

PanoSAMic: Panoramic Image Segmentation from SAM Feature Encoding and Dual View Fusion

Mahdi Chamseddine^{1,2}[0000-0003-4119-457X], Didier Stricker^{1,2}[0000-0002-5708-6023], and Jason Rambach¹[0000-0001-8122-6789]

¹ German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany

² RPTU Kaiserslautern-Landau, Kaiserslautern, Germany

firstname.lastname@dfki.de

Abstract. Existing image foundation models are not optimized for spherical images having been trained primarily on perspective images. PanoSAMic integrates the pre-trained Segment Anything (SAM) encoder to make use of its extensive training and integrate it into a semantic segmentation model for panoramic images using multiple modalities. We modify the SAM encoder to output multi-stage features and introduce a novel spatio-modal fusion module that allows the model to select the relevant modalities and best features from each modality for different areas of the input. Furthermore, our semantic decoder uses spherical attention and dual view fusion to overcome the distortions and edge discontinuity often associated with panoramic images. PanoSAMic achieves state-of-the-art (SotA) results on Stanford2D3DS for RGB, RGB-D, and RGB-D-N modalities and on Matterport3D for RGB and RGB-D modalities. <https://github.com/dfki-av/PanoSAMic>.

Keywords: Spherical images · Semantic segmentation · Modality fusion.

1 Introduction

Spherical images offer a new approach to sensing the environment. A single compact sensor captures a full 360° view of a scene and enables holistic understanding of the environment without the need for calibrating or aligning multiple input sources. Recent hardware developments have also gone beyond capturing only the spherical RGB data and started integrating depth information [1,4].

The unique characteristics of the spherical RGB and RGB-D (RGB + Depth) sensors have sparked interest in using them in many application areas such as robotics [5], extended and augmented vision [31], autonomous driving [17], and construction [12] among other fields.

Recent advances in deep learning algorithms and availability of massive data amounts have made it possible to train foundation models. These are models trained on vast datasets to tackle a wide range of tasks such as image and video segmentation [13] or depth estimation [32].

Even though several panoramic datasets [1,4,11] have been published to motivate further research, existing foundation models tend to under-perform

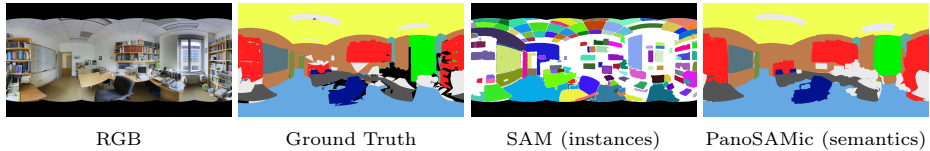


Fig. 1: SAM [13] is not trained for semantic segmentation and is unable to fully handle panoramic images (white is unsegmented). PanoSAMic uses the SAM pre-trained encoder and is tailored for semantic segmentation of panoramic images.

when used for spherical image processing as seen in Figure 1. Such discrepancy in performance between perspective and panoramic images is caused by the imbalance in the data used in training such foundation models.

Semantic segmentation is essential for scene understanding in various applications through dense pixel-level classification. Panoramic segmentation enables comprehensive scene understanding from a single frame. Existing work on panoramic image segmentation tried to solve the distortion associated with panoramic images with image projections [8,16], positional encoding [15], and deformable embeddings [9].

In this work, we use the encoder of the Segment Anything Model (SAM) [13], a pioneering foundation model for image segmentation, into a panoramic image segmentation model. We integrate the pre-trained encoder, benefiting from huge resources and data used in training it, and introduce a new panoramic decoder that can handle the spherical nature of the input. Additionally, we introduced a novel fusion and refinement module to fuse multiple input modalities: RGB, Depth, and Normals.

Our main contributions can be summarized by:

- Integrating the pre-trained SAM encoder into a novel panoramic image segmentation model.
- Introducing the dual-view fusion for handling the spherical nature and object separation on image borders of panoramic images.
- Developing the Moving Convolutional Block Attention Module (MCBAM) for spatio-modal fusion.
- Achieving State-of-the-Art results on the panoramic Stanford2D3DS and Matterport3D public datasets.

The rest of the paper is structured as follows: Section 2 presents an overview of the recent related works. Section 3 goes in depth explaining our contribution and implementation details. We then present our experiments and results in Section 4. Section 4.6 validates our contributions through various ablations. Finally, we conclude with our final remarks in Section 5.

2 Related Work

2.1 Panoramic Image Segmentation

Existing literature on panoramic and multimodal image segmentation explores techniques to handle distortions in panoramic images and the fusion of multiple modalities for improved segmentation performance.

Early methods adapted traditional perspective-based models for wide-field-of-view images using spherical polyhedrons or tangent image representations [8,16]. More recent techniques fall into distortion-aware and 2D geometry-aware approaches. Distortion-aware methods employ specialized convolutions, such as spherical convolutions [10], distortion-aware modules [24], and adaptive kernel fusion [42]. Transformer-based architectures, such as Trans4PASS+ [37], integrate Deformable Patch Embedding and Deformable Multi-Layer Perceptron for enhanced panoramic segmentation. SGAT4PASS [15] introduces a spherical geometry-aware transformer bridging the gap between 2D panoramic segmentation and 3D-aware scene perception.

Fusion-based segmentation techniques leverage multiple data sources such as RGB, Depth, and other sensor modalities. Previous research on RGB-D segmentation developed new layers to capture geometric properties [2] or architectures for multi-modal fusion [22,23]. SFSS-MMSI [9] combines the deformable characteristics of Trans4PASS+ with Cross Modal Fusion (CMX) [36] for panoramic semantic segmentation. 360BEV [26] on the other hand introduces transformer-based 360° panorama to bird’s eye view (BEV) semantic map.

2.2 Segment Anything Model

The Segment Anything Model (SAM) [13], is a foundation model for image segmentation trained on a large image dataset and designed to segment any object in an image. After its introduction, SAM has been rapidly adapted to diverse domains such as medical image segmentation [18], 3D segmentation [33], and quality monitoring and recycling [41]. Several approaches have extended it to semantic segmentation by combining its masks with a classifier [14] or injecting task-specific features through adapters [34], or coupling SAM with a domain-specific encoder [38]. Unlike existing work, we do not introduce additional encoders or adapters. Instead, we reuse the frozen encoder and extract and refine intermediate features.

