

This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Trans4Map: Revisiting Holistic Bird's-Eye-View Mapping from Egocentric Images to Allocentric Semantics with Vision Transformers

Chang Chen, Jiaming Zhang, Kailun Yang, Kunyu Peng, Rainer Stiefelhagen CV:HCI Lab, Karlsruhe Institute of Technology

Abstract

Humans have an innate ability to sense their surroundings, as they can extract the spatial representation from the egocentric perception and form an allocentric semantic map via spatial transformation and memory updating. However, endowing mobile agents with such a spatial sensing ability is still a challenge, due to two difficulties: (1) the previous convolutional models are limited by the local receptive field, thus, struggling to capture holistic long-range dependencies during observation; (2) the excessive computational budgets required for success, often lead to a separation of the mapping pipeline into stages, resulting the entire mapping process inefficient. To address these issues, we propose an end-to-end one-stage Transformer-based framework for Mapping, termed Trans4Map. Our egocentric-toallocentric mapping process includes three steps: (1) the efficient transformer extracts the contextual features from a batch of egocentric images; (2) the proposed Bidirectional Allocentric Memory (BAM) module projects egocentric features into the allocentric memory; (3) the map decoder parses the accumulated memory and predicts the topdown semantic segmentation map. In contrast, Trans4Map achieves state-of-the-art results, reducing 67.2% parameters, yet gaining a + 3.25% mIoU and a + 4.09% mBF1 improvements on the Matterport3D dataset.¹

1. Introduction

Holistic scene understanding has a crucial role in both indoor and outdoor applications, *e.g.*, autonomous driving [30, 42, 43, 45], indoor exploration and navigation [17, 22, 28, 31], as well as indoor and outdoor mapping [5, 6, 11]. These tasks are ordinary for humans with excellent spatial perception ability, as they can continuously extract information from the egocentric perspective, and construct the scene semantic map in the allocentric perspective through the memory and spatial transformation.



Figure 1: The egocentric-to-allocentric semantic mapping. Given a front-view image sequence of length N observed along a trajectory (the red dash line), Trans4Map performs the online extract-project-segment pipeline, yielding an allocentric semantic map in bird's eye view.

However, the semantic mapping is still difficult for an artificial intelligent mobile agent, particularly when exploring an unfamiliar environment. In this work, we focus on the image-based semantic mapping task, by predicting allocentric semantic segmentation from egocentric images. As the example shown in Fig. 1, given a trajectory in the scene, which is composed of a batch of first-view RGB images and the corresponding known camera pose, the mobile agent performs three steps: (1) extracting rich and compact contextual features; (2) projecting and updating the egocentric features in the online intermediate allocentric memory as spatial-semantic representation of the complex spaces; (3) parsing and predicting the final top-view semantic mapping via the decoder. The image-based egocentric-to-allocentric mapping pipeline is more in line with human intuition, and is able to perform mapping in an efficient way, avoiding the need for a time-consuming reconstruction phase [13].

In Fig. 2, image-based semantic mapping methods are divided into four pipelines. The project-then-segment pipeline (Fig. 2a) projects N high-resolution observations into Bird's Eye View (BEV), that hinders the small object segmentation due to the lack of fine visual information. The segment-then-project pipeline (Fig. 2b) depends heavily on the front-view segmentation performance and may accumulate errors from one to another stage. The offline project-

^{*}Correspondence: jiaming.zhang@kit.edu

¹Code at: https://github.com/jamycheung/Trans4Map.



Figure 2: Semantic mapping pipelines. The two-stage pipelines (a)(b)(c) differ from the projection locations, *i.e.*, early, late, and intermediate offline projection. The one-stage pipeline (d) avoids 2.5TB storage by using online projection and has a higher mIoU.

then-segment pipeline (Fig. 2c) requires large-scale local storage to save the feature map given by the pre-trained encoder of the first stage. It further demands huge GPU memory to reload offline features for the second stage training. Unlike two-stage pipelines above, our proposed on-line project-then-segment pipeline (Fig. 2d) performs online implicit projection and enables end-to-end and resource-friendly BEV semantic mapping. The one-stage pipeline is crucial, because it fits resource-limited platforms, *e.g.*, robots. Further, it helps mobile agents to quickly construct maps and get familiar with the unknown space.

However, a lightweight but effective backbone that requires few resources is the decisive factor to achieve successful one-stage semantic mapping. The vision transformer architecture [40] is able to capture long-distance contextual dependencies, forming a non-local representation. This mechanism naturally fits the semantic mapping task, since the mapping process demands a holistic understanding of scenes. This assumption leads us to revisit the top-down semantic mapping with a transformerbased model and put forward a novel end-to-end one-stage Trans4Map framework. It delivers two primary benefits: (1) the long-range feature modeling ability is advantageous to obtain a more comprehensive spatial representation during the egocentric observation process; (2) the efficient and lightweight model structure enables the one-stage endto-end mapping pipeline. Besides, unlike the previous method [3] using a single GRU cell to reload the offline features, we propose a novel Bidirectional Allocentric Memory (BAM) to combine features from both directions, which can avoid the occluded objects to be classified as other category, e.g., chairs under tables. Further, our BAM implicitly performs the efficient online projection, as another key point in implementing the one-stage mapping pipeline (Fig. 2d).

