

# Trans4Trans: Efficient Transformer for Transparent Object Segmentation to Help Visually Impaired People Navigate in the Real World

Jiaming Zhang, Kailun Yang\*, Angela Constantinescu, Kunyu Peng, Karin Müller and Rainer Stiefelhagen  
Karlsruhe Institute of Technology

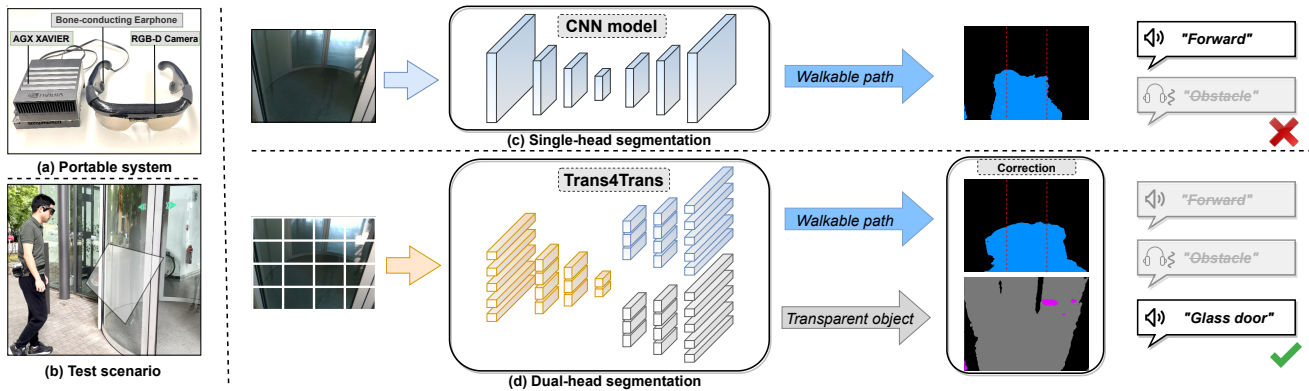


Figure 1: (a) the assistive system equipped with smart vision glasses and a portable GPU is tested (b) in front of a glass door. The input image is segmented as *Walkable Path* by (c) a single head CNN model, and is corrected as *Glass Door* by (d) our Transformer for Transparency (Trans4Trans) model, which is safety-critical for navigation. The user interface consists of vibration and voice feedback.

## Abstract

Common fully glazed facades and transparent objects present architectural barriers and impede the mobility of people with low vision or blindness, for instance, a path detected behind a glass door is inaccessible unless it is correctly perceived and reacted. However, segmenting these safety-critical objects is rarely covered by conventional assistive technologies. To tackle this issue, we construct a wearable system with a novel dual-head Transformer for Transparency (Trans4Trans) model, which is capable of segmenting general and transparent objects and performing real-time wayfinding to assist people walking alone more safely. Especially, both decoders created by our proposed Transformer Parsing Module (TPM) enable effective joint learning from different datasets. Besides, the efficient Trans4Trans model composed of symmetric transformer-based encoder and decoder, requires little computational expenses and is readily deployed on portable GPUs. Our Trans4Trans model outperforms state-of-the-art methods on the test sets of Stanford2D3D and Trans10K-v2 datasets and obtains mIoU of 45.13% and 75.14%, respectively. Through various pre-tests and a user study conducted in indoor and outdoor scenarios, the usability and reliability

of our assistive system have been extensively verified.

## 1. Introduction

Knowledge of glass architecture [6] and glass doors [44, 47] are particular important for visually impaired people, because transparent objects often present architectural barriers which hinder the mobility of people with low vision or blindness. For example, a path behind a glass door is not a free way to navigate (see Fig. 1) unless it is correctly recognized and reacted. However, most common vision-based navigation assistance systems [1, 60, 73] cannot handle transparent obstacles well, as 3D vision-based methods hardly recover the depth information of texture-less transparent surfaces [1, 73], whereas conventional image segmentation-based methods do not cover the categories of challenging transparent objects [38, 72]. In addition, guide dogs often get confused leading people with blindness to full-pane windows, and differentiation between doors, and large glass windows is difficult for people with residual sight [55]. A system that supports the recognition of landmarks such as doors is particularly appreciated by people with visual impairments, as finding a door before entering a building is difficult due to the inaccuracy of GPS [4, 55].

To address these issues, we propose a wearable system capable of real-time wayfinding and object segmentation to assist visually impaired individuals travel more

\*Corresponding author (e-mail: kailun.yang@kit.edu).

safely. We present *Trans4Trans*, precisely *Transformer for Transparency*, an efficient semantic segmentation architecture with dual heads, as shown in Fig. 1(d). As transparent objects are often texture-less or share similar content as the surroundings, it is essential to associate long-range visual concepts to robustly infer transparent regions. For this reason, *Trans4Trans* is established with both transformer-based encoder and decoder to fully exploit the long-range context modeling capacity of self-attention layers in transformers [59]. In particular, *Trans4Trans* features a novel *Transformer Paring Module (TPM)* to fuse multi-scale feature maps generated from embeddings of dense partitions, and the symmetric transformer-based decoder can consistently parse the feature maps from transformer-based encoder. Together with semantically predicting general things and stuff classes like walkable areas, the dual-head design allows to segment transparent objects accurately and completely, which are safety-critical for navigation.

*Trans4Trans* is integrated in our wearable system which comprises a pair of smart vision glasses and a mobile GPU processor, which delivers a holistic scene understanding swiftly and accurately thanks to the high efficiency of our model. With the complete semantic information, the user interface consists of a customized set of acoustic feedback via sonification of detected objects, walkable directions and warnings of the obstacles, which yields intuitive suggestions and no prior knowledge is needed. A comprehensive set of experiments has been conducted on multiple semantic segmentation datasets [2, 70]. In particular, the proposed model outperforms state-of-the-art methods on the test sets of Stanford2D3D [2] and Trans10K-v2 [70] datasets. Finally, a user study with visually impaired people and a variety of field tests demonstrate the usability and reliability of our assistive system for navigational perception in the wild. To the best of our knowledge, we are the first to use vision transformers for assisting people with visual impairment.

In summary, we deliver the following contributions:

- We present a wearable assistive system with a pair of smart vision glasses and a mobile GPU ported with vision transformers for visually impaired people.
- We propose an efficient semantic segmentation architecture *Transformer for Transparency (Trans4Trans)* with transformer-based encoder and decoder, a dual-head design to unify general object and challenging transparent object segmentation, and a *Transformer Parsing Module* to fuse multi-scale representations.
- *Trans4Trans*, maintaining a high efficiency, surpasses state-of-the-art CNN- and transformer-based methods on Stanford2D3D and Tran10K-v2 datasets.
- According to our designed system algorithm, we produce a customized set of acoustic feedback and conduct a user study and various field tests, demonstrating the usability and reliability of *Trans4Trans*.