2.3 Open-Vocabulary Segmentation

Recent methods have explored extending dense prediction to arbitrary categories. CAT-Seg [6] uses CLIP [19] to associate pixel-level features with different classes, and OpenSeeD [35] addresses panoptic segmentation in an open-vocabulary setting. Open Panoramic Segmentation (OOOPS) [39] adapts this idea specifically to panoramic images and SAM3 [3] delivers open-vocabulary segmentation for images and videos. While these approaches highlight the promise of open-vocabulary models, supervised methods still achieve stronger results on panoramic images.

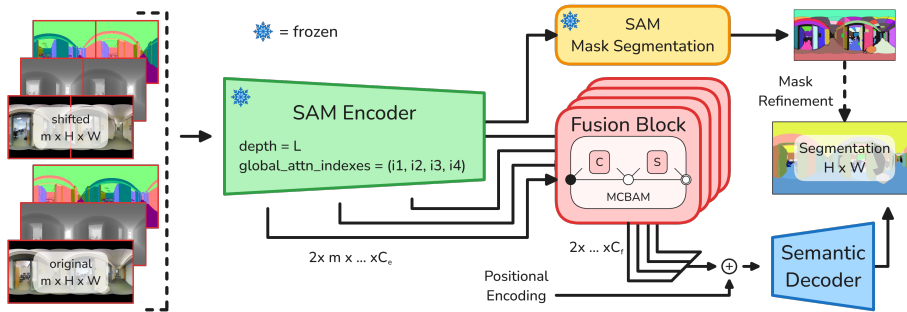


Fig. 2: PanoSAMic architecture. Two views of the same panoramic input are fed into the SAM encoder. The fusion block combines and refines the features from the processed input modalities. The features are then concatenated and passed into the decoder along with a horizontal positional encoding. The semantic decoder outputs the fused segmentation which is refined with mask prediction.

3 Methodology

The PanoSAMic architecture, seen in Figure 2, is comprised of the SAM [13] model with frozen weights and modified encoder, the feature fusion module with a feature fusion block per branch output, and the semantic decoder with dual-view fusion. The network processes two views of the same panoramic scene in parallel before fusing the segmentation outputs for the final segmentation output.

3.1 Input Modalities and Views

To achieve better scene understanding, we leverage different modalities for their respective contributions to the segmentation task. Whereas RGB provides color and context, normal images are ideal for detecting continuous surfaces as well as edges, and depth images help mitigate the ambiguity caused by perspective distortions by introducing scale information.

While equirectangular images represent data from a 360° scene, they introduce a discontinuity at the scene edges as shown in Figure 3 and thus reducing the segmentation quality. Existing approaches either ignore this limitation, or attempt to solve this problem through different projections [8,16] and complex encoding [15]. Such approaches require a complicated or multi-step process for merging the projections. To address those limitations and to exploit the pre-trained SAM encoder without the need for fine-tuning, we introduce the concept of *dual-view fusion* (shifted panoramas).

PanoSAMic processes the shifted input in parallel as a new batch and the feature maps are combined in a late-fusion step after reversing the shift to produce the combined segmentation output. Formally, we describe the shifting as $X_R(i, j) = X(i, (j + s) \bmod W)$, $\forall i \in H, j \in W$, where X and X_R are the original and shifted images respectively, (i, j) the pixel coordinates, H and W the height and width of the image, and s is the shift amount fixed to $W/2$. The shift ensures that an object is whole in at least one of the two views.

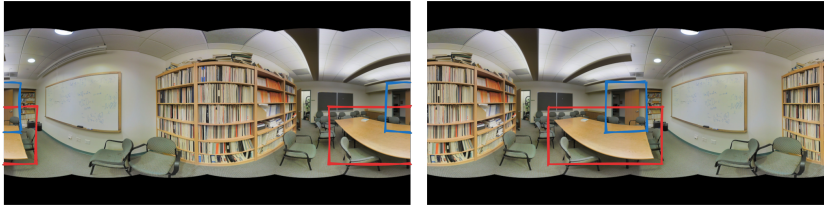


Fig. 3: Objects in panoramic images that are disconnected on the edges are processed as whole in the shifted view.

Horizontal Positional Encoding (HPE) To improve the *dual-view fusion*, we used a 1-D positional encoding. This positional encoding is added to the encoded inputs with shifted encoding added to the shifted input. Since a horizontal shift of the equirectangular images in reality represents a rotational shift, the positional encoding allows the model to align the features for the fused segmentation.

We used the positional encoding as described in [27] and defined as follows,

$$\begin{aligned}
 PE_{(w,2i)} &= \sin(w/10000^{2i/C}), \\
 PE_{(w,2i+1)} &= \cos(w/10000^{2i/C}), \\
 PE_R &= PE_{(w+W/2 \bmod W, \dots)},
 \end{aligned} \tag{1}$$

where w is the index of a column in the feature map, i is the index of the channel, C the total number of channels of the encoded features, and W the total width of the encoded features. PE_R is the shifted version of PE produced by rolling the values of PE along the width dimension. PE and PE_R are added to the fused features of both input views.

3.2 Encoder Modifications

The Segment Anything model (SAM) [13] uses a plain backbone based on the vision transformer (ViT) [7] architecture. The SAM encoder is a depth L ViT with window attention at every layer except for specific layers that use global attention instead. In SAM, the number of global attention layers was chosen as 4. The transformer blocks are followed by a convolutional block that outputs the final encoded features.

Kirillov *et al.* [13] used three different encoder sizes with an increasing number of transformer blocks and subsequently different positions of the global attention.

We extended the encoder to add a branch output after each global attention layer to allow the fusion of the features of different encoded modalities at multiple levels of encoding. This modification does not add any parameters to the encoder thus allowing for the use of the pre-trained weights.

3.3 Feature Fusion

The encoder processes the modalities independently as different batches and thus there is no interaction between their features. To add inter-modality feature fusion, we introduce a Fusion Block for each of the encoder branches.

3.4 Convolutional Block Attention Module

The Convolutional Block Attention Module (CBAM), proposed by Woo *et al.* [28], enhances CNN feature representation through sequential channel and spatial attention. The channel attention refines feature importance using global pooling and MLPs, while spatial attention captures spatial dependencies via pooling and convolution. CBAM is lightweight, easily integrated into CNNs, and showed performance improvements on classification and detection tasks.

