dataset comes with annotations for scene classification and semantic segmentation. Compared to NYUv2, the last three semantic filling classes, i.e., otherstructure, otherfurniture, and otherprop, are omitted and assigned to void. Moreover, some semantic annotations in the NYUv2 part have further been assigned to the void class, resulting in minor differences to the original NYUv2 dataset. Unfortunately, annotations for instance segmentation and instance orientation estimation are missing. However, fortunately, SUNRGB-D also provides 3D bounding boxes, each with a class label and orientation, that can be used for instance extraction. To obtain instances, we first created a mapping between semantic and box classes. Subsequently, we matched the box clusters with the semantic point cloud in 3D. During matching, a unique instance label was assigned to all pixels belonging to the same semantic class. This way, an instance mask as well as the orientation could be extracted for each bounding box. However, one limitation of this approach is that not all objects within a scene were annotated with a 3D bounding box, making the instance masks more sparse. To compensate this to some extent, we also merged the instance masks and orientations of NYUv2 back to SUNRGB-D. For panoptic segmentation, we consider the same semantic classes to belong to stuff as for NYUv2.

TABLE I

Overview about the datasets used for our multi-task approach.

	Split	# Images	# Instances	# Orientations
NYUv2	train	795	12,092	2,696
	test	654	9,874	2,069
SUNRGB-D	train	5,285	18,171	13,076
	test	5,050	16,961	12,440
Hypersim	train	57,443	3,009,566	-
	valid	7,286	261,677	-
	test	7,690	374,052	-

TABLE II

Distribution of scene classes for all datasets and their splits.

	NYUv2		SUNRGB-D		Hypersim		
	train	test	train	test	train	valid	test
void	29	0	745	507	1,370	0	0
bathroom	63	58	331	293	3,378	400	300
bedroom	192	191	558	526	4,386	499	400
dining room	66	55	311	296	1,894	200	100
discussion room	5	0	691	753	1,013	100	0
hallway	0	0	222	151	400	100	0
kitchen	125	110	288	297	6,221	400	600
living room	114	107	274	250	18,287	3,017	2,855
office	88	78	820	792	5,026	100	389
other indoor	113	55	1,041	1,180	15,068	2,370	2,846
stairs	0	0	4	5	400	100	200

*Hypersim*: Unlike SUNRGB-D and NYUv2, Hypersim [49] is a photo-realistic synthetic dataset. For its creation, virtual cameras were placed in 461 professionally rendered 3D scenes, resulting in 77,400 samples, of which we use 72,419. We blacklisted the remaining samples due to several scene or trajectory issues, i.e., void/single semantic label only, missing textures, or invalid depth. Each sample provides an RGB-D image, a mask for semantics and instances, instance orientations, and a scene label. However, the annotations for instance orientation are not consistent and, thus, cannot be used without further manual refinement. As Hypersim adopts the NYUv2 semantic classes, the same partitioning for stuff and thing can be applied for panoptic segmentation.

*Final remarks and further adjustments*: Due to the additional annotations, both NYUv2 and SUNRGB-D are suitable for training our full multi-task approach. HyperSim provides high-quality synthetic data and, thus, is well suited for pre-training. Tab. I summarizes important statistics for all datasets used for training and evaluating our multi-task approach. For scene classification, we further created an own spectrum of classes that unifies the classes in all datasets and accounts for similar classes. The resulting spectrum is tailored for indoor applications and is comprised of the classes listed in Tab. II. Note that the void class is used for images with unclear assignment that may disturb the learning process. Furthermore, images that show indoor scenes but cannot be assigned to one of the mentioned classes are considered as other indoor.

## V. EXPERIMENTS

We evaluate our approach in several settings on the indoor datasets NYUv2 and SUNRGB-D. First, we use the smaller NYUv2 dataset to elaborate suitable hyperparameters and task weights. We establish single-task baselines for each task and, subsequently, compare them to several multi-task settings.

Finally, we extend our studies to SUNRGB-D and Hypersim to examine the applicability to larger datasets and the relevance of synthetic data for pretraining.

### A. Implementation Details

Our architecture as well as the pipelines for training and evaluation are implemented using PyTorch [50]. We used pre-trained weights on ImageNet [39] to initialize both encoders and trained each network for 500 epochs with a batch size of 8. For optimization, we used SGD with momentum of 0.9 and a small weight decay of 0.0001. For determining a suitable learning rate, we performed a grid search with values of  $\{0.00125, 0.0025, 0.005, 0.01, 0.02, 0.03, 0.04, 0.08, 0.12\}$ . The learning rate was further adapted during training using a one-cycle learning rate scheduler. To increase the number of samples, we augmented images using random scaling, cropping, and flipping. For RGB images, we further applied slight color jittering in HSV space.