To succeed in the end-to-end one-stage Trans4Map framework, we investigate a vast number of advanced deep architectures [24, 25, 41, 47]. According our experiments



Figure 3: Semantic mapping scores (mIoU) evaluated by CNN- and Transformer-based backbones with different numbers of parameters (M). Trans4Map models achieve better results yet with fewer parameters.

in Fig. 3, we found that simply applying transformer-based backbones does not guarantee improvement. Thanks to the proposed framework and BAM module, our Trans4Map models have much fewer parameters, yet achieve surprising semantic mapping scores. The B2 version reduces 67.2% parameters compared to SMNet [3] and sets a new state-of-the-art of a >40\% mIoU on the Matterport3D [4] dataset.

To summarize, we present the following contributions:

- Rethink the top-down semantic mapping task in a onestage pipeline to fit resource-constrained platforms.
- Propose an end-to-end *Transformer for Mapping* (*Trans4Map*) framework to perform egocentric-to-allocentric semantic mapping, yielding a holistic dense understanding for indoor exploration.
- Put forward a novel *Bidirectional Allocentric Memory* (*BAM*) which accumulates and projects the egocentric feature to the allocentric spatial tensor, via combining online memories from two directions.
- Our framework outperforms state-of-the-art counterparts on both the Matterport3D dataset and the Replica dataset, while using a more lightweight model.



Figure 4: **The overview of the end-to-end Trans4Map framework.** There are a transformer-based encoder for extracting the egocentric features from the RGB images, a Bidirectional Allocentric Memory (BAM) to project and accumulate the extracted feature sequence to the allocentric feature map via the known depth and pose information, and a CNN-based decoder for parsing the accumulated feature and predicting the allocentric semantics.

2. Related Work

Semantic Mapping. Recently, a number of approaches have emerged around semantic mapping. Semantic SLAM pipelines [13, 32] forward the images into a segmentation network and then project predicted labels onto the topview map. These efforts follow the segment-then-project pipeline, which is particularly rigorous for depth information, *i.e.*, the global coordinates of each pixel in RGB images. Unfortunately, a slight error can lead to an offset in the projection, as well as an under-fitting of model training. The project-then-segment pipeline [36] loses a vast amount of visual information during the projection phase, which hinders the small object segmentation. In contrast, SMNet [3] performs offline project-then-segment pipeline, which trains the encoder and decoder separately in two stages and does not optimize the training process from the first-view inputs to the top-view semantics as a whole.

Lu *et al.* [26] proposed an end-to-end network encoding the front-view information of the driving scene utilizing a variational encoder-decoder network [19] and then decoding it into a 2D top-down view. Pan *et al.* [29] represented a cross-view network with a View Parsing Network (VPN) - an MLP that parses semantics across different views. Both methods predict a local semantic topdown map with an end-to-end network from egocentric observations. These methods do not encode depth information, so the objects on the semantic map do not reflect their geometry structure. Moreover, there are many Bird's-Eyeview (BEV) semantic segmentation approaches for driving scene perception [2, 8, 10, 12, 20, 34, 44] emerging in the field. BEVFormer [20] aggregates spatiotemporal cues from surround-view cameras, whereas ViT-BEVSeg [9] uses a spatial transformer decoder for generating semantic occupancy grid maps. Differing from these works, we revisit various backbones including CNN- or Transformerbased and propose a transformer-based end-to-end framework, which acts as a one-stage BEV semantic mapper for holistic *indoor* scene understanding. In addition, the alignment and geometric structure of objects in the generated semantic map are successfully obtained.

Allocentric Spatial Memory. Incrementally generating a top-view map from egocentric observations requires dynamically updating the allocentric memory, *i.e.*, aggregating information over time, such as a mobile agent traveling around an indoor scene. Visual SLAM pipelines [1, 27, 38] for this task involve multiple modules such as tracking, mapping, relocalization, loop closure, and graph optimization through bundle adjustment. MapNet [14] developed RNN to update memory and registers new observations via dense matching. Tung et al. [39] proposed Geometry-aware Recurrent Neural Networks (GRNNs) to segment objects in 3D. This work is very memory-demanding due to the highdimension features. The closest work to our approach is SMNet [3] which uses a GRU to update the projected tensor. Unlike previous works, we propose a bidirectional allocentric memory that can better accumulate information over time and segment occluded objects. This crucial design allows Trans4Map to perform implicit online projection, enabling the one-stage mapping pipeline and setting the new state-of-the-art in the large-scale indoor scenes.

3. Methodology

In this section, we revisit the allocentric semantic mapping task with vision transformers and introduce our proposed end-to-end one-stage framework, that can generate allocentric semantic maps from egocentric observations.