## 2. Related Work

**Semantic segmentation for visual assistance.** Whereas traditional assistance systems rely on multiple monocular detectors and depth sensors [1, 13, 14, 37, 60], semantic segmentation allows to solve many navigational perception problems at once and thereby has been quickly employed in visual assistance. Yang *et al.* [72] put forward seizing semantic segmentation to unify detection tasks and assist terrain awareness, whereas Mao *et al.* [43] argued for panoptic segmentation towards a holistic sensing. In [42, 77], instance-specific segmentation methods like Mask R-CNN [23] were directly applied for content-aware surrounding understanding. Semantic segmentation has also been used to address intersection perception like detection of crosswalks, sidewalks, and blind roads [8, 25]. Moreover, it has received increasing interests and various systems appear in the field [16, 18, 26, 38, 48]. Yet, both traditional sensor-based and segmentation-driven approaches cannot handle challenging transparent obstacles well.

**Transparent object sensing.** Classical visual assistance systems [3, 27] resort to multi-sensor fusion, *e.g.*, fusing RGB-D cameras and ultrasonic sensors, to overcome the difficulties in dealing with transparent obstacles like glass objects, French windows, French doors, etc. Chen *et al.* [11] design a multimodal stereo matching algorithm to improve the depth measurements of transparent objects with dual depth sensors. Polarization cues [31] and reflection priors [36] are also frequently explored for transparency perception. For example, Xiang *et al.* [68] propose a polarization-driven semantic segmentation architecture by adaptively bridging RGB and polarization dimensions, which significantly lift the performance of classes with polarization properties like *glass*.

Recently, Xie *et al.* [69, 70] built the Trans10K dataset and show that while the pure RGB-based transparent object segmentation is a largely unsolved task, it is promising for real-world usages with the increased data amount. This allows the community to go beyond traditional perception regimens relying on sensor fusion schemes and develop novel methods addressing transparent object segmentation. For example, AdaptiveASPP [7] is designed to extract rich features of multiple fields-of-view with appropriate importances, whereas EBLNet [22] incorporates an edge-aware graph convolution module to model global shape representations. Differing from most of these accuracy-oriented methods, we aim for a both efficient and robust semantic segmentation desirable for navigation assistance. We establish a transformer-based system to assist the detection of transparent objects in real-world scenes.

**Efficient transformers for dense prediction.** Due to the capacity to model long-range contextual correlations, attention in transformers [59] has been introduced in visual recognition tasks to learn inter-dependencies either in the

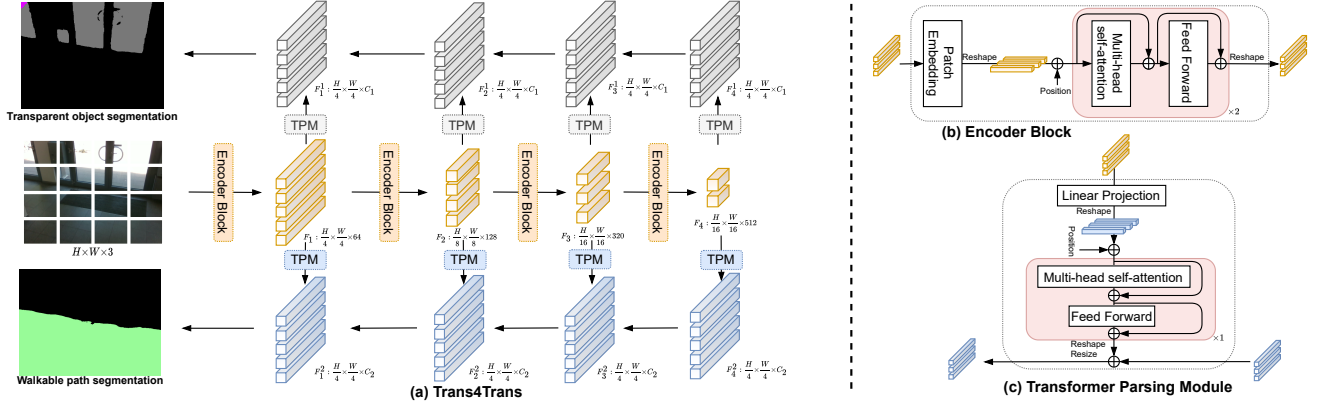


Figure 2: The architecture of (a) Trans4Trans model consists of shared encoder and dual decoders, while (b) and (c) are the general transformer-based encoder block and our proposed Transformer Parsing Module (TPM) for decoder, respectively.

channel or in the spatial dimension [19, 64, 79] by appending attention layers atop convolutional networks. To reduce the quadratic computation overhead of such non-local attention layers *w.r.t.*, the input size, their disentangled or asymmetric versions [28, 74, 76, 85] are constructed.

Recently, transformers are directly applied in vision tasks [9, 17, 84]. In ViT [17] and DeiT [58], a pure transformer is utilized to sequences of image patches for image recognition. For pixel-wise tasks, SETR [83] views semantic segmentation from a sequence-to-sequence perspective with vision transformers [17], whereas MaX-DeepLab [61] infers class masks with a dual-path transformer for panoptic segmentation. Inspired by their success, transformer architectures for dense prediction emerge [15, 40, 57, 66, 67, 71].

A vital set of these models is proposed with lightweight variants like Pyramid Vision Transformers (PVT) [63], ResT [80], and LeViT [20], aiming to optimize the accuracy-efficiency trade-off when porting transformers to real-world applications. In this work, we devise an efficient *Trans4Trans* framework with focus set on assisting navigation of visually impaired people in the wild. In contrast to existing works that either stack attention layers [19, 74] and encoder-decoder transformers on CNN backbones [70], or employ CNN-based decoders on top of transformer encoders [63, 83], in *Trans4Trans* both encoder and decoder are based on transformers, together with a novel Transformer Parsing Module design in our dual-head decoder.

### 3. System Architecture

#### 3.1. Trans4Trans

Inspired by the benefit of the ViT [17] transformer model in acquiring long-range dependencies, our dual-head Trans4Trans model is entirely composed of transformers, as shown in Fig. 2(a), while the single-head has only one decoder. The four-stage encoder is borrowed from PVT [63]. Different to PVT-based Trans2Seg [70] adopting CNN-decoder, both encoder and decoder of Trans4Trans

are symmetrically constructed by transformers for maintaining consistency in both feature extraction and feature parsing stages. Furthermore, different from CNN-based models [46, 53, 78, 79] learning the inductive bias, the transformer-based decoder is supposed to be more robust to parse unseen data captured in the wild. Yet, training a transformer model requires a large-scale dataset [17]. In order to solve the data-hunger problem and correct the misidentified walkable area through transparent objects segmentation, we designed a double-head model. Through the joint training of multiple datasets, it brings greater data diversity for learning a robust transformer-based model.

To construct such a lightweight decoder, we propose a *Transformer Parsing Module (TPM)* illustrated on Fig. 2(c). Each TPM contains only one single transformer-based layer, thus it only demands little computing resources and is flexible to be deployed on our portable hardware system. More precisely, our decoder consists of four symmetrical stages as encoder. Each stage has a TPM module and contains similar structure. As shown in Fig. 2(a), the pyramid features  $\{F_1, F_2, F_3, F_4\}$  from encoder are parsed consistently by the specific TPM module. Between two stages, re-size and element-wise addition will be performed for pyramid feature fusion. For balancing the capacity and computational demands, the feature resolution of each TPM is set as  $\frac{H}{4} \times \frac{W}{4} \times C$ , for which the default channel is 64.