Moving Convolutional Block Attention Module In its original design, CBAM applies global pooling before attention, which works well for classification but is less suited for semantic segmentation, where different image regions often correspond to different objects.

To overcome this limitation, we introduce Moving CBAM (MCBAM), which applies a sliding-window channel attention block followed by a sliding-window spatial attention block to each of the branch outputs of the encoder. Each local region of the feature map is refined separately, allowing tailored attention in different parts of the image. When windows overlap, the channel attention values are aggregated using per channel max pooling, while spatial attention values are summed up and passed through a sigmoid activation. This overlap-handling ensures a coherent feature map across window boundaries. Figure 4a shows a graphical representation of MCBAM.

The refined features are added to the input features through a feed-forward connection similar to [28] and then passed through a convolutional block and upscaling block.

Inter-Modality Fusion By applying MCBAM to the concatenated feature output of the encoder for all modalities, we allow the channel attention component to *select* the best features for each window location and from the most relevant modality for that window. Similarly, the channel attention highlights the most significant areas of each window.

After inter-modality fusion, the different branch outputs are concatenated and then the HPE is added to the concatenated features before being fed into the semantic decoder.

3.5 Semantic Decoder

For semantic segmentation, we adapt the lightweight decoder architecture of Xie *et al.* [30], chosen for its efficiency and reliability with transformer-based (ViT [7])

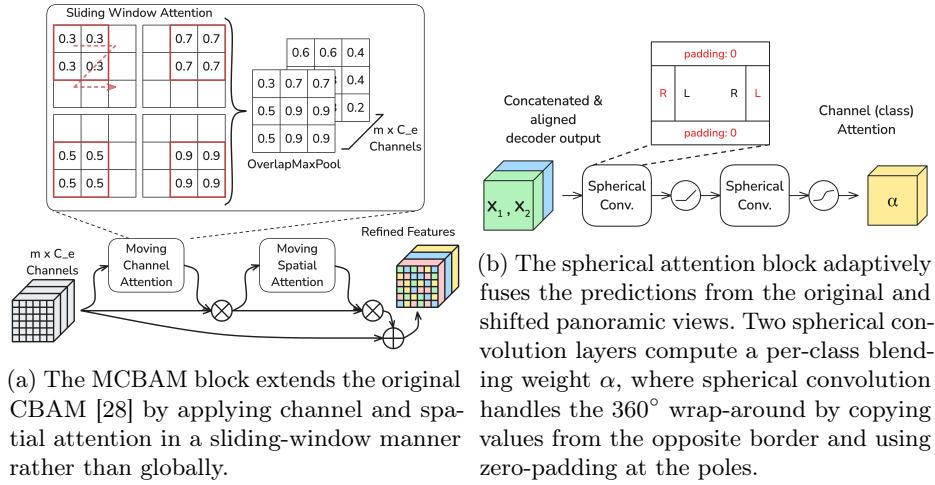


Fig. 4: Our novel blocks used for feature fusion and dual view fusion.

backbones. We extend it to handle the specific challenges of panoramic images. Both the original and shifted panoramas are decoded separately; the shifted view is then realigned to the original coordinate system. To merge the two predictions, we introduce a per-class blending mechanism computed with a novel spherical attention block: $x_{pred} = \alpha \cdot x_1 + (1 - \alpha) \cdot x_2$, where x_1 and x_2 are the aligned outputs of both views, and α is a class-wise attention weight in $[0, 1]$.

The spherical attention block, shown in Figure 4b, consists of two spherical convolutional layers separated by a non-linearity and followed by a sigmoid activation. Unlike standard convolution, spherical convolution handles the wrap-around of equirectangular panoramas by padding the left border with values from the right border and vice versa, while using zero-padding at the top and bottom. This ensures that features near the left-right boundary interact consistently across the full 360°.

3.6 Instance-Guided Semantic Refinement

We extended SAM’s instance segmentation to handle multi-modal panoramic inputs through three key modifications: (1) **Multi-modal fusion**: instance masks are generated independently from each modality and merged via greedy mask-based NMS (non-maximum suppression) to preserve complementary boundaries. (2) **Dual-view processing**: each modality is processed twice, with quality-aware NMS selecting higher-quality masks based on SAM’s predicted IoU scores. (3) **Semantic refinement**: instance masks refine semantic predictions via majority voting, each instance region is assigned its most frequent semantic class, while background pixels retain original predictions.

4 Experiments and Results

To evaluate the PanoSAMic architecture, we performed experiments with different modality inputs and compared to state-of-the-art methods, reporting both quantitative and qualitative results. In addition, we performed ablation experiments to evaluate the impact of specific design choices and proposed components.

4.1 Dataset

We used the Stanford2D3DS [1] and Matterport3D [4] for evaluating the model.

The **Stanford2D3DS** dataset contains 1413 samples with RGB, depth, and normal data as well as instance and semantic labels. The data is spread over six areas and provides 13 class labels and images at a 2048×4096 resolution. We used the 3-Fold cross-validation splits suggested by the authors and evaluated our model using mean intersection over union (mIoU) and mean accuracy (mAcc).

As for the **Matterport3D** dataset, we used the pre-processed data and splits provided by Teng *et al.* in BEV360 [26] in order to perform a fair comparison. The dataset contains 10615 samples with RGB and depth and a subset of 20 class labels down from 40 in the original dataset [4] to reduce the class imbalance.

4.2 Experiment Setup

For training the model, we used frozen SAM ViT-H backbone that has an encoder depth of 32 blocks, with blocks [8, 16, 24, 32] using global attention. We used a batch size of 8 and trained for 50 epochs, utilizing the Ranger21 optimizer [29] with a maximum learning rate of 0.0005 for Stanford2D3DS and 0.001 for Matterport3D. Image augmentation was done using random horizontal flipping, random horizontal rolling, and color permutation of the RGB input.

Similar to other approaches [8, 26, 21, 9], we resized the input to 512×1024 . We evaluated different modality inputs to the network: **RGB**, **RGB + Depth**, and **RGB + Depth + Normals**. When only evaluating the **RGB** modality, and following the procedure done in other methods, we masked the black area from the metrics computation due to the lack of data in those areas. For our MCBAM (Section 3.4) block we used a window size of 8×8 and stride of 4 and for our *SphericalAttention* (Section 3.5) we used a kernel size of 7×7 and a stride 1.

