

# Open Scene Understanding: Grounded Situation Recognition Meets Segment Anything for Helping People with Visual Impairments

Ruiping Liu<sup>1</sup>, Jiaming Zhang<sup>1,\*</sup>, Kunyu Peng<sup>1</sup>, Junwei Zheng<sup>1</sup>, Ke Cao<sup>1</sup>, Yufan Chen<sup>1</sup>,  
Kailun Yang<sup>2</sup>, Rainer Stiefelhagen<sup>1</sup>

<sup>1</sup>Karlsruhe Institute of Technology, <sup>2</sup>Hunan University

## Abstract

*Grounded Situation Recognition (GSR) is capable of recognizing and interpreting visual scenes in a contextually intuitive way, yielding salient activities (verbs) and the involved entities (roles) depicted in images. In this work, we focus on the application of GSR in assisting people with visual impairments (PVI). However, precise localization information of detected objects is often required to navigate their surroundings confidently and make informed decisions. For the first time, we propose an Open Scene Understanding (OpenSU) system that aims to generate pixel-wise dense segmentation masks of involved entities instead of bounding boxes. Specifically, we build our OpenSU system on top of GSR by additionally adopting an efficient Segment Anything Model (SAM). Furthermore, to enhance the feature extraction and interaction between the encoder-decoder structure, we construct our OpenSU system using a solid pure transformer backbone to improve the performance of GSR. In order to accelerate the convergence, we replace all the activation functions within the GSR decoders with GELU, thereby reducing the training duration. In quantitative analysis, our model achieves state-of-the-art performance on the SWiG dataset. Moreover, through field testing on dedicated assistive technology datasets and application demonstrations, the proposed OpenSU system can be used to enhance scene understanding and facilitate the independent mobility of people with visual impairments. Our code will be available at [OpenSU](#).*

## 1. Introduction

Scene understanding is a fundamental computer vision technology that interprets visual scenes and is widely used in various scenarios, such as autonomous driving [37], robotics [74], assistive technology [71], etc. In this work, we mainly explore scene understanding for helping People with Visual Impairments (PVI). PVI often have difficulty

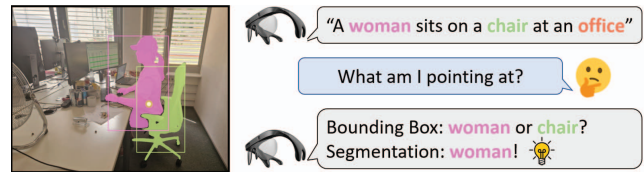


Figure 1: **OpenSU** generates image caption and segmentation maps of salient objects, while GSR outputs overlapping bounding boxes. The yellow point in the overlapping part of the bounding boxes represents the region of interest.

interpreting their surroundings correctly due to the lack of visual cues. Understanding a scene through tactile exploration proves insufficient and potentially hazardous. Often people need to recognize the whole scene at first glance, then gaze at each object, sort out their relationships, and react to the scene [25]. An assistive system providing scene descriptions and object locations could be very useful for visually impaired people to get a better understanding of the environment and thus lead towards more autonomy in many daily situations. Existing assistance systems [71, 74] focus on parsing the given image into pixel-level predictions with class names. However, only knowing the class name of objects – without having any description between objects – limits their ability to convey rich and comprehensive information about the scene. Grounded Situation Recognition (GSR) [48] is a verb-oriented task, which can predict the complete information of the scene, e.g., activity and entity information. Nevertheless, GSR models have limitations in object localization through bounding boxes, which lack fine-grained positional information and may result in ambiguous interaction, such as overlapping bounding boxes.

Recently, the advent of foundation models such as ChatGPT, GPT-3 [5] and Segment Anything Model (SAM) [26] has made significant contributions to the advancement of Artificial General Intelligence (AGI). However, there is still a lack of exploration regarding the application of these models in downstream tasks, particularly in the realm of assistive technologies for PVI. The research question arises: *how are foundation models applicable to assistance systems for helping people with visual impairments?*

\*Corresponding author (e-mail: [jiaming.zhang@kit.edu](mailto:jiaming.zhang@kit.edu)).