3.1. Trans4Map: Framework Overview

As shown in Fig. 4, our end-to-end Trans4Map framework includes three steps: (1) the incoming N egocentric images are fed into the transformer-based backbone (in Sec. 3.2), which extracts contextual feature and longrange dependency; (2) the Bidirectional Allocentric Memory (BAM) module (in Sec. 3.3) projects the extracted feature via the depth-based transformation index; (3) the lightweight CNN-based decoder parses the projected feature and predicts the allocentric semantics.

One-stage pipeline. An efficient one-stage semantic mapping pipeline is crucial for fast constructing map and is required to deploy mobile agents on resource-limited platforms. The previous SMNet [3] uses a two-branch CNNbased backbone to extract features from RGBD images. Due to the heavy dual-branch backbone, the data pipeline is divided into multiple stages: extracting RGBD feature maps by using a frozen encoder; storing the feature maps locally; and reloading them to fine-tune the decoder. Unlike such a multi-stage data flow, our framework shown in Fig. 4 operates in an one-stage end-to-end manner, benefiting from three designs: (1) a transformer-based backbone is leveraged to capture holistic features and long-range dependencies, instead of narrow-receptive-field CNN-based backbones; (2) a single branch structure for extracting RGB features makes the whole model more lightweight than the dual-branch one; (3) an online training pipeline from egocentric images to allocentric semantics is constructed. avoiding using the time-consuming two-stage process and feature maps storage. Following the one-stage pipeline, our framework can achieve superior allocentric semantic mapping, while maintaining efficiency.

3.2. Transformer backbone

To fully investigate the proposed Trans4Map framework, we explore different model architectures and learning modalities for the allocentric semantic mapping task, as shown in Fig. 5. The architectures are constructed by four stages, and each stage includes a series of convolutional blocks (see Fig. 5c) or self-attention blocks (see Fig. 5a). Different from the convolutional architecture, the transformer-based architecture is able to capture non-local features, thanks to the self-attention operation [40]. Considering that cross-modality complementary features are informative for predicting semantics [15, 16, 23], we leverage RGB-Depth inputs and a multimodal architecture (see Fig. 5b) is reformed by using efficient self-attention blocks.



For brevity, we describe the operation of the singlemodal process, while the bimodal process involves an additive fusion at each stage, in which the fusion block obtains the extracted contextual features and geometry features and then fuses them per pixel with the same dimension. Given a batch of RGB images of size $N \times H \times W \times 3$, the divided patches are passed through the four-stage transformer blocks, to obtain the hierarchical feature representation with downsampling rates of $\{\frac{1}{r_1}, \frac{1}{r_2}, \frac{1}{r_3}, \frac{1}{r_4}\}$ and increasing channels of $\{C_1, C_2, C_3, C_4\}$. Then, the multiscale features are concatenated by an MLP layer and followed by a convolution layer with 64 channels. So the hierarchical features are fused into an egocentric feature of size $N \times \frac{H}{r_1} \times \frac{W}{r_1} \times 64$. To study different semantic mapping architectures, the multi-scale features in this work are extracted with $\{\frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}\}$ downsampling rates and $\{64, 128, 320, 512\}$ channels.

To compare CNN- and transformer-based models, the RedNet backbone [16] used in SMNet and the ConvNeXt backbone are selected to form CNN-based mapping models, while transformer-based models include FAN [47], Swin [24], and SegFormer [41] backbones. Based on our experiments, we adopt SegFormer [41] as the default backbone of our visual encoder, as its simple and lightweight design can generate features ranging from high-resolution fine features to low-resolution coarse features. More ablation studies and discussions are unfolded in Sec. 4.4.

3.3. Bidirectional Allocentric Memory

After acquiring the egocentric features through the aforementioned transformer backbone, the projective index is needed to project representative contextual features into an allocentric memory map. In the Habitat simulator [35], we can directly obtain the state of the moving agent and then calculate the camera pose using relative orientation and position. In order to perform the online projection, we need to derive the 3D position of each pixel in the egocentric image, as presented in Eq. (1) and Eq. (2).

$$\begin{bmatrix} x\\ y\\ z \end{bmatrix}_{c} = K^{-1} \begin{bmatrix} u\\ v\\ d_{u,v} \end{bmatrix}_{i}$$
(1)

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{w} = R^{-1} \begin{bmatrix} x \\ y \\ z \end{bmatrix}_{c} - \overrightarrow{t}$$
(2)

K in Eq. (1) is the camera intrinsic parameter matrix, $[R|\vec{t}]$ in Eq. (2) are the rotation matrix and the translation matrix, respectively. First, in Eq. (1), using the pinhole camera model and the depth of each pixel $d_{u,v}$, the pixel coordinate (u, v) in the image coordinate system can be converted into the camera coordinate system. Then, in Eq. (2), the camera coordinates denoted as (x, y, z) of each point are converted to world coordinates denoted as (X, Y, Z) using the rotation matrix and the translation matrix.

Each pixel in the allocentric memory map represents a $2cm \times 2cm$ cell in the scene of the Matterport3D dataset [4], so the projective index $(i, j)_m$ can be calculated by dividing the world coordinates X and Z of each point by the resolution. Finally, we project egocentric features of N batch size onto the allocentric memory map using the calculated projective index.