For the loss, we used the Jaccard loss [20] for training on the Stanford2D3DS dataset and the alternating scheduled loss by Taubert *et al.* [25] (Cross-Entropy and Jaccard losses) for training on the Matterport3D dataset.

Since the SAM encoder was trained on images, RGB and normal maps were directly used as input to the encoder. The depth map was first preprocessed to create a pseudo-disparity image by cropping, and scaling:

$$D(i, j) = 1 - \frac{\min(D(i, j), d_t)}{d_t}, \quad (2)$$

where D is the depth map, (i, j) are the pixel coordinates, and d_t is the depth threshold representing 99.5% of the depth values in the training set rounded to the nearest 10 *cm*. We finally replicate D to create a 3-channel encoder input.

Table 1: Semantic segmentation results on the Stanford2D3DS [1] dataset using different input modalities and configurations (open vocabulary vs. supervised).

Method	Configuration / Modalities	3-fold Validation mIoU %	mAcc %	# Params (millions)
CAT-seg [6]	Open-vocabulary with RGB	39.60	–	59.5
OpenSeeD [35]		40.00	–	65.4
OOOPS [39]		42.60	–	8.7
SAM3 [3]		52.05	65.02	850
SFSS-MMSI [9]	RGB	52.87	63.96	40
HoHoNet [23]		51.99	62.97	70
PanoFormer [22]		52.35	64.31	20
CBFC [40]		52.20	65.60	–
Tangent [8]		45.60	65.20	–
Trans4PASS+ [37]		52.04	63.98	39
MultiPanoWise [21]		54.60	–	–
PanoSAMic (ours)		59.62	74.11	178
SFSS-MMSI [9]		RGB-D	55.49	68.57
HoHoNet [23]	56.73		68.23	70
PanoFormer [22]	57.03		68.08	20
CBFC [40]	56.70		70.80	–
Tangent [8]	52.50		70.10	–
360BEV [26]	54.30		–	27.7
PanoSAMic (ours)	60.90		73.95	184
SFSS-MMSI [9]	RGB-D-N	59.43	69.03	123
PanoSAMic (ours)		61.57	74.04	191

4.3 Quantitative Results

We compared the results of PanoSAMic in indoor semantic segmentation tasks and using different modality inputs with existing methods in Tables 1 and 2.

Stanford2D3DS For the CBFC [40], Tangent [8], 360BEV [26], and MultiPanoWise [21], we used the results reported by their respective authors in the original publications. For the rest of the methods, we used the results presented and reproduced by SFSS-MMSI [9]. As for the open vocabulary methods, we evaluated the pre-trained SAM3 [3] on all classes then merged and refined the predictions, while for the rest of the methods we report the values from OOOPS [39].

Table 1 shows that we achieve state-of-the-art results across all input modalities (RGB, RGB-D, and RGB-D-N) on the Stanford2D3DS dataset. Our method consistently outperforms prior supervised approaches, like MultiPanoWise [21] and SFSS-MMSI [9], with strong gains in mean accuracy across all modalities. We also note that supervised methods surpass most open-vocabulary approaches on RGB-based semantic segmentation.

While PanoSAMic has more trainable parameters than some prior methods, the overhead is modest compared to the SAM backbone itself. Since the encoder is frozen, training remains efficient, and the extra capacity comes mainly from lightweight fusion and decoding modules. Additionally, our model size increases only marginally when moving from RGB to RGB-D and RGB-D-N (+6M parameters), whereas SFSS-MMSI [9] grows by +41M between modalities.

Table 2: Semantic segmentation results on the Matterport3D [4] dataset using different input modalities.

Method	Configuration / Modalities	mIoU %
CAT-seg [6]	Open-vocabulary with RGB	31.10
OpenSeeD [35]		31.60
OOOPS [39]		32.50
SAM3 [3]		42.29
Trans4PASS+ [37]	RGB	42.60
HoHoNet [23]		44.10
SegFormer [30]		45.53
PanoSAMic (ours)		46.59
360BEV [26]		46.35
PanoSAMic (ours)	RGB-D	48.43

Matterport3D We use all results as reported in the 360BEV [26] experiments. For the open-vocabulary methods, we report results from OOOPS [39] and ensure fair comparison by using the same pre-processed data as 360BEV and OOOPS. SAM3 [3] results are evaluated similar to Stanford2D3DS.

On Matterport3D, Table 2 shows overall results are lower than on Stanford2D3DS, but PanoSAMic still outperforms prior methods for RGB-only segmentation and achieves state-of-the-art performance on RGB-D, surpassing 360BEV [26]. As with Stanford2D3DS, supervised methods clearly outperform open-vocabulary approaches.

4.4 Qualitative Results

In addition to the quantitative results, we visualize our segmentation results on the Stanford2D3DS [1] dataset for a qualitative analysis. Figure 5 compares the different configurations of our PanoSAMic model on different scenes.

We observe that with **RGB** input the segmentation shows better quality results on near objects than further away which appear smaller in the image. Furthermore, we see that with **RGB-D** input, more objects are segmented with finer detail, however, we notice that there was a tendency to mistake walls or doors with bookcases in some cases (**Scene 1** and **Scene 1**). Finally, and as confirmed by the quantitative results in Table 1, we see that the **RGB-D-N** configuration results in the highest fidelity segmentation for both near and far objects in the scenes.

Overall, the high quality of the segmentation for all tested modalities aligns with the SotA results presented in Table 1.

4.5 Ground Truth Analysis and Generalization

While analyzing the results of our evaluation, we observed more closely some of the worse performing samples (low mIoU) on the validation set of the different folds. This resulted in finding some samples in the dataset that had various levels of ground truth imprecision. Figure 6 shows different examples of these imprecisions as well as their respective segmentation results.