For postprocessing instance centers, we first apply a threshold of 0.1 and max-pooling with pooling size of 17 to perform keypoint non-maximum suppression, and finally filter the top-64 instances. The pooling size results in the network not being able to predict instance centers closer than 8 pixels away from each other. However, for both NYUv2 and SUNRGB-D, this decision affects less than 1% of the instances. For further details and other hyperparameters, we refer to our implementation available on GitHub.

### B. Metrics

As we focus on fast inference, we do not apply any testing tricks such as horizontal flipping or multiscale inputs. Before obtaining any performance metric for dense predictions, we resize the predictions to the full resolution of the ground truth.

*Semantic Segmentation (Sem):* As common for this task, we use the mean intersection over union (mIoU).

*Panoptic Segmentation:* The common metric for panoptic segmentation is the panoptic quality (PQ) [17]. The panoptic quality  $PQ_c$  for each class  $c$  is determined by the product of the recognition quality ( $RQ_c$ ), i.e. the percentage of correctly detected instances similar to  $F_1$ -score, and the segmentation quality ( $SQ_c$ ), i.e., the segmentation accuracy similar to mIoU but only for matched segments. These metrics are typically averaged across all stuff (st) and thing (th) classes, resulting in  $RQ^{st}$ ,  $SQ^{st}$ ,  $PQ^{st}$  and  $RQ^{th}$ ,  $SQ^{th}$ ,  $PQ^{th}$ , respectively. It is also common to average the three metrics independently of stuff and things across all classes, resulting in RQ, SQ, and PQ. Note that determining the metrics this way results in PQ to be typically not equal to  $RQ \cdot SQ$ . For SUNRGB-D, we further ignore the classes floor mat and shower curtain as no instances of these classes occur in the test split.

*Instance Segmentation (Ins):* For instance segmentation, we also stick to panoptic quality instead of reporting the average precision (AP) as this would require assigning a confidence score to each instance [51]. Moreover, AP and PQ track closely, which is why the latter also evaluates instance segmentation in a meaningful way [17].

*Orientation Estimation (Or):* For evaluating instance orientations, we use the mean absolute angular error (MAAE) in degrees similar to [36], [38], i.e., in contrast to the other metrics, lower is better, and the maximum error is 180. We report the MAAE for two settings: 1) independently of other tasks, i.e., using the ground-truth instances, and 2) for matched instances after panoptic merging. Note that we do not penalize unmatched instances for the latter setting.

*Scene Classification (Sce):* As the scene class labels are imbalanced, we evaluate scene classification with the balanced accuracy (bAcc).

### C. Single-task Setting

We aim at solving multiple tasks at once using a single neural network. To be able to elaborate whether this diminishes or boosts the performance of individual tasks, we first conducted experiments in a single-task setting. Furthermore, the goal was to examine how the new tasks behave for different modalities, and how the additional network parts affect inference time on a mobile platform. Note that performing instance segmentation solely requires semantics and a foreground mask. We use the ground-truth semantic segmentation in this case, leading to  $RQ^{st}$ ,  $SQ^{st}$ , and  $PQ^{st}$  to always be equal to 1. Moreover, as instance orientation estimation requires instance masks, we rely on the ground-truth instances as well.

Fig. 3 shows the results of this set of experiments. It becomes obvious that all tasks are able to benefit from incorporating complementary information processed in an additional depth branch instead of processing RGB only. For semantic and instance segmentation, the results coincide with our findings in [2] that, whenever possible, processing both RGB and depth should be preferred over applying a more sophisticated single-modality RGB encoder. The results further show that depth is crucial for estimating the orientation more accurately, and that RGB is essential for scene classification. Finally, when comparing different backbones, it becomes obvious that, for all tasks except of scene classification, backbones featuring the NonBottleneck1D block (NBt1D) lead to better results in terms of both performance and inference throughput compared to their counterparts with BasicBlock. Even more, ResNet34 NBt1D often competes with the more complex ResNet50 while enabling faster inference. Therefore, we stick to ResNet34 NBt1D for the remaining experiments.