To enhance the long-range content dependency and aggregate the incoming information completely, we propose Bidirectional Allocentric Memory (*BAM*), in which we transmit projected features via Bi-directional GRU (Bi-GRU), which are used to update and accumulate incoming observations from two directions. Specifically, BAM adds a reverse GRU unit that performs feature parsing and accumulation. As shown in Fig. 6, the upper GRU unit processes the allocentric memory tensor in a forward direction from the M^{t-1} to the M^t feature and the lower GRU unit in a backward direction from the M^t to the M^{t-1} feature. A simple yet effective convolutional layer is applied to fuse two projected allocentric memory features. The computation of updated spatial memory tensor is formulated as:

$$M_{i,j}^{t} = GRU(F_{i,j}^{t}, M_{i,j}^{t-1});$$
(3)

$$M_{i,j}^{t-1} = GRU(F_{i,j}^{t-1}, M_{i,j}^t);$$
(4)

$$T = Conv(M_{i,j}^{t}, M_{i,j}^{t-1}).$$
(5)

 $M_{i,j}^t$ and $M_{i,j}^{t-1}$ are the current time step spatial memory and previous time step spatial memory, respectively. The fused spatial memory tensors T are accessible to the decoding step for the final semantic top-down map prediction. Thanks to the bidirectional parsing process, BAM is able to accumulate the observations per each time step in both direction in parallel, thus, it can better avoid occluded objects



Figure 6: Bidirectional Allocentric Memory (BAM).

being wrongly-classified. With BAM, Trans4Map can produce a more meaningful allocentric representation which combines bidirectional projected features.

4. Experiment

4.1. Datasets

Matterport3D. The Matterport3D dataset [4] contains photo-realistic scans of 90 building-scale environments. It provides RGBD images and 3D annotations with 40 categories, which are crucial for scene understanding tasks. We follow the same dataset split as SMNet [3], consisting of 61 training scenes, 7 validation scenes, and 17 test scenes, which mainly focuses on 12 object categories: chair, table, cushion, cabinet, shelving, sink, dresser, plant, bed, sofa, counter, fireplace. Other rare objects and floor surfaces are masked as void classes. Given a trajectory, we sample 50 sets of N consecutive waypoints in each unique scene via Habitat simulator [35]. The N consecutive RGBD image sequences as a batch are forwarded into the model as input. **Replica**. The Replica dataset [37] contains 18 various highly photo-realistic indoor environments. It provides dense-mesh, high-resolution RGBD images and also a large range of instance annotations with 88 categories. We follow the same setup as in the Matterport3D dataset, *i.e.*, focusing on 12 object categories. We test our model on the Replica using the weight that was trained on Matterport3D dataset, so the 18 scenes are included in the test split.

4.2. Implementation details

Performing the online project-segment paradigm, we train our model in an end-to-end manner. For the ablation study, we revisit the transformer-based mapping and investigate several visual encoders in our model which are pre-trained on the ImageNet [33] and ADE20K [46], and the BAM module and the decoder are randomly initialized. We use the AdamW optimizer [18] to train our model with four 1080Ti GPUs on the Matterport3D dataset. The learning

Table 1: Allocentric semantic mapping results on the Matterport3D dataset. Model[†] is our implementation.

Method	Acc	mRecall	mPrecision	mIoU mBF1
Seg. $GT \rightarrow Proj.$	89.49	73.73	74.58	59.73 54.05
Two-stage Proj. \rightarrow Seg.	83.18	27.32	35.30	19.96 17.33
Two-stage Seg. \rightarrow Proj.	88.06	40.53	58.92	32.76 33.21
Two-stage Semantic SLAM	85.17	37.51	51.54	28.11 31.05
Two-stage SMNet	88.14	47.49	58.27	36.77 37.02
Two-stage SMNet †	89.14	46.34	56.98	36.16 35.95
One-stage Trans4Map	89.02	54.50	56.20	40.02 41.11

Table 2: Allocentric semantic mapping results on the Replica dataset. Model[†] is our implementation. Note that the last two rows are evaluated on the partially available Replica [37] dataset, while the others have all data as [3].

Method	Acc	mRecall	mPrecision	mIoU	mBF1
Seg. $GT \rightarrow Proj.$	96.83	83.84	94.05	79.76	86.89
Two-stage Seg. \rightarrow Proj.	88.61	48.11	65.20	40.77	45.86
Two-stage Semantic SLAM	88.30	45.80	62.41	37.99	46.71
Two-stage SMNet	89.26	53.37	64.81	43.12	45.18
Two-stage SMNet †	87.69	58.88	34.85	27.68	42.67
One-stage Trans4Map	86.19	65.27	34.91	29.15	48.66

rate is initialized with $6e^{-5}$ and then scheduled by LambdaLR. We use Cross-Entropy as the loss function. Training with 100 epochs will take about 30 hours. Similar to [3], the evaluation metrics include the pixel-wise accuracy (Acc), the pixel recall (mRecall) and precision (mPrecision) scores, the intersection-over-union (mIoU) score, and the boundary F1 (mBF1) score [7].