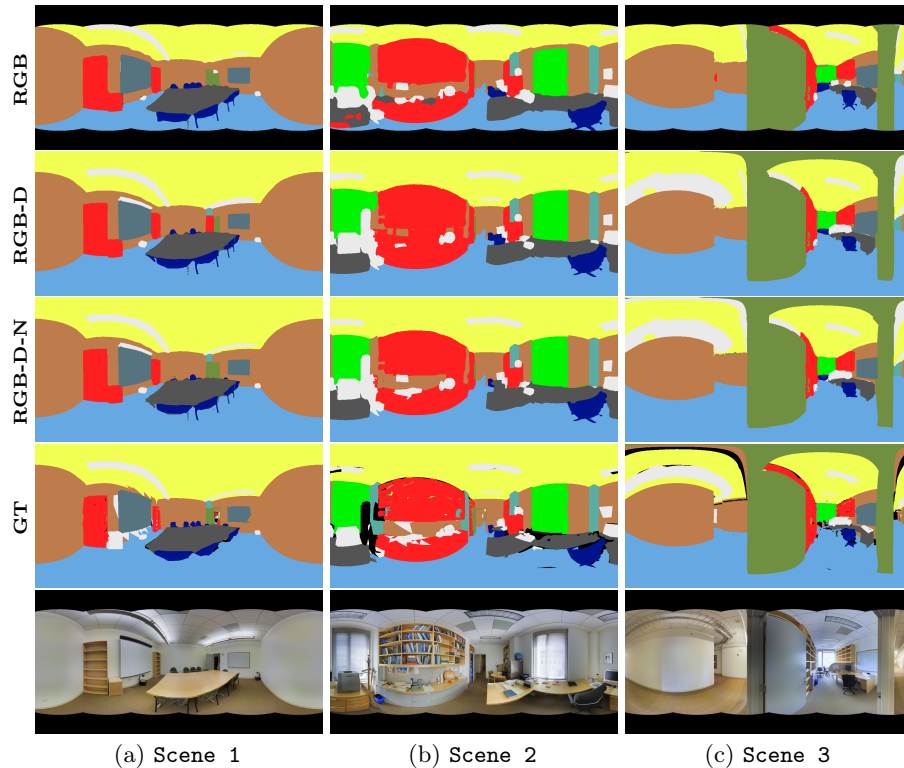


Fig. 5: Comparison of the qualitative segmentation results of our PanoSAMic model for different scenes from the Stanford2D3DS dataset [1] and using different input configurations.

Such imprecisions can be caused by, quantization in mesh reprojection (Stanford2D3DS labels are projected from 3D to 2D), mislabeled instances, or missing annotations as can be seen in the outdoor scene in Figure 6c.

Comparing the images in Figure 6 to the segmentation output shows that PanoSAMic segmented scenes well and often produced better labels than the provided ground truth. Given that Stanford2D3DS is primarily an indoor dataset, the segmentation result in Figure 6c clearly demonstrates the generalization capability of the model. The model, however, mislabeled the sky, which is expected given the absence of this class in training and the visual similarity between an overcast sky and a white wall.

4.6 Ablation Studies

We validate our contributions through ablations on different model configurations. We follow the setup in Section 4.2 using **RGB-D-N** input on the standard Stanford2D3DS 1-Fold (**Area 5**) validation. As baseline, we use the vanilla SAM encoder with a convolutional decoder with a **single view** input.

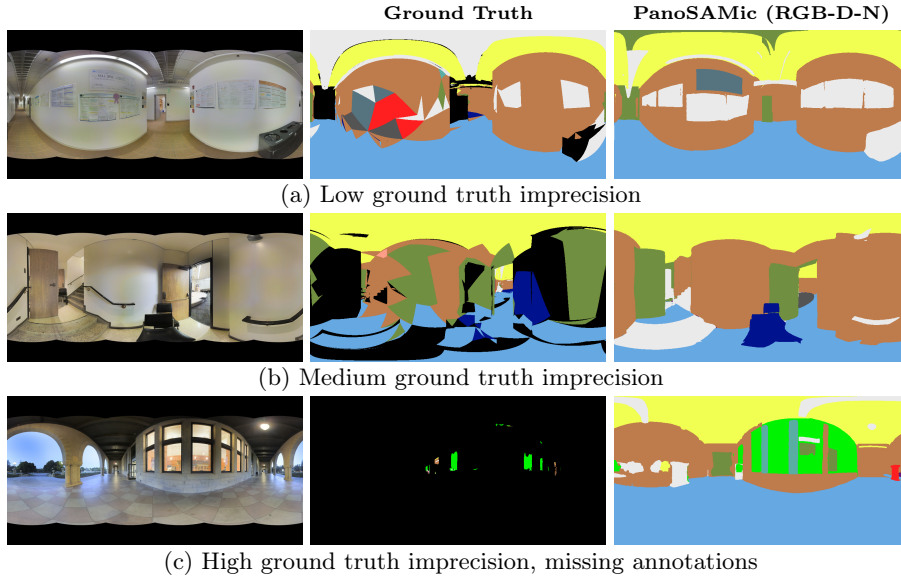


Fig. 6: Some examples of scenes with imprecise ground truth labels from Stanford2D3DS [1]. Our PanoSAMic model performs very well on the different scenes and locations showing high generalization.

The different configurations evaluated in Table 3 explore alternative feature refinement strategies within a **single-view** segmentation setting. Introducing the encoder modification alone already yields a substantial improvement over the baseline, resulting in a large gain in mIoU and placing the model firmly within state-of-the-art performance for Stanford2D3DS [1].

As shown in Table 3, adding encoder branches without attention (B) yields a substantial improvement over the baseline configuration (A). Introducing plain Channel Attention on top of this setup (C) leads to a slight performance drop compared to no attention, indicating that channel-wise reweighting alone is not sufficient for dense semantic segmentation. Similarly, applying the standard CBAM module (D) does not consistently improve performance, suggesting limitations of its global spatial attention when operating on fine-grained, densely labeled scenes. In contrast, the proposed Moving CBAM (MCBAM) (E) recovers and

Table 3: Ablations of different model configurations and their effect on the segmentation results of Stanford2D3DS [1].

Model Configuration	mIoU %
A: (baseline) SAM Encoder + Conv. Decoder	53.29
B: A + Enc. Branches + <i>no attention</i> (ours)	61.20
C: B + Channel Attention (ours)	60.82
D: B + CBAM (ours)	61.11
E: B + MCBAM (ours)	61.61
F: E + Instance Refinement (ours)	62.43

Table 4: Comparison of the segmentation of the edges of input **RGB** between single-view and dual-view configurations.

Edge Ratio	Single View		Dual View	
	mIoU %	mAcc %	mIoU %	mAcc %
-	57.95	71.79	59.62	74.11
0.5	57.79↓	72.14	59.86↑	74.83
0.3	57.51↓	72.17	60.23↑	75.56
0.1	55.52↓	70.54	60.39↑	76.21

further improves the mIoU by enabling localized, region-aware feature refinement. Finally, incorporating the instance refinement strategy (F) yields the best overall performance, demonstrating its complementary benefit on top of MCBAM.