Benefiting from our proposed TPM, the amount of GFLOPs and parameters of this dual-head structure is largely reduced compared to deploying two separate models. Also importantly, diverse features can be learned from various datasets. Thereby, the dual-head model maintains lightweight and is robust in terms of preventing overfitting when testing in real-world scenarios. The decoder composed of our TPM module can be flexibly applied with various CNN- or transformer-based encoder structures as well. For multi-task learning, mounting decoder heads robustifies the feature learned via the shared encoder, and the entire model will not be computationally overburdened.

### 3.2. Portable System

Our entire portable system consists of two hardware components: a pair of smart vision glasses and a portable GPU, *e.g.*, NVIDIA AGX Xavier or a lightweight laptop.

The smart vision glasses integrate a RealSense R200 RGB-D sensor to capture RGB and depth images at the resolution of  $640 \times 480$  in real time, and a pair of bone-conducting earphones for generating acoustic feedback to visually impaired people. This is critical as visually impaired people often use the sounds from their surrounding environments for orientation and bone-conducting headphones will not block their ears when using the system. The integrated RealSense R200 sensor leverages a combination of active speckle projecting and passive stereo matching, and thereby it can work in both indoor and outdoor scenes. In texture-less indoor scenes, the projected infrared speckles will augment the environments, which are beneficial for stereo matching algorithms to yield dense depth estimation. In sunny outdoor scenes, whereas projected patterns would be overwhelmed by sunlight, the infrared components of natural light shine on the scene to form well-textured infrared image pairs, thus producing robust depth sensing. In our system, depth information is used to perform the obstacle avoidance function and can be used for prioritizing near-range objects over mid-/long-range objects.

### 3.3. System Algorithm

Our software components are the aforementioned dual-head Trans4Trans and a user interface as described in Algorithm 1. Starting from the input data and to guarantee the timely capture of the facing environment, the frame rate of RGB-D stream is set to 60. Once the system starts, it repeats image segmentation every  $n$  seconds. According to our experiments, the time interval setting as 2 seconds can effectively prevent cognitive overload, especially in cases of complex scenes containing a large number of objects. Still, it is adjustable depending on the demands of users, *e.g.*, a short interval for more feedback to explore unknown space.

**Obstacle avoidance.** When moving in a relatively restricted indoor space, the building materials or densely-arranged objects will impede the flexibility of merely using white cane as the aid tool for avoiding obstacles. In order to tackle the collision issue and balance indoor and outdoor scenarios, our system presets the highest priority for obstacle avoidance. In other words, if the average value of the depth information is smaller than the preset distance threshold  $\theta_{obstacle}$ , the user will be immediately notified in the form of *vibration*. To minimize the uncertainty of vibrations and the cognitive load, only one single default threshold is set to 1 meter, instead of setting various vibration frequencies for different distances. Another purpose is to preclude the chaotic and low-confidence segmentation from the less-textured images when users walk too close and face to the

---

#### Algorithm 1: Assistive system

---

**Data:** RGB-D as  $X \in \mathcal{R}^{H \times W \times 3}$  and  $Y \in \mathcal{R}^{H \times W}$ .  
**Result:** General segmentation  $G \in \mathcal{R}^{H \times W \times 13}$ ;  
 Transparency  $T \in \mathcal{R}^{H \times W \times 11}$ ;

- 1 initialize walkable rate:  $R_l, R_f, R_r$ , parameters:  
 $\theta_{obstacle}, \theta_{trans}, \theta_{walkable}$ ;
- 2 **while** system start and each  $n$  seconds **do**
- 3   RGB-D update and Trans4Trans segmentation:
- 4    $G_{path} \in \mathcal{R}^{H \times W}, G_{object} \in \mathcal{R}^{H \times W \times 12}$ ;
- 5    $T_{stuff} \in \mathcal{R}^{H \times W \times 3}, T_{thing} \in \mathcal{R}^{H \times W \times 8}$ ;
- 6   partition  $\{R_l, R_f, R_r\} \leftarrow G_{path}$ ;
- 7   **if**  $\bar{Y} < \theta_{obstacle}$  **then**
- 8     | *vibration* as obstacle warning;
- 9   **else if**  $\max\{\bar{T}_i\} \in T_{stuff} > \theta_{trans}$  **then**
- 10    | *speech*  $\leftarrow \operatorname{argmax}\{\bar{T}_i\} \in T_{stuff}$ ;
- 11   **else if**  $\max\{R_l, R_f, R_r\} > \theta_{walkable}$  **then**
- 12    | *speech*  $\leftarrow \operatorname{argmax}\{R_l, R_f, R_r\} \in \{left, forward, right\}$ ;
- 13   **else**
- 14    | *speech*  $\leftarrow \operatorname{nearest}\{T_{thing}, G_{object}\}$ ;
- 15   **end**
- 16 **end**

---

object surface, such as images from white wall or doors.

**(Transparent) object segmentation.** After receiving the RGB image  $X \in \mathcal{R}^{H \times W \times 3}$ , our efficient Trans4Trans model outputs two segmentation predictions, which are general object segmentation  $G \in \mathcal{R}^{H \times W \times 13}$  and transparent object segmentation  $T \in \mathcal{R}^{H \times W \times 11}$ , respectively. The general object segmentation is divided into  $G_{path}$  for *walkable path* and  $G_{object}$  for other *objects*. Afterwards, the walkable mask is further partitioned into three regions as  $\{left, forward, right\}$  directions for orientation. In order to correct the wrongly-segmented walkable area by the high-confidence transparency perception, the transparent object segmentation is divided into two disjoint sets as:  $T_{stuff} \in \mathcal{R}^{H \times W \times 3}$  with  $\{window, glass\ door, glass\ wall\}$ , and  $T_{things} \in \mathcal{R}^{H \times W \times 8}$  with  $\{shelf, jar/tank, freezer, eye-glass, cup, bowl, bottle, box\}$ .

**Walkable path detection.** After achieving object segmentation, the local ratio of walkable area  $G_{path}$ , *e.g.*, *floor* category from Stanford2D3D, is further horizontally divided into three different directions as  $\{R_l, R_f, R_r\} \leftarrow G_{path}$ . Then, an intuitive and effective strategy is to prompt the direction that has the largest walkable area, only when its local ratio is greater than the preset threshold  $\theta_{walkable}$  for safety. According to our test, this orientation approach guarantees anti-veering in a straight path outdoors and indoors. Furthermore, it can also accurately predict the best instantaneous turning direction during walking at an intersection, so as to constantly yield a safe direction suggestion.