### D. Multi-task Setting

Learning multiple tasks using a single neural network is challenging, as the tasks may influence each other. Thus, tuning the task weights, which balance the losses to each other, is crucial. With uncertainty weighting [52], GradNorm [53], Dynamic Weight Average [54], and VarNorm [55], several approaches for determining the weights automatically have been proposed. Unfortunately, none of them led to good performance in our scenario. Therefore, we performed extensive experiments to determine suitable task weights. We first combined two tasks at a time to elaborate essential relations between tasks. With these findings at hand, we were able to

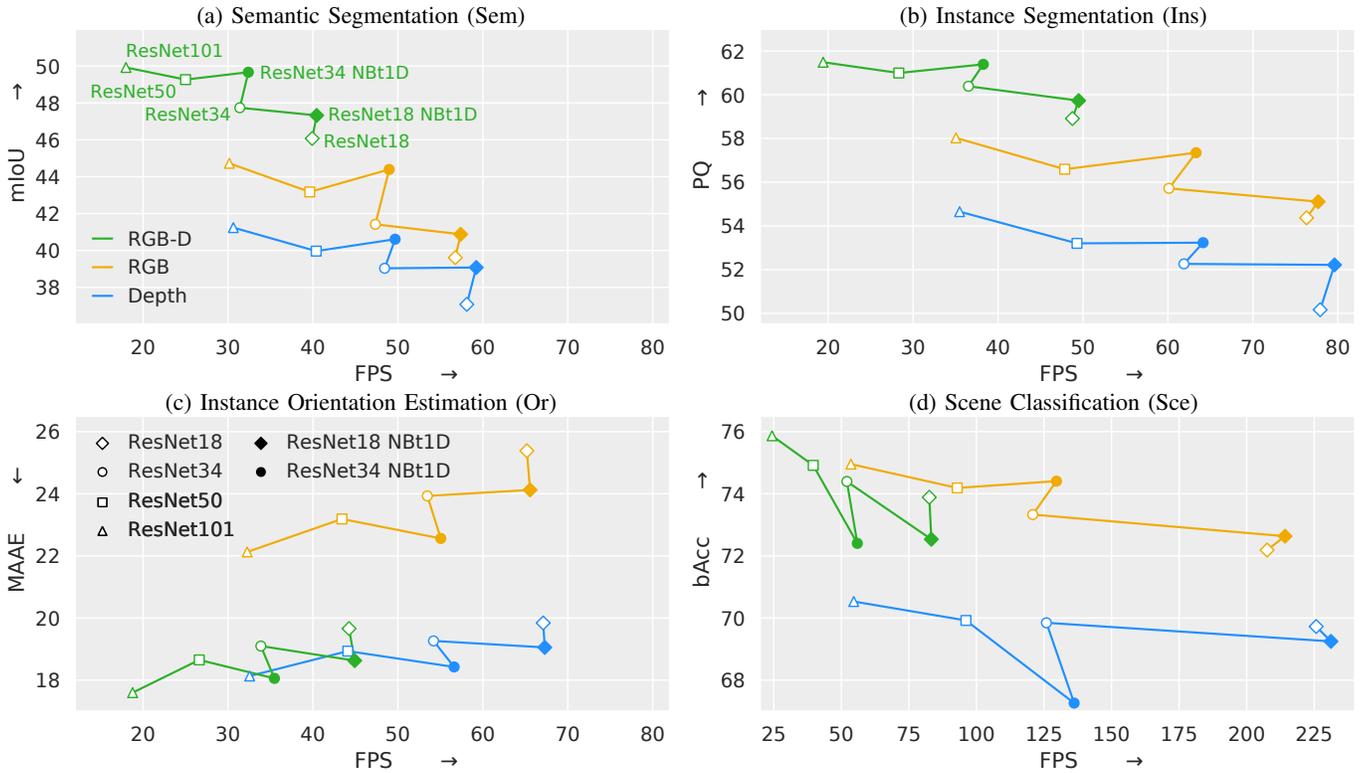


Fig. 3. Results on NYUv2 test split when performing each task in a single-task setting with various backbones over the inference throughput in frames per second measured on an NVIDIA Jetson AGX Xavier (Jetpack 4.6, TensorRT 8, Float16). See Sec. V-B for metric abbreviations.

restrict the search space for the full multi-task setting. Tab. III summarizes the best results and compares them to their single-task counterparts of Fig. 3. Furthermore, it lists the applied task weights, the learning rate, and the achieved frames per second on an NVIDIA Jetson AGX Xavier.

*Sem + Sce:* As shown in Tab. III (MT I), combining both tasks requires a much larger weight for semantic segmentation to reach its single-task performance. However, even when putting more weight on semantic segmentation, scene classification benefits from such a setting, already closing the gap shown in Fig. 3 (d) for ResNet34 NBt1D and ResNet34. This shows that knowledge about the individual parts of a scene is shared and helps to classify the scene.