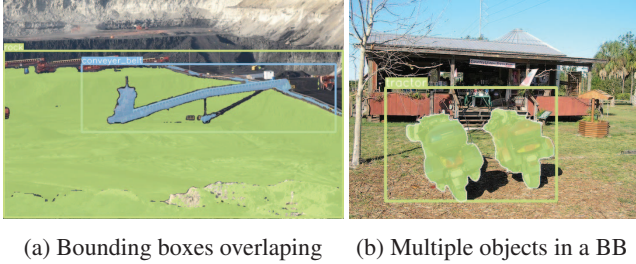


Figure 2: **Challenging cases** of using bounding boxes (BBs) for object localization.

Based on the aforementioned viewpoints, we introduce **Open Scene Understanding (OpenSU)** system (in Fig. 1), which consists of Grounded Situation Recognition (GSR) [48] and Segment Anything (SAM) [26]. Our OpenSU system can generate the scene description message as GSR while additionally producing disjoint segmentation masks of the objects to alleviate the confusion of object localization. Specifically, GSR outputs structured semantic descriptions of images, including the main activity, entities of the predefined roles (e.g., agent, tool, place), and the bounding-box groundings of entities. FrameNet [4] is the lexical-semantic resource that maps every verb to a set of semantic roles. Additionally, a preliminary frame is provided for each verb to facilitate the integration of roles, allowing for the subsequent insertion of predicted nouns into those roles. Besides, to achieve a wide range of open-set scene description capabilities, SWiG [48] dataset with over 500 verb and 10K entity classes is applied to train the GSR model. Compared to image captioning, GSR only considers the significant words, while image captioning also considers the word order and nonessential words, like the articles. Therefore, GSR yields more efficient descriptions and does not face the evaluation challenges [2, 59] that image captioning encounters.

As mentioned before, ambiguous interaction occurs when people attempt to localize the objects with bounding boxes generated by GSR models, as shown in Fig. 2. The positional relationship between objects can lead to overlapping bounding boxes (in Fig. 2(a)), or a single bounding box may encompass multiple similar objects (in Fig 2(b)). These challenging cases may cause confusion about object localization for PVI. Furthermore, the coarse bounding boxes to the object outlines result in the inclusion of background elements and an incomplete representation of the objects. In the depicted situations, the system suffers from confusion to extract the semantic information from the bounding box annotated scenes. To eliminate the localization limitation, SAM takes the bounding boxes as prompt and produces the accurate segmentation masks of the instances. Fig. 3 showcases the visual and linguistic outputs of SAM, GSR, and OpenSU. SAM generates precise masks without semantic information or text messages. GSR




Figure 3: **Comparison among scene understanding methods**, including Segment Anything Model (SAM) [26], Grounded Situation Recognition (GSR) [13], and our OpenSU. Our system provides the description: “A *man* rides the *motorcycle* at a *road*”, and semantic masks for more precise localization than overlapping bounding boxes.

predicts the activities, related objects, and bounding boxes. OpenSU yields both an image description and segmentation masks, enhancing the overall scene understanding.

In our approach, first, the description based on GSR and the segmentation masks of salient objects is presented by our OpenSU system. The user is then asked to specify a region of interest to get further information. To this end, we propose various potential region indication methods. Subsequently, the system returns the relevant semantic information corresponding to the defined region, so as to assist the user to localize the object of interest.

About the experiment, we validate the performance of our GSR model on the SWiG dataset [48], demonstrating a notable improvement of +3.47% over the previous state-of-the-art GSRFormer model [12]. The vanilla SAM suffers from large model sizes and latency. We explore different SAM variations [26, 70] in order to generate segmentation maps efficiently. Compared to the counterpart Grounded-SAM<sup>1</sup>, a stack of foundation models, our system realizes real-time Open Scene Understanding, having 0.7s vs. 74.04s processing time per image of the baseline.