Network	Encoder	Decoder	GFLOPs	MParams	Stanford2D3D	Trans10K-v2
Trans2Seg-T	R18 [24]	Transformer [70]	16.96	17.87	42.07	64.20
Trans2Seg-T	R34 [24]	Transformer [70]	30.26	27.98	42.91	68.84
Trans2Seg-S	R50 [24]	Transformer [70]	40.98	30.53	43.83	69.20
PVT-T	PVT-T [63]	Transformer [70]	10.16	13.11	41.00	64.60
PVT-S	PVT-S [63]	Transformer [70]	19.58	24.36	41.89	68.47
PVT-M	PVT-M [63]	Transformer [70]	49.00	56.20	42.49	72.10
Trans4Trans-T	PVT-T [63]	TPM / Single-head	10.45	12.71	41.28(+0.28)	68.63(+4.03)
Trans4Trans-S	PVT-S [63]	TPM / Single-head	19.92	23.95	44.47(+3.04)	74.15(+5.68)
Trans4Trans-M	PVT-M [63]	TPM / Single-head	34.38	43.65	45.73(+3.24)	75.14(+3.04)
Trans4Trans-T	PVT-T [63]	TPM / Dual-head	11.22	13.10	40.44(-0.56)	69.84(+5.24)
Trans4Trans-S	PVT-S [63]	TPM / Dual-head	20.69	24.34	43.45(+1.56)	74.57(+6.10)
Trans4Trans-M	PVT-M [63]	TPM / Dual-head	35.17	44.04	45.15(+2.66)	74.98(+2.88)

Table 1: Comparison with state-of-the-art methods on Stanford2D3D [2] and Trans10K-v2 [70]. #MParams, #GFLOPs are calculated with the input size of  $512 \times 512$ .

## 4. Experiments

### 4.1. Datasets and settings

**Trans10K-v2** [70] contains 10,428 images with a  $835 \times 1113$  resolution and 11 object categories: *shelf, jar or tank, freezer, window, glass door, eyeglass, cup, wall, glass bow, water bottle, storage box*. It is divided as 5,000, 1,000 and 4,428 images for training, validation and testing.

**Stanford2D3D** [2] contains 70,496 images with a  $1080 \times 1080$  resolution and 13 object categories: *beam, board, bookcase, ceiling, chair, clutter, column, door, floor, sofa, table, wall, window*. We follow the fold-1 setting as [2] which splits Area 1, 2, 3, 4 and 6 as training set, Area 5a and 5b as validation and testing set.

**Implementation details.** We implement the model with PyTorch 1.8.0 and CUDA 11.2. Learning rate is initialized as  $1e-4$  and is scheduled by poly strategy [78] with power 0.9 in 100 epochs. The Adam [32] with epsilon  $1e-8$  and weight decay  $1e-4$  is used as the optimizer. Batch size is set as 4 on each of four 1080Ti GPUs. To maintain the shape of position embedding, the images are resized in the resolution of  $512 \times 512$  for all experiments. For a fair comparison with [70], some tricks such as OHEM, auxiliary or class-weighted loss are not applied in our experiments.

### 4.2. Segmentation accuracy

In this subsection, we present analysis for the experiments on different datasets and the results for the combination of CNN/Transformer to materialize varied encoder-decoder structures. Experimental results on computation complexity in GFLOPs and segmentation accuracy are presented and compared with state-of-the-art methods.

**Results.** As shown in Table 1, four main encoder-decoder structures are utilized for comparison. Unlike ResNet encoder-based Trans2Seg [70] and PVT [63], our Trans4Trans uses both Transformer-based encoder and decoder with a TPM design. It can be seen that single-head Trans4Trans-Medium has achieved the best performance in mIoU on both Stanford2D3D (45.73%) and Trans10K-v2 (75.14%), exceeding by more than 3% on the challenging transparent object segmentation benchmark *w.r.t* PVT-Medium. Meanwhile, it has clearly the smaller compu-

tation complexity in GFLOPs compared to PVT-M and ResNet50-based Trans2Seg. Trans4Trans-Tiny and -Small also achieve higher performances on Trans10K-v2 than the state-of-the-art structures. Dual-head Trans4Trans consistently improves the performance on Trans10K-v2 by incorporating more general knowledge when learning jointly with supervision from Stanford2D3D, which is more suitable for real-world navigational perception, as it highly prevents overfitting and reduces false positives of transparent obstacle warning observed in our field tests. Overall, these results verify the superiority and efficiency of Trans4Trans for transparent and general object segmentation.

**Combination of CNN/Transformer.** As shown in Table 2, varied combinations of CNN-/Transformer-based encoder and decoder are compared, where FCN [41] and OCNet [79] are composed of only CNN, whereas Trans2Seg is composed of CNN-based encoder and transformer-based decoder. The proposed Trans4Trans is a fully transformer-based encoder-decoder structure. It outperforms both these competitive architectures and PVT, another transformer-based encoder-decoder architecture. Yet, our Trans4Trans keeps smaller GFLOPs while being more accurate, demonstrating its suitability for transparent object segmentation.

Method	Trans. Enc.	CNN Enc.	Trans. Dec.	CNN Dec.	GFLOPs↓	mIoU (%)↑
FCN [41]		✓		✓	42.2	62.7
OCNet [79]		✓		✓	43.3	66.3
Trans2Seg [70]		✓	✓		40.9	69.2
PVT [63]	✓		✓		49.0	72.1
Trans4Trans (ours)	✓		✓		<b>34.3</b>	<b>75.1</b>

Table 2: Effect of CNN/Transformer combination. Models are evaluated on the Trans10K-v2 dataset.

**Comparison to state-of-the-art models.** Following [70], we compare with both accuracy- and efficiency-oriented semantic segmentation models as shown in Table 3. Compared with both CNNs and transformer-based methods like Trans2Seg [70], the superiority of Trans4Trans is further confirmed. Our Trans4Trans-M model outperforms the state-of-the-art method Trans2Seg by 2.99% in mIoU and 0.87% in ACC, while requiring much less GFLOPs. For category-wise accuracy, our Trans4Trans model achieves the state-of-the-art IoU on the classes *background, jar or tank, window, door, cup, wall, bottle* and *box*. These experimental results show the efficacy of transparent object segmentation of the proposed Trans4Trans architecture.

**Channel of TPM.** Since one of our critical designs lies in the TPM, we now analyze the effect of the numbers of embedding channels applied in the decoder of Trans4Trans, as shown in Table 4. It can be seen that performance increases as the number of channels increases until 256, and it drops at 512 where the decoder overfits the encoded feature (see Fig. 3) and the computation complexity becomes exceedingly large. For the response time-critical wearable system, we adopt 64 channels when deploying Trans4Trans due to its high efficiency and good performance.