*Sem + Ins:* This setting allows obtaining panoptic results with predicted semantic segmentation for the first time. As the semantic segmentation provides the semantics and the foreground mask for instance segmentation, combining these two tasks is crucial for our multi-task system. As shown in Tab. III (MT II), the best PQ is achieved by putting more focus on instance segmentation. The mIoU obtained for the semantic decoder indicates that the network further benefits from performing both tasks in conjunction. When keeping in mind, that the PQ for the instance decoder is computed using the ground-truth semantic segmentation, it is reasonable that the PQ for the panoptic results is lower. With ground-truth semantic segmentation, the network reaches an RQ of 70.15 and a SQ of 85.78 (not listed in Tab. III). This shows that the drop in PQ is mainly due to a loss in RQ. We observed that this is mostly caused by small instances, which are not part

of the predicted foreground mask.

*Ins + Or:* The results in Tab. III (MT III) reveal that both tasks can be performed using a single decoder. Compared to the single-task baseline, this setting slightly improves orientation estimation, almost surpassing the reachable level for annotating orientations in 3D. However, even when putting more weight on instance segmentation, we always observed a slight drop in PQ.

*Sem + Sce + Ins + Or:* With the findings of the aforementioned dual-task experiments at hand, we combined all tasks in a single neural network. The best result is presented in Tab. III (MT IV). It becomes obvious that both semantic segmentation and scene classification greatly benefit from the entire multi-task setting. Instance segmentation and instance orientation estimation almost reach the same level of accuracy as when performed in single-task settings. The panoptic results, i.e., after merging semantic and instance predictions, are similar to the multi-task setting MT II. The mIoU obtained after merging is slightly lower than before merging but still at a similar level. This indicates that the applied merging of both predictions with focus on instances does not diminish the semantic segmentation result. The detailed breakdown of the IoUs in Fig. 4 shows that this holds true for almost all classes. Finally, when taking a look at the results for orientation estimation, it can be seen that the orientation error after merging is lower than the error for the instance decoder. However, this does not necessarily indicate better results, as the MAAE after panoptic merging only represents instances that could be matched.

TABLE III

Results obtained on NYUv2 test split when training EMSANet with ResNet34 NBt1D backbone in various multi-task settings. See Sec. V-B for details on the reported metrics. Panoptic results are obtained after merging semantic and instance prediction. Legend: italic: metric used for determining the best checkpoint, \*: best result within the same run, Lr: learning rate, pre.: additional pretraining on Hypersim, FPS: frames per second on an NVIDIA Jetson AGX Xavier (Jetpack 4.6, TensorRT 8, Float16).

	Task Weights	Lr	Semantic Decoder	Instance Decoder		Scene Head	Panoptic Results (after merging)					FPS $\uparrow$	
			mIoU $\uparrow$	PQ $\uparrow$	MAAE $\downarrow$	bAcc $\uparrow$	mIoU $\uparrow$	PQ $\uparrow$	RQ $\uparrow$	SQ $\uparrow$	MAEE $\downarrow$		
ST	Semantic Segmentation (Sem)	—	0.04	49.66	—	—	—	—	—	—	—	32.4	
	Instance Segmentation (Ins)	—	0.08	—	61.39	—	—	—	—	—	—	38.3	
	Orientation Estimation (Or)	—	0.02	—	—	18.06	—	—	—	—	—	35.5	
	Scene Classification (Sce)	—	0.02	—	—	—	72.40	—	—	—	—	55.9	
MT I	Sem + Sce	3 : 1	0.02	<i>49.57</i>	—	—	74.29	—	—	—	—	32.2	
	Sem + Sce *			<i>49.57</i>	—	—	77.30	—	—	—	—		
MT II	Sem + Ins	1 : 3	0.03	50.22	60.90	—	—	50.24	<i>43.74</i>	52.46	82.43	—	25.8
	Sem + Ins *			50.22	61.61	—	—	50.54	<i>43.74</i>	52.50	82.63	—	
MT III	Ins + Or	3 : 1	0.04	—	<i>59.72</i>	17.66	—	—	—	—	—	—	35.5
	Ins + Or *			—	<i>59.72</i>	17.53	—	—	—	—	—	—	
MT IV	Sem + Sce + Ins + Or	1 : 0.25 : 3 : 1	0.04	50.97	61.35	19.01	76.46	50.54	<i>43.56</i>	52.20	82.48	16.38	24.5
	Sem + Sce + Ins + Or *			51.15	61.53	18.93	78.18	51.31	<i>43.56</i>	52.27	82.70	15.76	
MT V	Sem + Sce + Ins + Or (pre.)	1 : 0.25 : 3 : 1	0.01	53.34	64.41	18.84	75.25	53.79	<i>47.38</i>	55.95	83.74	15.91	24.5
	Sem + Sce + Ins + Or (pre.) *			53.55	64.98	18.27	76.98	54.00	<i>47.38</i>	55.99	84.08	15.56	