4.3. Allocentric semantic mapping results

As shown in Table 1 and Table 2, a set of results are conducted on the Matterport3D and the Replica datasets. The experiments assess four pipelines discussed in Fig. 2, *i.e.*, two-stage project-then-segment, segment-then-project, offline project, and our one-stage online Trans4Map.

Matterport3D. The results on Matterport3D are in Table 1. Following the segment-project paradigm, the result obtained by using the label data is the upper bound performance. As in Table 1, the segment-project baseline performs much better than the project-segment one, since part of information will be lost in the process of converting an egocentric image into a top-down view. The semantic SLAM in [13] also uses the segment-project method but achieves worse performance than the image-based segment-project baseline. SMNet [3] follows the offline project-segment paradigm and adds a spatial memory update module. Here, we reproduce the experiment using the released code under the same condition and obtain the results with a mIoU score of 36.16% and a mBF1 value of 35.95%. Com-

Table 3: **Comparison of training resources** between twostage and one-stage semantic mapping pipeline. The metrics include the local storage (TB), one epoch training time (h:hour), one epoch data loading time (h:hour), RAM requirement (GB), #Param (M), and mIoU (%).

Method	Storage (TB)	Train (h)	Load (h)	RAM (GB)	#Param (M	l) mIoU (%)
Two-stage	e 2.5	6.00	2.00	256	83.9	36.77
One-stage	e 0	0.33	0.01	18	27.5	40.88
Change	-100%	-94.5%	-99.5%	-93.0%	-67.2%	+11.2%

Table 4: **Comparison of different backbones** including the number of parameters and performance on Matterport3D. Except for two-stage SMNet, all methods are using one-stage Trans4Map framework, so as to ablate the backbone.

Method	Backbone	#Param (M)	mIoU (%)
SMNet	RedNet [3]	83.9	36.77
ConvNeXt	ConvNeXt-T [25]	31.1	35.91
ConvNeXt	ConvNeXt-S [25]	52.7	36.49
FAN	FAN-T [47]	09.8	31.07
FAN	FAN-S [47]	31.3	34.62
Swin	Swin-T ^[24]	30.8	34.19
Swin	Swin-S [24]	52.1	36.80
Trans4Map	MiT-B1 [41]	15.4	34.38
Trans4Map	MiT-B2 [41]	27.5 (-56.4)	40.02 (+3.25)
Trans4Map	MiT-B3 [41]	47.4 (-36.5)	39.98 (+3.21)
Trans4Map	MiT-B4 [41]	64.1 (-19.8)	40.88 (+4.11)

pared to SMNet, our Trans4Map model achieves significant improvements in terms of mIoU (40.02%) and mBF1 (41.11%) on the Matterport3D dataset, which proves the effectiveness of our proposed allocentric mapping framework. **Replica**. The results on Replica are in Table 2. All models are trained on the Matterport3D dataset and tested on the Replica dataset. Note that the trajectories and labels of the Replica dataset are partially available at the moment, thus, the results are tested on the constrained data of the Replica dataset. Nonetheless, under the same condition with the same label data, our Trans4Map outperforms the baseline SMNet with a 1.47% mIoU and a 5.99% mBF1 improvements, respectively. The results indicate that our Trans4Map framework achieves consistent improvements across different datasets.

4.4. Ablation study

Analysis of One-stage Pipeline. To analyze the efficiency of different semantic mapping pipelines, we present the training process of the two-stage [3] and our one-stage pipeline in Table 3. The two-stage method saves the intermediate feature maps offline and reloads them for a secondstage fine-tuning. Compared to the two-stage method, our one-stage method performing online mapping does not require extra local storage (0TB vs. 2.5TB). Thanks to our efficient transformer-based model, the one-stage pipeline achieves faster training (0.33h vs. 6h) and loading (0.01h vs. 2h) processes than the two-stage one. Fur-

Table 5: Ablation study of the BAM in Trans4Map. Methods are based on Mit-B4 and four sampling points.

	Method	mIoU
	GRU (SMNet [3])	36.77
1 2 3 4 5 6 7	GRU + Conv GRU + 2 GRUCells BiGRU + Concatenate BiGRU + gMLP [21] + Concatenate BiGRU + gMLP [21] + Conv Fusion BiGRU + 2 GRUCells + Conv Fusion BiGRU + Conv Fusion (our BAM)	37.86 37.67 36.73 37.49 40.15 40.00 40.44

ther, our RAM requirement (18GB vs. 256GB) and #Param (27.5M vs. 83.9M) are much lower, which is crucial for resource-limited mobile platforms. Surprisingly, the one-stage pipeline surpasses the two-stage one with +4.11% mIoU gains. There are sufficient improvements while saving resources, demonstrating the effectiveness of our proposed one-stage semantic mapping pipeline.