Method	GFLOPs ↓	ACC ↑	mIoU ↑	Category IoU ↑											
				Background	Shelf	Jar/Tank	Freezer	Window	Door	Eyeglass	Cup	Wall	Bowl	Bottle	Box
FPENet [39]	0.76	70.31	10.14	74.97	0.01	0.00	0.02	2.11	2.83	0.00	16.84	24.81	0.00	0.04	0.00
ESPNetv2 [46]	0.83	73.03	12.27	78.98	0.00	0.00	0.00	0.00	6.17	0.00	30.65	37.03	0.00	0.00	0.00
ContextNet [52]	0.87	86.75	46.69	89.86	23.22	34.88	32.34	44.24	42.25	50.36	65.23	60.00	43.88	53.81	20.17
FastSCNN [53]	1.01	88.05	51.93	90.64	32.76	41.12	47.28	47.47	44.64	48.99	67.88	63.80	55.08	58.86	24.65
DFANet [34]	1.02	85.15	42.54	88.49	26.65	27.84	28.94	46.27	39.47	33.06	58.87	59.45	43.22	44.87	13.37
ENet [51]	2.09	71.67	8.50	79.74	0.00	0.00	0.00	0.00	0.00	0.00	0.00	22.25	0.00	0.00	0.00
DeepLabv3+MBv2 [56]	2.62	88.39	54.16	89.95	31.79	48.29	46.18	41.39	43.42	61.97	69.48	61.65	54.89	63.47	37.36
HRNet_w18 [62]	4.20	89.58	54.25	92.47	27.66	45.08	40.53	45.66	45.00	68.05	73.24	64.86	52.85	62.52	33.02
HardNet [10]	4.42	90.19	56.19	92.87	34.62	47.50	42.40	49.78	49.19	62.33	72.93	68.32	58.14	65.33	30.90
DABNet [33]	5.18	77.43	15.27	81.19	0.00	0.09	0.00	4.10	10.49	0.00	36.18	42.83	0.00	8.30	0.00
LEDNet [65]	6.23	86.07	46.40	88.59	28.13	36.72	32.45	43.77	38.55	41.51	64.19	60.05	42.40	53.12	27.29
Trans4Trans-T	10.45	93.23	68.63	94.44	48.39	61.89	61.86	61.14	54.83	73.60	83.03	75.20	74.69	75.26	59.19
ICNet [81]	10.64	78.23	23.39	83.29	2.96	4.91	9.33	19.24	15.35	24.11	44.54	41.49	7.58	27.47	3.80
BiSeNet [78]	19.91	89.13	58.40	90.12	39.54	53.71	50.90	46.95	44.68	64.32	72.86	63.57	61.38	67.88	44.85
Trans4Trans-S	19.92	94.57	74.15	95.60	<b>57.05</b>	71.18	<b>70.21</b>	63.95	61.25	81.67	87.34	78.52	<b>77.13</b>	81.00	64.88
DenseASPP [75]	36.20	90.86	63.01	91.39	42.41	60.93	64.75	48.97	51.40	65.72	75.64	67.93	67.03	70.26	49.64
DeepLabv3+ [12]	37.98	92.75	68.87	93.82	51.29	64.65	65.71	55.26	57.19	77.06	81.89	72.64	70.81	77.44	58.63
FCN [41]	42.23	91.65	62.75	93.62	38.84	56.05	58.76	46.91	50.74	82.56	78.71	68.78	57.87	73.66	46.54
OCNet [79]	43.31	92.03	66.31	93.12	41.47	63.54	60.05	54.10	51.01	79.57	81.95	69.40	68.44	78.41	54.65
RefineNet [35]	44.56	87.99	58.18	90.63	30.62	53.17	55.95	42.72	46.59	70.85	76.01	62.91	57.05	70.34	41.32
Trans2Seg [70]	49.03	94.14	72.15	95.35	53.43	67.82	64.20	59.64	60.56	<b>88.52</b>	86.67	75.99	73.98	82.43	57.17
TransLab [69]	61.31	92.67	69.00	93.90	54.36	64.48	65.14	54.58	57.72	79.85	81.61	72.82	69.63	77.50	56.43
DUNet [29]	123.69	90.67	59.01	93.07	34.20	50.95	54.96	43.19	45.05	79.80	76.07	65.29	54.33	68.57	42.64
U-Net [54]	124.55	81.90	29.23	86.34	8.76	15.18	19.02	27.13	24.73	17.26	53.40	47.36	11.97	37.79	1.77
DANet [19]	198.00	92.70	68.81	93.69	47.69	66.05	70.18	53.01	56.15	77.73	82.89	72.24	72.18	77.87	56.06
PSPNet [82]	187.03	92.47	68.23	93.62	50.33	64.24	70.19	51.51	55.27	79.27	81.93	71.95	68.91	77.13	54.43
Trans4Trans-M	34.38	<b>95.01</b>	<b>75.14</b>	<b>96.08</b>	55.81	<b>71.46</b>	69.25	<b>65.16</b>	<b>63.96</b>	83.84	<b>88.21</b>	<b>80.29</b>	76.33	<b>83.09</b>	<b>68.09</b>

Table 3: Computation complexity in GFLOPs and category-wise accuracy evaluation and comparison with state-of-the-art semantic segmentation methods on the Trans10K-v2 dataset [70].

Channel	MParams	GFLOPs	Acc (%)	mIoU (%)
64	10.45	12.71	93.46	67.89
128	12.46	14.02	93.49	68.88
256	20.41	17.82	93.58	69.51
512	51.50	33.82	92.37	63.33

Table 4: Effect of embedding channel in TPM. All tiny Trans4Trans are trained on Trans10K-v2 at  $512 \times 512$  on one GPU.




Figure 3: mIoU and pixel accuracy curves under different embedding channels in TPM.

### 4.3. Real-time performance

To calculate the inference speed of our different versions of dual-head Trans4Trans model, 300 samples from the Trans10K-v2 test set with a batch size of 1 and a resolution of  $512 \times 512$  are tested on three different GPUs, *i.e.*,

a mobile NVIDIA AGX Xavier in the MAXN mode, an NVIDIA GeForce MX350 from a lightweight laptop and an RTX 2070 from a workstation. As shown in Table 5, the computation costs of our tiny Trans4Trans model on three GPUs are considerably lower than the other two, meanwhile the performances of the three models on both datasets are suitable for our system. In real applications, the more timely response of the navigation system is beneficial for assisting users with a similar prediction accuracy on each frame. Hence, the tiny version is selected in our user study.

Network	NVIDIA Xavier (ms) ↓	MX350 (ms) ↓	RTX 2070 (ms) ↓
Trans4Trans-M	115.9(±1.1) / 202.8(±1.1)	186.1(±0.3) / 243.2(±0.3)	22.9(±0.3) / 36.6(±0.8)
Trans4Trans-S	95.3(±0.6) / 158.6(±1.8)	140.6(±0.3) / 188.4(±0.4)	17.1(±0.3) / 27.7(±0.5)
Trans4Trans-T	75.8(±0.7) / 122.7(±0.7)	101.5(±0.3) / 141.7(±1.6)	12.8(±0.5) / 20.3(±0.5)

Table 5: Inference time (ms/frame) of dual-head Trans4Trans is tested in half-/single-precision on various GPUs at  $512 \times 512$ .