### E. Results on SUNRGB-D

After elaborating suitable multi-task parameters, we also applied our approach to the larger SUNRGB-D dataset. However, as instances in SUNRGB-D are more sparse, we put less weight on the instance decoder and used the mIoU for determining the best checkpoint. The results are listed in Tab. IV. Compared to the single-task baselines, our multi-task approach reaches slightly better performance for semantic segmentation and scene classification. The result for orientation estimation is still suitable for real-world application. As shown in Fig. 5, the network generalizes well for instance segmentation even though annotations are much more sparse.

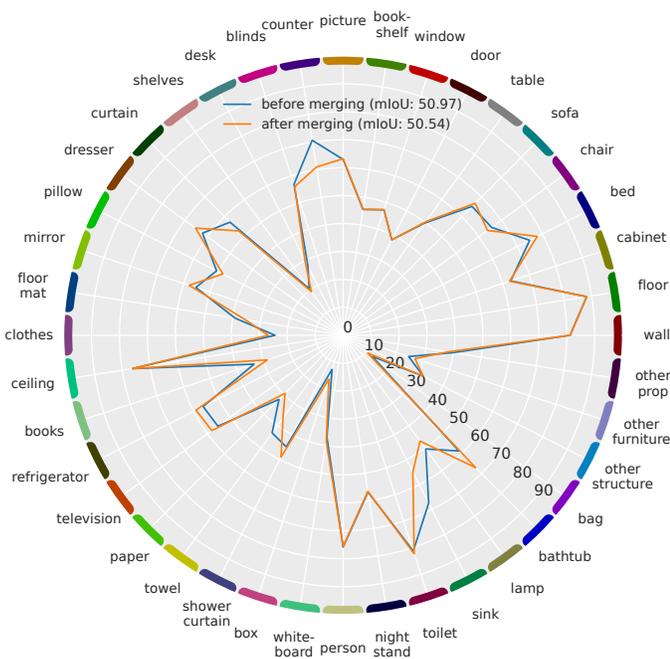


Fig. 4. Semantic IoUs on NYUv2 test split for the full multi-task network (MT IV) before and after merging semantic and instance segmentation.

Note that the panoptic quality does not account for these areas as they are labeled as void class in the ground truth. However, as shown in the last row, missing instance centers can still lead to assigning pixels far away to the same instance, lowering the PQ and mIoU after merging. We already observed a great benefit when masking centers based on the ground-truth instance mask during training, as proposed in Sec. III-A. For real-world application, we further tackle this issue by thresholding predicted offsets after shifting and assign an unknown instance label if they are too far away from a center.

### F. Additional Pretraining on Hypersim

Finally, we examined how pretraining on the synthetic Hypersim dataset affects the performance of our derived multi-task setting for both NYUv2 and SUNRGB-D. For further details on the pretraining, we refer to our implementation. The results are shown in Tab. III (MT V) and Tab. IV (MT V). It turns out that, for NYUv2, especially mIoU and PQ greatly benefit from additional pretraining, while, for SUNRGB-D, only the performance of instance-related tasks is improved. This can be deduced to the fact that SUNRGB-D alone is already much larger than NYUv2. Fig. 6 shows qualitative

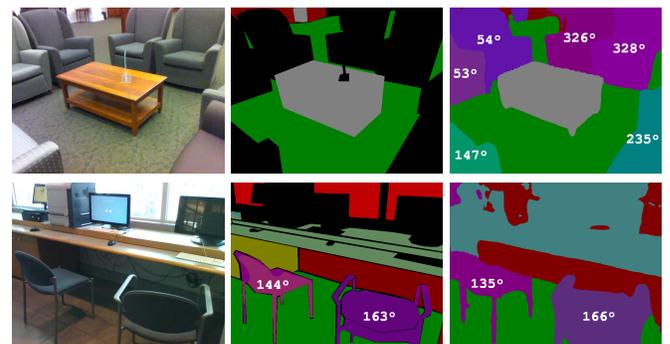


Fig. 5. Qualitative results on SUNRGB-D test split highlighting faced challenges (Sec V-E) as RGB image, ground-truth panoptic segmentation with orientations, and predicted panoptic segmentation with orientations.