To further assess the impact of dual-view fusion, we evaluated segmentation performance specifically on edge regions of the images using the 3-Fold **RGB** validation of Stanford2D3DS. Table 4 reports results for different edge ratios, defined as the fraction of pixels near the left and right borders (*e.g.* an edge ratio of 0.5 corresponds to $0.25W$ on each side, where W is the image width).

For single-view segmentation, performance decreases as the edge ratio narrows. In contrast, dual-view fusion shows the opposite trend. The mIoU difference between single and dual view grows from 1.67% for full images to 4.87% at the edges, confirming that dual-view fusion effectively mitigates boundary discontinuities.

5 Conclusion

In this work, we presented *PanoSAMic*, our multi-modal panoramic segmentation model using the SAM feature encoder and leveraging its pre-trained weights on large amounts of data. We extended the existing SAM architecture with our dual-view fusion to handle edge discontinuity of spherical images and introduced an improved self attention block (MCBAM) for multi-modal fusion and segmentation.

We evaluated our model on public data and achieved State-of-the-Art results by a large margin. We tested our model with different modalities including depth and normals and show that we achieve SotA with all input combinations. We further validated the importance of our architecture contributions through multiple ablations and proved that our Moving CBAM refines features for semantic segmentation tasks unlike the original CBAM module tailored for classification. We also showed that our dual view fusion successfully addresses the edge discontinuity of panoramic images. Our results demonstrate strong generalization capabilities and show that SAM can be adapted for semantic segmentation.

Acknowledgements

This research was funded by the European Union as part of the projects: HumanTech (Grant Agreement 101058236) and ShieldBOT (Grant Agreement 101235093).

References

1. Armeni, I., Sax, S., Zamir, A.R., Savarese, S.: Joint 2d-3d-semantic data for indoor scene understanding. *arXiv:1702.01105* (2017)
2. Cao, J., Leng, H., Lischinski, D., Cohen-Or, D., Tu, C., Li, Y.: Shapeconv: Shape-aware convolutional layer for indoor rgb-d semantic segmentation. In: *ICCV* (2021)
3. Carion, N., Gustafson, L., Hu, Y.T., Debnath, S., Hu, R., Suris, D., Ryali, C., et al.: Sam 3: Segment anything with concepts. *arXiv:2511.16719* (2025)
4. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., et al.: Matterport3d: Learning from rgb-d data in indoor environments. *arXiv:1709.06158* (2017)
5. Chaplot, D.S., Salakhutdinov, R., Gupta, A., Gupta, S.: Neural topological slam for visual navigation. In: *CVPR* (2020)
6. Cho, S., Shin, H., Hong, S., Arnab, A., Seo, P.H., Kim, S.: Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. In: *CVPR* (2024)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR* (2021)
8. Eder, M., Shvets, M., Lim, J., Frahm, J.M.: Tangent images for mitigating spherical distortion. In: *CVPR* (2020)
9. Guttikonda, S., Rambach, J.: Single frame semantic segmentation using multi-modal spherical images. In: *WACV* (2024)
10. Jiang, C.M., Huang, J., Kashinath, K., Prabhat, Marcus, P., Niessner, M.: Spherical CNNs on unstructured grids. In: *ICLR* (2019)
11. Kanayama, H., Chamseddine, M., Guttikonda, S., Okumura, S., Yokota, S., Stricker, D., Rambach, J.: Tof-360-a panoramic time-of-flight rgb-d dataset for single capture indoor semantic 3d reconstruction. In: *CVPRW* (2025)
12. Kaufmann, F., Chamseddine, M., Guttikonda, S., Glock, C., Stricker, D., Rambach, J.: Ontology-based semantic labeling for rgb-d and point cloud datasets. In: *EC3*. vol. 4. European Council on Computing in Construction (2023)
13. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., et al.: Segment anything. In: *ICCV* (2023)
14. Kweon, H., Yoon, K.J.: From sam to cams: Exploring segment anything model for weakly supervised semantic segmentation. In: *CVPR* (2024)
15. Li, X., Wu, T., Qi, Z., Wang, G., Shan, Y., Li, X.: Sgat4pass: spherical geometry-aware transformer for panoramic semantic segmentation. In: *IJCAI* (2023)
16. Li, Y., Guo, Y., Yan, Z., Huang, X., Duan, Y., Ren, L.: Omnifusion: 360 monocular depth estimation via geometry-aware fusion. In: *CVPR* (2022)
17. Ma, C., Zhang, J., Yang, K., Roitberg, A., Stiefelhagen, R.: Densepass: Dense panoramic semantic segmentation via unsupervised domain adaptation with attention-augmented context exchange. In: *ITSC*. IEEE (2021)
18. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1) (2024)
19. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., et al.: Learning transferable visual models from natural language supervision. In: *ICML*. PmLR (2021)
20. Rahman, M.A., Wang, Y.: Optimizing intersection-over-union in deep neural networks for image segmentation. In: *ISVC*. Springer (2016)
21. Shah, U., Tukur, M., Alzubaidi, M., Pintore, G., Gobbetti, E., Househ, M., Schneider, J., et al.: Multipanowise: holistic deep architecture for multi-task dense prediction from a single panoramic image. In: *CVPR* (2024)