Analysis of Encoder. Based on the one-stage pipeline, we analyze the effects of using CNN-based or transformerbased backbones to extract features. Table 4 presents the comparison between the two-stage SMNet and the other methods that are all based on the one-stage Trans4Map framework. The baseline SMNet using CNN-based backbone has the largest number of parameters but achieves only 36.77% in mIoU. We found that simply applying an advanced CNN-based backbone (ConvNeXt [25]) or transformer-based backbones (FAN [47] and Swin [24]) does not lead to sufficient improvement. When their backbones are much more lightweight, their performances are even a bit inferior to SMNet. In contrast, our Trans4Map equipped with the MiT-B2 backbone [41] reduces 67.2%parameters compared to SMNet, i.e., from 83.9M to 27.5M, but the performance has made a sufficient leap in mIoU (+3.25%). The MiT-B3 and -B4 backbone with 47.4M and 64.1M parameters bring similar mIoU gains (+3.21% and +4.11%) as the MiT-B2 backbone, because a larger transformer model often requires more training data to obtain a desired boost. The results of our proposed Trans4Map confirm that transformer-based models fit the semantic mapping task and can maintain the desired performance with a good trade-off of accuracy and efficiency.

Analysis of BAM. Apart from the encoder, we further analyze different structures of our *BAM* module. As shown in Table 5, compared the baseline GRU in SMNet and **1**2, simply stacking more GRU cells or convolutional layers does not improve model performance. Our proposed convolutional fusion is better than the concatenation fusion, yield-ing +2.77% gains ($\mathfrak{S} \rightarrow \mathfrak{O}$) and +2.66% gains ($\mathfrak{C} \rightarrow \mathfrak{S}$), respectively. The method (\mathfrak{S}) is to follow an advanced gMLP block [21] for token information mixing after the Bi-GRU module and yields a mIoU of 40.15%. The method

Table 6: Analysis of model complexities, data sources, sampling points in a trajectory, and data modalities.

	Method	Pre-train	#Points	Modality	mIoU
	SMNet [3]	SUN-RGBD	20	RGBD	36.77
0	B2	ImageNet	4	RGB	37.86
2	B2	SUN-RGBD	4	RGBD	40.27
3	B2	ADE20K	4	RGBD	40.15
4	B2	ADE20K	4	RGB	37.71
6	B2	ADE20K	20	RGB	40.02
6	B3	ADE20K	4	RGB	38.78
0	B3	ADE20K	20	RGB	39.98
8	B4	ADE20K	4	RGB	40.44
9	B4	ADE20K	20	RGB	40.88

(**6**) equipped with two GRU cells obtains 40.00% in mIoU. Our *BAM* (**7**) applies BiGRU and a convolutional layer to process the spatial tensors and achieves the best performance with 40.44% in mIoU. The ablation study demonstrates that our BAM is crucial for semantic mapping.

Analysis of pre-train sources. Further, we ablate three different pre-train sources. Comparing 2 and 3 in Table 6, Trans4Map model benefits more from the SUN-RGBD dataset than from the ADE20K dataset, since the ADE20K weights are shared between RGB and depth branches, while the SUN-RGBD ones are separated. Between 1 and 3, both Trans4Map models pre-trained on ImageNet and ADE20K have comparable results. It indicates that our single-modal Trans4Map is stable and robust across different pre-training settings.

Analysis of sampling points. The number of sampling points is a key factor in obtaining dense observations. Theoretically, as the sequence of input images increases, the richer the resulting map will be. In (2) and (5) of Table 6, increasing the number of sampling points from 4 to 20, Trans4Map with MiT-B2 benefits from the dense observations and obtains a gain of +2.31% in mIoU. Comparing (3) and (9), the consistent improvement is achieved by Trans4Map with MiT-B4.

Analysis of data modalities. To ablate the effect of data modalities, RGB and RGBD inputs are compared by using single-modal and dual-modal Trans4Map. Compared between the RGBD baseline SMNet and *Q* in Table 6, our dual-modal Trans4Map has a 3.5% mIoU gain, while both are trained on SUN-RGBD. When using 4 sampling points, the dual-modal Trans4Map performs better than the single-modal one, as compared in *S* and *G*. It indicates the effectiveness of our dual-modal Trans4Map in harvesting cross-stream complementary features for boosting performance.

Analysis of model complexities. To inspect the trade-off between efficiency and performance, we analyze the model complexity by using three different backbones, *i.e.*, MiT-B2, -B3 and -B4. They have respective 27.5M, 47.4M and 64.1M parameters, yet all are more lightweight than the baseline SMNet with 83.9M, as presented in Table. **4**. Com-



Figure 7: Allocentric semantic mapping visualizations. There are two indoor scenes from the Matterport3D test set. From left to right are the predicted results of SMNet, the results of our Trans4Map and the ground truth. Zoom in for better view.



Figure 8: Visualizations of the challenging case. Zoom in for better view.

paring three backbones (**16(3)**, larger models have better results, since more training data is available based on 4 sampling points. In the case of 20 sampling points (**57)**, the three models achieve competitive results compared to the baseline model. This experiment demonstrates the effectiveness of our Trans4Map framework, achieving consistent improvements based on backbones of different sizes.