TABLE IV

Results obtained on SUNRGB-D test split when training EMSANet with ResNet34 Nbt1D backbone in both single-task and multi-task settings. See Sec. V-B for details on the reported metrics. Panoptic results are obtained after merging semantic and instance prediction. Legend: italic: metric used for determining the best checkpoint, \*: best result within the same run, Lr: learning rate, pre.: additional pretraining on Hypersim.

	Task	Weights	Lr	Semantic Decoder	Instance Decoder		Scene Head	Panoptic Results (after merging)					
				mIoU <sup>↑</sup>	PQ <sup>↑</sup>	MAAE <sup>↓</sup>	bAcc <sup>↑</sup>	mIoU <sup>↑</sup>	PQ <sup>↑</sup>	RQ <sup>↑</sup>	SQ <sup>↑</sup>	MAEE <sup>↓</sup>	
ST	Semantic Segmentation (Sem)	—	0.005	48.23	—	—	—	—	—	—	—	—	—
	Instance Segmentation (Ins)	—	0.01	—	60.99	—	—	—	—	—	—	—	—
	Orientation Estimation (Or)	—	0.005	—	—	13.68	—	—	—	—	—	—	—
	Scene Classification (Sce)	—	0.001	—	—	—	58.66	—	—	—	—	—	—
MT IV	Sem + Sce + Ins + Or	1 : 0.25 : 2 : 0.5	0.005	48.39	60.60	16.81	61.83	45.56	50.15	58.14	84.85	14.24	—
	Sem + Sce + Ins + Or *	—	—	48.39	61.48	16.70	62.66	45.66	50.53	58.66	85.20	14.15	—
MT V	Sem + Sce + Ins + Or (pre.)	1 : 0.25 : 2 : 0.5	0.0025	48.47	64.24	18.40	57.22	44.18	52.84	60.67	86.01	14.10	—
	Sem + Sce + Ins + Or (pre.)*	—	—	48.47	64.82	17.94	59.39	45.04	53.35	61.31	86.25	14.04	—

results from pretraining on Hypersim and subsequent training on NYUv2. The latter represents our best network for NYUv2.

### G. Comparison to State of the Art

Comparing our approach to the state of the art is challenging as tasks such as orientation estimation or scene classification have not been considered so far in related work due to missing annotations and a deviating class spectrum. Moreover, panoptic segmentation has not yet been attempted on NYUv2 or SUNRGB-D. Therefore, we first established comprehensive single-task baselines (see Fig. 3) covering common backbones ranging from sophisticated backbones to more efficient backbones that also enable mobile application. Beyond that, we further trained well-known approaches for panoptic segmentation and scene classification on NYUv2, as shown in Tab. V. For Panoptic DeepLab, we applied the same parameters as described in Sec. III-A for keypoint non-maximum suppression to postprocess instance centers. To summarize, our proposed lightweight EMSANet achieves comparable or even better results than other approaches. Moreover, larger backbones do not necessarily improve performance but significantly increase resource requirements.

## VI. CONCLUSION

In this paper, we have proposed an efficient RGB-D multi-task approach for panoptic segmentation, instance orientation estimation, and scene classification, called EMSANet. For

TABLE V

Comparison to other state-of-the-art approaches on NYUv2 test split without test-time augmentation. Legend: italic: metric used for determining the best checkpoint, \*: (re)training on our enriched NYUv2 dataset, †: ten-crop evaluation at 224×224. pre.: additional pretraining on Hypersim.

	Backbone	Mod.	mIoU <sup>↑</sup>	PQ <sup>↑</sup>	bAcc <sup>↑</sup>
ESANet [2]	2×R34-Nbt1d	RGB-D	50.30	—	—
ShapeConv [16]	ResNext101_32x8d	RGB-D	50.20	—	—
Panoptic DeepLab* [23]	R50	RGB	39.42	30.99	—
	R101	RGB	42.55	35.32	—
MobileNetV2* [41]	$\alpha=1$	RGB	—	—	69.30 <sup>†</sup>
EfficientNet* [42]	B0	RGB	—	—	70.83 <sup>†</sup>
Mod. ResNet34 (ours)	R34-Nbt1d	RGB	—	—	70.25 <sup>†</sup>
<b>EMSANet (ours)</b>	R34-Nbt1d	RGB	44.66	37.69	70.88
	2×R34-Nbt1d	RGB-D	50.97	43.56	76.46
	2×R101-Nbt1d	RGB-D	50.83	45.12	77.41
<b>EMSANet pre. (ours)</b>	2×R34-Nbt1d	RGB-D	53.34	47.38	75.25

training and evaluation, we have enriched the annotations of the common RGB-D indoor datasets NYUv2 and SUNRGB-D, which we also make publicly available. To the best of our knowledge, we are the first to provide results in such a comprehensive multi-task setting for indoor scene analysis. We have shown that all tasks can be solved using a single multi-task network. Moreover, the individual tasks can benefit from each other when trained together. Due to the efficient design, our approach enables fast inference, i.e. 24.5 FPS on an NVIDIA Jetson AGX Xavier and, thus, is well suited for mobile robotic applications.