22. Shen, Z., Lin, C., Liao, K., Nie, L., Zheng, Z., Zhao, Y.: Panoformer: panorama transformer for indoor 360° depth estimation. In: ECCV. Springer (2022)
23. Sun, C., Sun, M., Chen, H.T.: Hohonet: 360 indoor holistic understanding with latent horizontal features. In: CVPR (2021)
24. Tateno, K., Navab, N., Tombari, F.: Distortion-aware convolutional filters for dense prediction in panoramic images. In: ECCV (2018)
25. Taubert, O., Götz, M., Schug, A., Streit, A.: Loss scheduling for class-imbalanced image segmentation problems. In: ICMLA. IEEE (2020)
26. Teng, Z., Zhang, J., Yang, K., Peng, K., Shi, H., Reiß, S., Cao, K., et al.: 360bev: Panoramic semantic mapping for indoor bird’s-eye view. In: WACV (2024)
27. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., et al.: Attention is all you need. *NeurIPS* **30** (2017)
28. Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: ECCV (2018)
29. Wright, L., Demeure, N.: Ranger21: a synergistic deep learning optimizer. arXiv:2106.13731 (2021)
30. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. *NeurIPS* **34** (2021)
31. Xu, Y., Zhang, Z., Gao, S.: Spherical dnns and their applications in 360° images and videos. *TPAMI* **44**(10) (2021)
32. Yang, L., Kang, B., Huang, Z., Xu, X., Feng, J., Zhao, H.: Depth anything: Unleashing the power of large-scale unlabeled data. In: CVPR (2024)
33. Yang, Y., Wu, X., He, T., Zhao, H., Liu, X.: Sam3d: Segment anything in 3d scenes. arXiv:2306.03908 (2023)
34. Yao, B., Deng, Y., Liu, Y., Chen, H., Li, Y., Yang, Z.: Sam-event-adapter: Adapting segment anything model for event-rgb semantic segmentation. In: ICRA. IEEE (2024)
35. Zhang, H., Li, F., Zou, X., Liu, S., Li, C., Yang, J., Zhang, L.: A simple framework for open-vocabulary segmentation and detection. In: ICCV (2023)
36. Zhang, J., Liu, H., Yang, K., Hu, X., Liu, R., Stiefelhagen, R.: Cmx: Cross-modal fusion for rgb-x semantic segmentation with transformers. *T-ITS* **24**(12) (2023)
37. Zhang, J., Yang, K., Shi, H., Reiß, S., Peng, K., Ma, C., Fu, H., et al.: Behind every domain there is a shift: Adapting distortion-aware vision transformers for panoramic semantic segmentation. *TPAMI* (2024)
38. Zhang, J., Ma, K., Kapse, S., Saltz, J., Vakalopoulou, M., Prasanna, P., Samaras, D.: Sam-path: A segment anything model for semantic segmentation in digital pathology. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer (2023)
39. Zheng, J., Liu, R., Chen, Y., Peng, K., Wu, C., Yang, K., Zhang, J., et al.: Open panoramic segmentation. In: ECCV. Springer (2024)
40. Zheng, Z., Lin, C., Nie, L., Liao, K., Shen, Z., Zhao, Y.: Complementary bi-directional feature compression for indoor 360deg semantic segmentation with self-distillation. In: WACV (2023)
41. Zhou, Y., Thielmann, P., Chamoli, A., Mirbach, B., Stricker, D., Rambach, J.: Particlesam: Small particle segmentation for material quality monitoring in recycling processes. arXiv:2508.03490 (2025)
42. Zhuang, C., Lu, Z., Wang, Y., Xiao, J., Wang, Y.: Acdnet: Adaptively combined dilated convolution for monocular panorama depth estimation. In: AAAI. vol. 36 (2022)

PanoSAMic: Panoramic Image Segmentation from SAM Feature Encoding and Dual View Fusion

Supplementary Material

Mahdi Chamseddine^{1,2}, Didier Stricker^{1,2}, and Jason Rambach¹

¹ German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany

² RPTU Kaiserslautern-Landau, Kaiserslautern, Germany

firstname.lastname@dfki.de

A Instance Guided Refinement

PanoSAMic extends SAM’s [13] instance segmentation capabilities to multi-modal panoramic scenes. The model produces two complementary outputs: (1) instance segmentation masks from SAM’s frozen components, and (2) semantic segmentation logits from trainable fusion and decoder modules. Instance generation uses SAM’s automatic mask generator with a 32×32 grid of point prompts, processed in batches of 64.

A.1 Multi-Modal Instance Generation

Modality-wise segmentation: Unlike SAM which processes only RGB, we generate instance masks from all available modalities (RGB, depth, normals). For each modality, we independently apply SAM’s prompt encoder and mask decoder, producing separate sets of instance proposals. This multi-modal approach captures complementary boundaries: RGB excels at texture edges, depth captures geometric discontinuities, and normals detect surface orientation changes.

Post-processing: Each modality’s predictions undergo SAM’s standard post-processing pipeline: stability score filtering, predicted IoU filtering, box-based NMS (non-maximum suppression), and small region removal.

Cross-modality fusion: After individual modality processing, we merge masks from all modalities using greedy mask-based NMS. Masks are sorted by quality score (predicted IoU), and lower-quality masks overlapping with higher-quality ones are removed. This preserves the best boundaries from each modality.

A.2 Dual-View Panoramic Fusion

Masks from the rotated view are *unshifted* back to original coordinates. When masks from both views overlap significantly, we select the higher-quality mask based on predicted IoU. If quality scores differ by a small margin, we use mask area as a tiebreaker, preferring larger masks. This ensures we keep the best representation of each object regardless of which view captured it better.

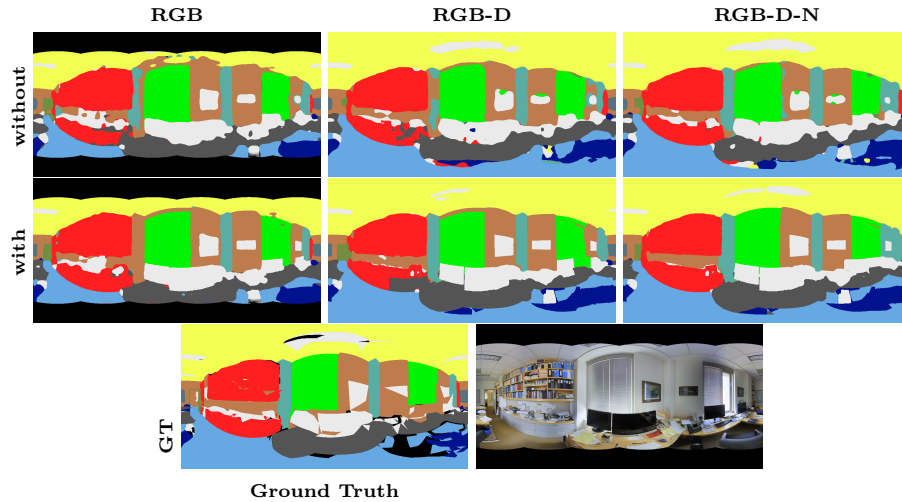


Fig. 7: Comparison of the qualitative segmentation results before and after refinement on different modality inputs.