### 4.4. Qualitative segmentation analysis

Fig. 4 visualizes qualitative comparisons between our tiny Trans4Trans and the previous state-of-the-art method Trans2Seg [70]. Fig. 4(a) shows failed recognition cases of both models, but our model can yield a clearly better




Figure 4: Qualitative analysis on Trans10K-v2 test set. (a) shows some negative predictions from both models. In (b), our Trans4Trans can correctly segment those cases failed by Trans2Seg. In (c) and (d), our results are more precise.

boundary. Fig. 4(b) shows that examples where our model predicts the correct label, whereas Trans2Seg is confused. In Fig. 4(c)(d), it can be seen that our model is not only effective for detecting navigation-related *glass door* and *glass window*, but can also predict more refined segmentation of small objects like *jar/tank* and *glass cup*.

We further perform field tests by navigating around the university campus and capturing real-world scenes with our smart vision glasses. The collected RGB-D images and corresponding predictions are shown in Fig. 5. The glass door in the first row captured at a moderate distance can be correctly identified, whereas in the other rows they are mis-classified as walkable paths by general object segmentation models. As it can be observed, transparent surfaces are often texture-less and the infrared patterns projected by the glasses will transmit the glass regions, and thereby the depth information are often sparse, noisy or even lost, which makes it challenging for 3D vision-based systems [1, 60, 73] to help avoid hazards even though the obstacles are close. In contrast, our Trans4Trans accurately and completely segments those transparent objects, meanwhile covers general objects, which is ideally suitable for safety-critical navigation assistance.

## 5. User Study

We conducted a qualitative study with 5 participants in order to assess the acceptance of our prototype and draw design conclusions [49, 50].

**Methodology.** The hardware used during the test consisted of the smart vision glasses and a backpack with a lightweight laptop and a battery pack inside. The system’s battery life under these conditions was approximately 4 hours.




Figure 5: Visualization of real-world scenes. From left to right are RGB and depth image, segmentation as *walkable path* by single-head model trained on Stanford2D3D, and as transparent objects (glass door or *glass wall*) corrected by our dual-head Trans4Trans model.

Participants tried the system inside 2 buildings, and the blind participant also on a 700m route outdoors - see Fig. 6. The study lasted about 2 hours. As Corona-protective measures, everyone wore FFP2 or surgery masks throughout the study and the prototype was disinfected several times.

After a short introduction, all participants agreed to participation and recordings and signed the data protection statement. The participants put on the system and walked around the rooms, thinking out loud [30]. The study was recorded with an action camera and voice recorder. At the end, demographics and NASA Raw Task Load Index (RTLX) [21] questionnaires were filled in.

**Participants.** Due to COVID-19 restrictions, only one of the participants (P1B) belonged to the target user group, being early blind. The other 4 participants were sighted (P2-P5). Age and gender of the participants were fairly balanced - see Table 6. When asked if they can see glass objects during the day, P1B said he can sometimes see closed windows,




Figure 6: Incidences of participants using the system for navigation outdoors and indoors.



due to the light-dark contrast. Windows that open inside the room, however, are very dangerous, as one can get serious head injuries (P1B). All sighted participants said they can see glass objects most of the time, but some objects, like bottles and glass cups (P2), glass doors (P3,P5), glass walls and windows (P5), can be challenging under certain light conditions, *e.g.*, backlighting.

Gender		Age Range				Hearing loss
Male	Female	18-25	26-35	36-45	46-55	No
3	2	1	2	1	1	5

Table 6: Aggregated demographics of participants.

**Cognitive load.** The raw task load index, averaged over all participants, was 14.7 with a standard deviation of 4.1. This score is enough to keep the user motivated, while not burdening too much [45]. The blind participant, P1B, had the second lowest score. According to the individual ratings (Fig. 7), effort and physical demand were slightly higher, while frustration was the lowest subscale. This might suggest that users enjoyed the experience of using our system, but a further reduction of hardware would be welcome.




Figure 7: Cognitive load (raw task load index, 0-100).

**User comments.** A thematic analysis [5] performed on the comments made by users yielded the following insights:

- All users found the system useful and were impressed by its functionality (P1B-P5) and smooth running (P5). Users also praised the fact that it works both indoor and outdoor (P2,P3), is easy to use and interact (P2), the several functions are well integrated and battery life is high (P5). P1B said: “*for the first time, I had the feeling that artificial intelligence can be useful [...] It’s just cool!*”.
- All 5 users liked the fact that the system recognizes so many object classes, including glass objects. Objects recommended to be included in future versions were: trash cans, city scouters (P1B) and constructions site fences (P5). Detection of some false positives was mentioned by P1B and P5.
- Users found the hardware light (P1B,P2,P3,P5), comfortable (P1B,P3) and good looking (P5). For a commercial system, however, the laptop should be replaced

by a belt (P1B) or smartphone (P5) and the glasses should connect wirelessly (P1B,P5).

- The direction of detected objects should be announced (P1B,P2,P5) as clock directions or stereo sound (P1B).
- Synthetic voice was very much appreciated by P1B, as “*it clearly stands out from background sounds*”.
- The 1 – 2s delay setting was mentioned by P1B,P4,P5. P1B said the delay is acceptable for object segmentation, as he only uses this function when he is unsure. P5 only mentioned the delay with respect to obstacle avoidance. P4 said the delay is fine.

**Augmented reality for partially sighted people.** Since transparent obstacles are often a threat for people with low vision in everyday navigation and even challenging for sighted people in some cases, we further test our Trans4Trans method with a HoloLens 2 device by capturing real-world data around our computer vision laboratory. As shown in Fig. 8, the transparent objects like *glass door*, *transparent wall*, and *glass window* can be completely and consistently segmented, and the colored segmentation mask can be easily overlaid and naturally projected onto the original RGB image captured by the glasses for rendering augmented reality or mixed reality. This field test demonstrates that the proposed Trans4Trans framework is not only helpful for assisting blind people, but can be potentially useful for partially sighted people.



Figure 8: Augmented results with a HoloLens 2 device.

## 6. Conclusion

We look into the perception of transparent object segmentation via Trans4Trans, an efficient transformer architecture established with both transformer-based encoder and decoder. With a novel Transformer Parsing Module (TPM) integrated in the dual-head, Trans4Trans precisely segments general and transparent objects. It attains state-of-the-art performances on Stanford2D3D and Trans10K-v2 datasets, meanwhile being swift and robust to support online navigational perception. The learned efficient vision transformer is ported in our wearable system with a pair of smart vision glasses designed to help visually impaired people travel and explore real-world scenes, where transparent objects are omnipresent. Extensive results from a user study and various field tests show that the proposed assistive system is reliable and with low cognitive load.