Fig. 6. Qualitative results as RGB image overlaid with predicted panoptic segmentation, predicted scene class, and, for NYUv2, estimated orientations.

## REFERENCES

- [1] D. Seichter, P. Langer, T. Wengefeld, B. Lewandowski, D. Höchmer, and H.-M. Gross, "Efficient and robust semantic mapping for indoor environments," in *Proc. of ICRA*, 2022.
- [2] 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. 13 525–13 531.
- [3] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor Segmentation and Support Inference from RGBD Images," in *Proc. of ECCV*, 2012.
- [4] S. Song, S. P. Lichtenberg, and J. Xiao, "SUN RGB-D: A RGB-D Scene Understanding Benchmark Suite," in *Proc. of CVPR*, 2015, pp. 567–576.
- [5] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. of CVPR*, 2017, pp. 2881–2890.
- [6] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Semantic image segmentation with deep convolutional nets and fully connected crfs," in *Proc. of ICLR*, 2015.
- [7] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking Atrous Convolution for Semantic Image Segmentation," *arXiv preprint arXiv:1706.05587*, 2017.
- [8] 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.
- [9] E. Romera, J. M. Alvarez, L. M. Bergasa, and R. Arroyo, "ERFNet: Efficient Residual Factorized ConvNet for Real-Time Semantic Segmentation," *ITS*, pp. 263–272, 2018.
- [10] M. Oršić, I. Krešo, P. Bevandić, and S. Šegvić, "In Defense of Pre-trained ImageNet Architectures for Real-time Semantic Segmentation of Road-driving Images," in *Proc. of CVPR*, 2019, pp. 12 607–12 616.
- [11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of CVPR*, 2016, pp. 770–778.
- [12] 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.
- [13] X. Hu, K. Yang, L. Fei, and K. Wang, "ACNet: Attention Based Network to Exploit Complementary Features for RGBD Semantic Segmentation," in *Proc. of ICIP*, 2019.
- [14] A. Valada, R. Mohan, and W. Burgard, "Self-supervised model adaptation for multimodal semantic segmentation," *IJCV*, 2019.
- [15] 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.
- [16] J. 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 CVPR*, 2021, pp. 7088–7097.
- [17] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár, "Panoptic segmentation," in *Proc. of CVPR*, 2019, pp. 9404–9413.
- [18] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," *Proc. of ICCV*, pp. 2961–2969, 2017.
- [19] A. Kirillov, R. Girshick, K. He, and P. Dollár, "Panoptic feature pyramid networks," in *Proc. of CVPR*, 2019, pp. 6399–6408.
- [20] Y. Xiong *et al.*, "Upsnet: A unified panoptic segmentation network," in *Proc. of CVPR*, 2019, pp. 8818–8826.
- [21] T.-J. Yang *et al.*, "Deeplab: Single-shot image parser," *arXiv preprint arXiv:1902.05093*, 2019.
- [22] N. Gao *et al.*, "Ssap: Single-shot instance segmentation with affinity pyramid," in *Proc. of ICCV*, 2019, pp. 642–651.
- [23] B. Cheng *et al.*, "Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation," in *Proc. of CVPR*, 2020, pp. 12 475–12 485.
- [24] H. Wang, Y. Zhu, H. Adam, A. Yuille, and L.-C. Chen, "Max-deeplab: End-to-end panoptic segmentation with mask transformers," in *Proc. of CVPR*, 2021, pp. 5463–5474.
- [25] A. Vaswani *et al.*, "Attention is all you need," *Proc. of NeurIPS*, vol. 30, 2017.
- [26] R. Mohan and A. Valada, "Efficientps: Efficient panoptic segmentation," *IJCV*, vol. 129, no. 5, pp. 1551–1579, 2021.
- [27] W. Hong, Q. Guo, W. Zhang, J. Chen, and W. Chu, "Lpsnet: A lightweight solution for fast panoptic segmentation," in *Proc. of CVPR*, 2021, pp. 16 746–16 754.
- [28] C.-Y. Chang, S.-E. Chang, P.-Y. Hsiao, and L.-C. Fu, "Epsnet: efficient panoptic segmentation network with cross-layer attention fusion," in *Proc. of ACCV*, 2020.