4.5. Qualitative analysis

Semantic map visualizations. We visualize the semantic map results from the test set of the Matterport3D dataset, as shown in Fig. 7. Thanks to the extracted non-local features and long-distance dependencies, Trans4Map has much better segmentation results. In the first scene in Fig. 7, Trans4Map is better at segmenting the bed. Further, Trans4Map is able to successfully classify the fireplace, while the baseline model fails and predicts it as a cabinet. In the second scene in Fig. 7, Trans4Map delivers semantic mapping accurately, such as on *cabinet* and *chair* categories, while SMNet misclassifies them as tables. Besides, SMNet yields incomplete *chair* segmentation results. Challenging case analysis. The semantic mapping in the large-scale indoor scenes is still a challenging task. Fig. 8 demonstrates that both models struggle with semantic mapping in such a large-scale scene. SMNet has difficulty segmenting the four *chairs* completely as shown in Fig. 8 (a). Trans4Map performs slightly better, but there is still a large improvement space. All the six *sinks* in the lower part are not predicted correctly. The reason is that the depth information of this scene is less reliable, while the whole scene covers about 2500 square meters on the allocentric map. One potential solution is to obtain more observations to alleviate the error caused by the noisy depth measurements.

5. Conclusion

In this paper, we propose an end-to-end transformerbased framework, termed *Trans4Map*, in order to revisit the egocentric-to-allocentric mapping from the frontview images to bird's-eye-view semantics. Based on the transformer-driven backbone and the *Bidirectional Allocentric Memory (BAM)* updater, Tran4Map sets the new state of the art on both Matterport3D and Replica datasets, while using a more lightweight architecture and having fewer parameters as compared to the previous work. In the future, we will further explore domain adaptation methods to transfer the mapping models trained on the synthetic datasets to real-world scenes. Based on the constructed semantic map, the downstream tasks such as indoor navigation and path planning will also be interesting research directions.

References

- Carlos Campos, Richard Elvira, Juan J. Gómez Rodríguez, José M. M. Montiel, and Juan D. Tardós. ORB-SLAM3: An accurate open-source library for visual, visual–inertial, and multimap SLAM. *T-RO*, 2021. 3
- [2] Yigit Baran Can, Alexander Liniger, Danda Pani Paudel, and Luc Van Gool. Structured bird's-eye-view traffic scene understanding from onboard images. In *CVPR*, 2021. 3
- [3] Vincent Cartillier, Zhile Ren, Neha Jain, Stefan Lee, Irfan Essa, and Dhruv Batra. Semantic MapNet: Building allocentric semantic maps and representations from egocentric views. In AAAI, 2021. 2, 3, 4, 5, 6, 7
- [4] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. In *3DV*, 2017. 2, 5
- [5] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural SLAM. In *ICLR*, 2020. 1
- [6] Hao Chen, Weijian Hu, Kailun Yang, Jian Bai, and Kaiwei Wang. Panoramic annular SLAM with loop closure and global optimization. AO, 2021. 1
- [7] Gabriela Csurka, Diane Larlus, Florent Perronnin, and France Meylan. What is a good evaluation measure for semantic segmentation?. In *BMVC*, 2013. 6
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3
- [9] Pramit Dutta, Ganesh Sistu, Senthil Yogamani, Edgar Galván, and John McDonald. ViT-BEVSeg: A hierarchical transformer network for monocular birds-eye-view segmentation. In *IJCNN*, 2022. 3
- [10] Isht Dwivedi, Srikanth Malla, Yi-Ting Chen, and Behzad Dariush. Bird's eye view segmentation using lifted 2D semantic features. In *BMVC*, 2021. 3
- [11] Georgios Georgakis, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, and Kostas Daniilidis. Learning to map for active semantic goal navigation. In *ICLR*, 2022. 1
- [12] Nikhil Gosala and Abhinav Valada. Bird's-eye-view panoptic segmentation using monocular frontal view images. *RA-L*, 2022. 3
- [13] Margarita Grinvald, Fadri Furrer, Tonci Novkovic, Jen Jen Chung, Cesar Cadena, Roland Siegwart, and Juan Nieto. Volumetric instance-aware semantic mapping and 3D object discovery. *RA-L*, 2019. 1, 3, 6
- [14] João F. Henriques and Andrea Vedaldi. MapNet: An allocentric spatial memory for mapping environments. In *CVPR*, 2018. 3
- [15] Xinxin Hu, Kailun Yang, Lei Fei, and Kaiwei Wang. ACNet: Attention based network to exploit complementary features for RGBD semantic segmentation. In *ICIP*, 2019. 4
- [16] Jindong Jiang, Lunan Zheng, Fei Luo, and Zhijun Zhang. RedNet: Residual encoder-decoder network for indoor rgb-