## References

- [1] Aitor Aladren, Gonzalo López-Nicolás, Luis Puig, and Josechu J. Guerrero. Navigation assistance for the visually impaired using RGB-D sensor with range expansion. *IEEE Systems Journal*, 2016. 1, 2, 7
- [2] Iro Armeni, Sasha Sax, Amir Roshan Zamir, and Silvio Savarese. Joint 2D-3D-semantic data for indoor scene understanding. *arXiv*, 2017. 2, 5
- [3] Jinqiang Bai, Shiguo Lian, Zhaoxiang Liu, Kai Wang, and Dijun Liu. Smart guiding glasses for visually impaired people in indoor environment. *IEEE Transactions on Consumer Electronics*, 2017. 2
- [4] Bruno Berenguel-Baeta, Manuel Guerrero-Viu, A. Nova, Jesus Bermudez-Cameo, Alejandro Pérez-Yus, and Josechu J. Guerrero. Floor extraction and door detection for visually impaired guidance. In *ICARCV*, 2020. 1
- [5] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 2006. 8
- [6] F. M. Butera. Glass architecture: is it sustainable. *Passive and Low Energy Cooling for the Built Environment*, 2005. 1
- [7] Yang Cao, Zhengqiang Zhang, Enze Xie, Qibin Hou, Kai Zhao, Xiangui Luo, and Jian Tuo. FakeMix augmentation improves transparent object detection. *arXiv*, 2021. 2
- [8] Zhengcai Cao, Xiaowen Xu, Biao Hu, and MengChu Zhou. Rapid detection of blind roads and crosswalks by using a lightweight semantic segmentation network. *IEEE Transactions on Intelligent Transportation Systems*, 2020. 2
- [9] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 3
- [10] Ping Chao, Chao-Yang Kao, Yu-Shan Ruan, Chien-Hsiang Huang, and Youn-Long Lin. HardNet: A low memory traffic network. In *ICCV*, 2019. 6
- [11] Hao Chen, Kaiwei Wang, and Kailun Yang. Improving RealSense by fusing color stereo vision and infrared stereo vision for the visually impaired. In *ICISS*, 2018. 2
- [12] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 6
- [13] Ruiqi Cheng, Kaiwei Wang, Kailun Yang, Ningbo Long, Jian Bai, and Dong Liu. Real-time pedestrian crossing lights detection algorithm for the visually impaired. *Multimedia Tools and Applications*, 2018. 2
- [14] Ruiqi Cheng, Kaiwei Wang, Kailun Yang, Ningbo Long, Weijian Hu, Hao Chen, Jian Bai, and Dong Liu. Crosswalk navigation for people with visual impairments on a wearable device. *Journal of Electronic Imaging*, 2017. 2
- [15] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting spatial attention design in vision transformers. *arXiv*, 2021. 3
- [16] Angela Constantinescu, Karin Müller, Monica Haurilet, Vanessa Petrausch, and Rainer Stiefelhausen. Bring the environment to life: A sonification module for people with visual impairments to improve situation awareness. In *ICMI*, 2020. 2
- [17] 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
- [18] Ping-Jung Duh, Yu-Cheng Sung, Liang-Yu Fan Chiang, Yung-Ju Chang, and Kuan-Wen Chen. V-Eye: A vision-based navigation system for the visually impaired. *IEEE Transactions on Multimedia*, 2021. 2
- [19] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, 2019. 3, 6
- [20] Ben Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. LeViT: a vision transformer in ConvNet's clothing for faster inference. In *ICCV*, 2021. 3
- [21] Sandra G. Hart. NASA-Task load index (NASA-TLX); 20 years later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 2006. 7
- [22] Hao He, Xiangtai Li, Guangliang Cheng, Jianping Shi, Yunhai Tong, Gaofeng Meng, Véronique Pinet, and Lubin Weng. Enhanced boundary learning for glass-like object segmentation. *arXiv*, 2021. 2
- [23] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 2
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [25] I-Hsuan Hsieh, Hsiao-Chu Cheng, Hao-Hsiang Ke, Hsiang-Chieh Chen, and Wen-June Wang. Outdoor walking guide for the visually-impaired people based on semantic segmentation and depth map. In *ICPAI*, 2020. 2
- [26] Yi-Zeng Hsieh, Shih-Syun Lin, and Fu-Xiong Xu. Development of a wearable guide device based on convolutional neural network for blind or visually impaired persons. *Multimedia Tools and Applications*, 2020. 2
- [27] Zhiming Huang, Kaiwei Wang, Kailun Yang, Ruiqi Cheng, and Jian Bai. Glass detection and recognition based on the fusion of ultrasonic sensor and RGB-D sensor for the visually impaired. In *SPIE*, 2018. 2
- [28] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. CCNet: Criss-cross attention for semantic segmentation. In *ICCV*, 2019. 3
- [29] Qiangguo Jin, Zhaopeng Meng, Tuan D. Pham, Qi Chen, Leyi Wei, and Ran Su. DUNet: A deformable network for retinal vessel segmentation. *Knowledge-Based Systems*, 2019. 6
- [30] Christopher J. Johnstone, Nicole A. Bottsford-Miller, and Sandra J. Thompson. Using the think aloud method (cognitive labs) to evaluate test design for students with disabilities and english language learners. Technical report, National Center on Educational Outcomes, University of Minnesota, 2006. 7
- [31] Agastya Kalra, Vage Taamazyan, Supreeth Krishna Rao, Kartik Venkataraman, Ramesh Raskar, and Achuta Kadambi. Deep polarization cues for transparent object segmentation. In *CVPR*, 2020. 2