A.3 Semantic Refinement

Given semantic logits $\mathbf{L} \in \mathbb{R}^{C \times H \times W}$ from the decoder and filtered instance masks, we refine predictions as follows: for each instance mask, we compute the most frequent semantic class within that instance region (using initial argmax predictions), then assign all pixels in that instance to this majority class. Background pixels (not covered by any instance) retain their original semantic predictions.

A.4 Evaluation

Figure 7 shows the effect of refinement on the prediction of the different modality inputs. The refinement step improves the overall quality of the segmentation, reduces “blobiness”, and enhances the edges. In rare cases, the refinement can have a negative effect by removing the correct class if it is not well segmented in the data: *e.g.* the column in the RGB example, or the clutter on the bookshelf in the RGB-D and RGB-D-N examples.

B SAM3 Evaluation

SAM3 [3] is a text-promptable segmentation model that generates instance masks conditioned on natural language class descriptions. We used the official SAM3 checkpoint and code and evaluation is performed on all three folds of Stanford2D3DS and the single validation split of Matterport3D and results reported in Tables 1 and 2.

B.1 Evaluation Procedure

For each test image, we perform per-class prompting by sequentially querying SAM3 with text prompts corresponding to each semantic class name (*e.g.*, “wall”, “floor”, “ceiling”). Each prompt generates a set of scored instance masks. We fuse these per-class predictions into a unified semantic segmentation map by constructing a class score tensor of shape $C \times H \times W$. For each class c at pixel (x, y) , we compute:

$$s_c(x, y) = \max_i m_i(x, y) \cdot \sigma_i, \quad (1)$$

where m_i is the mask logit from prediction i for class c , and σ_i is its confidence score. Only masks with confidence $\sigma_i \geq 0.25$ are retained. The final semantic label assigns each pixel to the class with the highest score.

B.2 Clutter Class Handling

Both datasets include a catch-all class for miscellaneous objects: “clutter” in Stanford2D3DS and “objects” in Matterport3D. For pixels with no coverage (all class scores are zero), we assign them to the clutter class. Additionally, pixels where the maximum class score falls below 0.05 are also assigned to clutter. This effectively creates a low-confidence region classifier.

B.3 Spatial Smoothing

To reduce speckle artifacts, we apply 3×3 majority filtering after obtaining the initial label predictions. For each pixel, we replace its label with the most frequent label in its connected neighborhood. This smoothing is applied after clutter assignment to ensure it operates on the final label space including the clutter class.

C Evaluation of Encoder Size

Table 5 shows the results of testing the different encoder sizes pretrained by SAM [13] on the Stanford2D3DS [1] dataset. While the ViT-L encoder performs slightly worse than ViT-H, the ViT-B encoder shows significant result degradation.

Table 5: Comparing the segmentation results (3-Fold) with respect to different encoder depths.

	Encoder Depth	mIoU %	mAcc %
ViT-B	12	56.68	70.49
ViT-L	24	60.90	73.09
ViT-H	32	61.57	74.04

Overall, our model still delivers competitive even state-of-the-art semantic segmentation results regardless of the SAM encoder used.

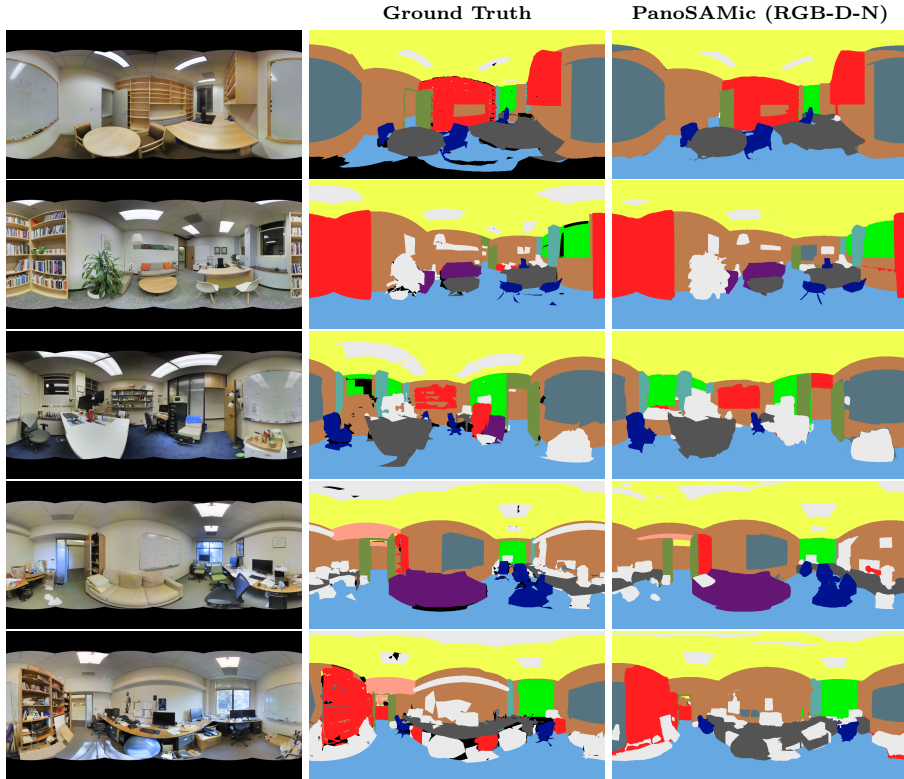


Fig. 8: More qualitative results on the Stanford2D3DS dataset [1].

D Model Parameters and Efficiency

Our model uses the frozen SAM encoder for its backbone. The number of parameters for SAM are as reported in their paper [13]: 91M, 308M, and 636M parameters for the ViT-B, ViT-L, and ViT-H backbones respectively. Our trainable parameters and model FLOPs shown in Table 6 for the different modality inputs and lower and upper bounds using the ViT-B and ViT-H encoders.

Table 6: Trainable parameters and full model FLOPs.

Modalities	# M. Params	TFLOPS
RGB	139 – 178	1.4 – 6.3
RGB-D	141 – 184	2.7 – 12.7
RGB-D-N	144 – 191	4.1 – 19.0

E More Qualitative Results

Figure 8 shows some more qualitative results of our **RGB-D-N** model configuration. The results agree with our reported quantitative results in Table 1.

The comparison also shows that our model even surpasses the ground truth in some places: pillow on a sofa is classified as clutter, better detection of label edges, reliable class prediction for missing ground truth areas, etc.