- [29] R. Hou *et al.*, "Real-time panoptic segmentation from dense detections," in *Proc. of CVPR*, 2020, pp. 8523–8532.
- [30] D. de Geus, P. Meletis, and G. Dubbelman, "Fast panoptic segmentation network," *RAL*, vol. 5, no. 2, pp. 1742–1749, 2020.
- [31] 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.
- [32] 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.
- [33] Y. Zhang, J. Lu, and J. Zhou, "Objects are different: Flexible monocular 3d object detection," in *Proc. of CVPR*, 2021, pp. 3289–3298.
- [34] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2d pose estimation using part affinity fields," in *Proc. of CVPR*, 2017, pp. 7291–7299.
- [35] L. Beyer, A. Hermans, and B. Leibe, "Biternion nets: Continuous head pose regression from discrete training labels," in *Proc. of GCPR*. Springer, 2015, pp. 157–168.
- [36] B. Lewandowski, D. Seichter, T. Wengefeld, L. Pfennig, H. Drumm, and H.-M. Gross, "Deep orientation: Fast and robust upper body orientation estimation for mobile robotic applications," in *Proc. of IROS*, 2019, pp. 441–448.
- [37] T. Wengefeld, B. Lewandowski, D. Seichter, L. Pfennig, and H.-M. Gross, "Real-time person orientation estimation using colored point-clouds," in *Proc. of ECMR*, 2019.
- [38] D. Seichter, B. Lewandowski, D. Höchmer, 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. 10 497–10 504.
- [39] O. Russakovsky *et al.*, "ImageNet Large Scale Visual Recognition Challenge," in *IJCV*, 2015, pp. 211–252.
- [40] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. of CVPR*, 2018, pp. 7132–7141.
- [41] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proc. of CVPR*, 2018, pp. 4510–4520.
- [42] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *Proc. of ICML*, 2019, pp. 6105–6114.
- [43] T. Standley, A. Zamir, D. Chen, L. Guibas, J. Malik, and S. Savarese, "Which tasks should be learned together in multi-task learning?" in *Proc. of ICML*. PMLR, 2020, pp. 9120–9132.
- [44] S. Vandenhende, S. Georgoulis, W. Van Gansbeke, M. Proesmans, D. Dai, and L. Van Gool, "Multi-task learning for dense prediction tasks: A survey," *TPAMI*, 2021.
- [45] S. Vandenhende, S. Georgoulis, and L. V. Gool, "Mti-net: Multi-scale task interaction networks for multi-task learning," in *Proc. of ECCV*. Springer, 2020, pp. 527–543.
- [46] Z. Zhang, Z. Cui, C. Xu, Z. Jie, X. Li, and J. Yang, "Joint task-recursive learning for semantic segmentation and depth estimation," in *Proc. of ECCV*, 2018, pp. 235–251.
- [47] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proc. of ICCV*, 2015, pp. 1026–1034.
- [48] P. Goyal *et al.*, "Accurate, large minibatch sgd: Training imagenet in 1 hour," *arXiv preprint arXiv:1706.02677*, 2017.
- [49] M. Roberts *et al.*, "Hypersim: A Photorealistic Synthetic Dataset for Holistic Indoor Scene Understanding," in *Proc. of ICCV*, 2021.
- [50] A. Paszke *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Proc. of NeurIPS*. Curran Associates, Inc., 2019, pp. 8024–8035.
- [51] T.-Y. Lin *et al.*, "Microsoft coco: Common objects in context," in *Proc. of ECCV*, 2014.
- [52] A. Kendall, Y. Gal, and R. Cipolla, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *Proc. of CVPR*, 2018, pp. 7482–7491.
- [53] Z. Chen, V. Badrinarayanan, C.-Y. Lee, and A. Rabinovich, "Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks," in *Proc. of ICML*. PMLR, 2018, pp. 794–803.
- [54] S. Liu, E. Johns, and A. J. Davison, "End-to-end multi-task learning with attention," in *Proc. of CVPR*, 2019, pp. 1871–1880.
- [55] V. R. Kumar *et al.*, "Omnidet: Surround view cameras based multi-task visual perception network for autonomous driving," *RAL*, vol. 6, no. 2, pp. 2830–2837, 2021.