- [32] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [33] Gen Li, Inyoung Yun, Jonghyun Kim, and Joongkyu Kim. DABNet: Depth-wise asymmetric bottleneck for real-time semantic segmentation. In *BMVC*, 2019. 6
- [34] Hanchao Li, Pengfei Xiong, Haoqiang Fan, and Jian Sun. DFANet: Deep feature aggregation for real-time semantic segmentation. In *CVPR*, 2019. 6
- [35] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian D. Reid. RefineNet: Multi-path refinement networks for high-resolution semantic segmentation. In *CVPR*, 2017. 6
- [36] Jiaying Lin, Zebang He, and Rynson W. H. Lau. Rich context aggregation with reflection prior for glass surface detection. In *CVPR*, 2021. 2
- [37] Shufei Lin, Kaiwei Wang, Kailun Yang, and Ruiqi Cheng. KrNet: A kinetic real-time convolutional neural network for navigational assistance. In *ICCHP*, 2018. 2
- [38] Yimin Lin, Kai Wang, Wanxin Yi, and Shiguo Lian. Deep learning based wearable assistive system for visually impaired people. In *ICCVW*, 2019. 1, 2
- [39] Mengyu Liu and Hujun Yin. Feature pyramid encoding network for real-time semantic segmentation. *arXiv*, 2019. 6
- [40] 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. 3
- [41] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 5, 6
- [42] Ningbo Long, Kaiwei Wang, Ruiqi Cheng, Weijian Hu, and Kailun Yang. Unifying obstacle detection, recognition, and fusion based on millimeter wave radar and RGB-depth sensors for the visually impaired. *Review of Scientific Instruments*, 2019. 2
- [43] Wei Mao, Jiaming Zhang, Kailun Yang, and Rainer Stiefel-hagen. Panoptic lntention network: Towards efficient navigational perception for the visually impaired. In *RCAR*, 2021. 2
- [44] Magdalena Maringer, Nico Hauck, and Ardeshtir Mahdavi. Suitability evaluation of visual indicators on glass walls and doors for visually impaired people. *Applied Mechanics and Materials*, 2019. 1
- [45] Manuel Martinez, Kailun Yang, Angela Constantinescu, and Rainer Stiefel-hagen. Helping the blind to get through COVID-19: Social distancing assistant using real-time semantic segmentation on RGB-D video. *Sensors*, 2020. 8
- [46] Sachin Mehta, Mohammad Rastegari, Linda Shapiro, and Hannaneh Hajishirzi. ESPNetV2: A light-weight, power efficient, and general purpose convolutional neural network. In *CVPR*, 2019. 3, 6
- [47] Haiyang Mei, Xin Yang, Yang Wang, Yuanyuan Liu, Shengfeng He, Qiang Zhang, Xiaopeng Wei, and Rynson W. H. Lau. Don't hit me! Glass detection in real-world scenes. In *CVPR*, 2020. 1
- [48] Ondrej Miksik, Vibhav Vineet, Morten Lidegaard, Ram Prasaath, Matthias Nießner, Stuart Golodetz, Stephen L. Hicks, Patrick Pérez, Shahram Izadi, and Philip H. S. Torr. The semantic paintbrush: Interactive 3D mapping and recognition in large outdoor spaces. In *CHI*, 2015. 2
- [49] Jakob Nielsen. Estimating the number of subjects needed for a thinking aloud test. *International Journal of Human-Computer Studies*, 1994. 7
- [50] Jakob Nielsen. How many test users in a usability study? <https://www.nngroup.com/articles/how-many-test-users/>, 2012. Accessed: 2021-06-30. 7
- [51] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. ENet: A deep neural network architecture for real-time semantic segmentation. *arXiv*, 2016. 6
- [52] Rudra P. K. Poudel, Ujwal Bonde, Stephan Liwicki, and Christopher Zach. ContextNet: Exploring context and detail for semantic segmentation in real-time. In *BMVC*, 2018. 6
- [53] Rudra P. K. Poudel, Stephan Liwicki, and Roberto Cipolla. Fast-SCNN: Fast semantic segmentation network. In *BMVC*, 2019. 3, 6
- [54] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 6
- [55] Manaswi Saha, Alexander J. Fiannaca, Melanie Kneisel, Edward Cutrell, and Meredith Ringel Morris. Closing the gap: Designing for the last-few-meters wayfinding problem for people with visual impairments. In *ASSETS*, 2019. 1
- [56] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 6
- [57] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *ICCV*, 2021. 3
- [58] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. 3
- [59] 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
- [60] Hsueh-Cheng Wang, Robert K. Katzschmann, Santani Teng, Brandon Araki, Laura Giarre, and Daniela Rus. Enabling independent navigation for visually impaired people through a wearable vision-based feedback system. In *ICRA*, 2017. 1, 2, 7
- [61] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. MaX-DeepLab: End-to-End panoptic segmentation with mask transformers. In *CVPR*, 2021. 3
- [62] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 6
- [63] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, 2021. 3, 5

- [64] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 3
- [65] Yu Wang, Quan Zhou, Jia Liu, Jian Xiong, Guangwei Gao, Xiaofu Wu, and Longin Jan Latecki. LEDNet: A lightweight encoder-decoder network for real-time semantic segmentation. In *ICIP*, 2019. 6
- [66] Sitong Wu, Tianyi Wu, Fangjian Lin, Shengwei Tian, and Guodong Guo. Fully transformer networks for semantic image segmentation. *arXiv*, 2021. 3
- [67] Yu-Huan Wu, Yun Liu, Xin Zhan, and Ming-Ming Cheng. P2T: Pyramid pooling transformer for scene understanding. *arXiv*, 2021. 3
- [68] Kaite Xiang, Kailun Yang, and Kaiwei Wang. Polarization-driven semantic segmentation via efficient attention-bridged fusion. *Optics Express*, 2021. 2
- [69] Enze Xie, Wenjia Wang, Wenhai Wang, Mingyu Ding, Chunhua Shen, and Ping Luo. Segmenting transparent objects in the wild. In *ECCV*, 2020. 2, 6
- [70] Enze Xie, Wenjia Wang, Wenhai Wang, Peize Sun, Hang Xu, Ding Liang, and Ping Luo. Segmenting transparent object in the wild with transformer. In *IJCAI*, 2021. 2, 3, 5, 6
- [71] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. SegFormer: Simple and efficient design for semantic segmentation with transformers. *arXiv*, 2021. 3
- [72] Kailun Yang, Kaiwei Wang, Luis Miguel Bergasa, Eduardo Romera, Weijian Hu, Dongming Sun, Junwei Sun, Ruiqi Cheng, Tianxue Chen, and Elena López Guillén. Unifying terrain awareness for the visually impaired through real-time semantic segmentation. *Sensors*, 2018. 1, 2
- [73] Kailun Yang, Kaiwei Wang, Xiangdong Zhao, Ruiqi Cheng, Jian Bai, Yongying Yang, and Dong Liu. IR stereo RealSense: Decreasing minimum range of navigational assistance for visually impaired individuals. *Journal of Ambient Intelligence and Smart Environments*, 2017. 1, 7
- [74] Kailun Yang, Jiaming Zhang, Simon Reiß, Xinxin Hu, and Rainer Stiefelhagen. Capturing omni-range context for omnidirectional segmentation. In *CVPR*, 2021. 3
- [75] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. DenseASPP for semantic segmentation in street scenes. In *CVPR*, 2018. 6
- [76] Minghao Yin, Zhuliang Yao, Yue Cao, Xiu Li, Zheng Zhang, Stephen Lin, and Han Hu. Disentangled non-local neural networks. In *ECCV*, 2020. 3
- [77] Ervin Yohannes, Timothy K. Shih, and Chih-Yang Lin. Content-aware video analysis to guide visually impaired walking on the street. In *IVIC*, 2019. 2
- [78] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. BiSeNet: Bilateral segmentation network for real-time semantic segmentation. In *ECCV*, 2018. 3, 5, 6
- [79] Yuhui Yuan, Lang Huang, Jianyuan Guo, Chao Zhang, Xilin Chen, and Jingdong Wang. OCNet: Object context for semantic segmentation. *International Journal of Computer Vision*, 2021. 3, 5, 6
- [80] Qinglong Zhang and Yubin Yang. ResT: An efficient transformer for visual recognition. *arXiv*, 2021. 3
- [81] Hengshuang Zhao, Xiaojuan Qi, Xiaoyong Shen, Jianping Shi, and Jiaya Jia. ICNet for real-time semantic segmentation on high-resolution images. In *ECCV*, 2018. 6
- [82] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. 6
- [83] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip H. S. Torr, and Li Zhang. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *CVPR*, 2021. 3
- [84] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable DETR: Deformable transformers for end-to-end object detection. In *ICLR*, 2021. 3
- [85] Zhen Zhu, Mengde Xu, Song Bai, Tengpeng Huang, and Xiang Bai. Asymmetric non-local neural networks for semantic segmentation. In *ICCV*, 2019. 3