d semantic segmentation. *arXiv preprint arXiv:1806.01054*, 2018. 4

- [17] Kapil D. Katyal, Adam Polevoy, Joseph L. Moore, Craig Knuth, and Katie M. Popek. High-speed robot navigation using predicted occupancy maps. In *ICRA*, 2021. 1
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [19] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 3
- [20] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. BEVFormer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. In ECCV, 2022. 3
- [21] Hanxiao Liu, Zihang Dai, David R. So, and Quoc V. Le. Pay attention to MLPs. In *NeurIPS*, 2021. 7
- [22] Huayao Liu, Ruiping Liu, Kailun Yang, Jiaming Zhang, Kunyu Peng, and Rainer Stiefelhagen. HIDA: Towards holistic indoor understanding for the visually impaired via semantic instance segmentation with a wearable solid-state lidar sensor. In *ICCVW*, 2021. 1
- [23] Huayao Liu, Jiaming Zhang, Kailun Yang, Xinxin Hu, and Rainer Stiefelhagen. CMX: Cross-modal fusion for RGB-X semantic segmentation with transformers. arXiv preprint arXiv:2203.04838, 2022. 4
- [24] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 2, 4, 6, 7
- [25] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A ConvNet for the 2020s. In *CVPR*, 2022. 2, 6, 7
- [26] Chenyang Lu, Marinus Jacobus Gerardus van de Molengraft, and Gijs Dubbelman. Monocular semantic occupancy grid mapping with convolutional variational encoderdecoder networks. *RA-L*, 2019. 3
- [27] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: A versatile and accurate monocular SLAM system. *T-RO*, 2015. 3
- [28] Wenyan Ou, Jiaming Zhang, Kunyu Peng, Kailun Yang, Gerhard Jaworek, Karin Müller, and Rainer Stiefelhagen. Indoor navigation assistance for visually impaired people via dynamic SLAM and panoptic segmentation with an RGB-D sensor. In *ICCHP-AAATE*, 2022. 1
- [29] Bowen Pan, Jiankai Sun, Ho Yin Tiga Leung, Alex Andonian, and Bolei Zhou. Cross-view semantic segmentation for sensing surroundings. *RA-L*, 2020. 3
- [30] Kunyu Peng, Juncong Fei, Kailun Yang, Alina Roitberg, Jiaming Zhang, Frank Bieder, Philipp Heidenreich, Christoph Stiller, and Rainer Stiefelhagen. MASS: Multi-attentional semantic segmentation of LiDAR data for dense top-view understanding. *T-ITS*, 2022. 1
- [31] Santhosh K Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Occupancy anticipation for efficient exploration and navigation. In ECCV, 2020. 1
- [32] Teng Ran, Liang Yuan, Jianbo Zhang, Dingxin Tang, and Li He. RS-SLAM: A robust semantic SLAM in dynamic environments based on RGB-D sensor. *IEEE Sensors Journal*, 2021. 3

- [33] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV*, 2015. 5
- [34] Avishkar Saha, Oscar Mendez Maldonado, Chris Russell, and Richard Bowden. Translating images into maps. In *ICRA*, 2022. 3
- [35] Manolis Savva, Jitendra Malik, Devi Parikh, Dhruv Batra, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, and Vladlen Koltun. Habitat: A platform for embodied AI research. In *ICCV*, 2019. 4, 5
- [36] Suriya Singh, Anil Batra, Guan Pang, Lorenzo Torresani, Saikat Basu, Manohar Paluri, and C. V. Jawahar. Selfsupervised feature learning for semantic segmentation of overhead imagery. In *BMVC*, 2018. 3
- [37] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard A. Newcombe. The replica dataset: A digital replica of indoor spaces. arXiv preprint arXiv:1906.05797, 2019. 5, 6
- [38] Sebastian Thrun. Probabilistic robotics. *Communications of the ACM*, 2002. **3**
- [39] Hsiao-Yu Fish Tung, Ricson Cheng, and Katerina Fragkiadaki. Learning spatial common sense with geometry-aware recurrent networks. In *CVPR*, 2019. 3
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2, 4
- [41] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. SegFormer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, 2021. 2, 4, 6, 7
- [42] Weixiang Yang, Qi Li, Wenxi Liu, Yuanlong Yu, Yuexin Ma, Shengfeng He, and Jia Pan. Projecting your view attentively: Monocular road scene layout estimation via crossview transformation. In CVPR, 2021. 1
- [43] Jiaming Zhang, Kailun Yang, Chaoxiang Ma, Simon Reiß, Kunyu Peng, and Rainer Stiefelhagen. Bending reality: Distortion-aware transformers for adapting to panoramic semantic segmentation. In CVPR, 2022. 1
- [44] Yun Zhao, Yu Zhang, Zhan Gong, and Hong Zhu. Scene representation in bird's-eye view from surrounding cameras with transformers. In *CVPR*, 2022. **3**
- [45] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In CVPR, 2022. 1
- [46] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ADE20K dataset. In *CVPR*, 2017. 5

[47] Daquan Zhou, Zhiding Yu, Enze Xie, Chaowei Xiao, Anima Anandkumar, Jiashi Feng, and Jose M. Alvarez. Understanding the robustness in vision transformers. In *ICML*, 2022. 2, 4, 6, 7