To summarize, we present the following contributions:

- We design an *Open Scene Understanding (OpenSU)* system, allowing people with visual impairments to perceive the entirety of the scene and retrieve object information in a specified direction.
- To achieve real-time OpenSU, we propose the state-of-the-art GSR model which combines the Swin transformer and CoFormer architecture, and uses GELU to shorten the convergence duration.
- We for the first time, adopt and test two variations of the Segment Anything Model (SAM) in assistive technology for helping people with visual impairments.
- The effectiveness and efficiency of the OpenSU system are verified by various experiments and field tests.

<sup>1</sup><https://github.com/IDEA-Research>

## 2. Related Work

### 2.1. Semantic Segmentation

Semantic segmentation aims to partition a given image into several regions according to the semantic meaning, which tends to classify each pixel into its corresponding semantic class. Diverse research fields, *e.g.*, autonomous driving [11, 19, 47, 54, 66, 67, 75], medical image analysis [31, 50, 65], and assistance for visually impaired people [18, 30, 36, 56, 71, 74, 77], rely heavily on semantic segmentation to provide a deep understanding of the scenario to achieve better assistance. Most of the existing semantic segmentation works [36, 66, 71, 75] are close-set, where the prediction space is limited by the class types provided in the training set. With the rapid development of vision-language model [17, 28], *e.g.*, CLIP [24, 51] and BLIP [29], and segmentation foundation models, *e.g.*, SAM [26], open-set semantic segmentation [33, 40, 46, 58, 68] has played an increasingly important role in the computer vision community, which benefits from the ability to classify the pixels even for the class type which is not involved during the training and thereby relies less on the expensive pixel-level annotation on the new classes. Unlike existing systems, we for the first time, introduce SAM-based scene understanding for helping visually impaired people, which largely enhances the practicability of vision-based assistance in real-world unconstrained everyday scenarios.

### 2.2. Scene Understanding

Scene understanding concentrates on comprehending the existing elements and their corresponding relationships within a given image, which results in the analysis of the status of different objects, the attributes of the agents, and their interactions. Scene understanding enables the deep learning model to interpret the visual data, which benefits diverse practical usages [3, 34, 41, 63, 76], *e.g.*, assistance for visually impaired [32, 36, 56, 71] and autonomous driving [8, 44, 53], where Grounded Situation Recognition (GSR) [13] is one of the most essential scene understanding tasks, which predicts the entities, the activities and the bounding boxes groundings for the corresponding entities. SituFormer [61] makes use of a coarse-to-fine verb model and a transformer-based Noun model to achieve GSR. Cho *et al.* [13] propose a collaborative glance-gaze transformer where the glance transformer is used to predict the main activity and the gaze transformer is designed for the estimation of entities. Cheng *et al.* [12] propose GSRFormer to harvest structural-semantic description of a given image. In this work, we for the first time propose a grounded situation segmentation system to achieve dense segmentation of the entities inside their bounding boxes. We leverage Swin as the feature extraction backbone and SAM as the segmentation mask generator to achieve a superior scene

understanding. The convergence is accelerated by using the GELU activation function in the whole framework.

### 2.3. Visual Assistance Systems

Visual assistance systems are used to provide essential environmental information to achieve the navigation of the visually impaired through wearable sensors [10, 27, 35] and environment perceiving sensors [18, 36, 56, 71]. Semantic segmentation is a well-established technique used to assist the visually impaired. Lin *et al.* [35] propose a learning-based wearable system to achieve the navigation for visually impaired. Apart from semantic segmentation techniques, V-Eye [18] makes use of a global localization method to pursue a better scene understanding, while the outdoor walking guide system [23] leverages depth information. Liu *et al.* [36] propose HIDA to make use of instance semantic segmentation techniques with LiDAR sensor. The assistance for the visually impaired is likely to be involved in different traffic situations. Tian *et al.* [56] and Li *et al.* [32] concentrate on handling the crosswalk situation, which involves the segmentation of objects and prediction of the traffic light status, while Zou *et al.* [77] focus on real-time passable area segmentation. Most of the above-mentioned works rely on Convolutional Neural Networks (CNN) as feature extractors. Zhang *et al.* [71] propose Trans4Trans to leverage transformers to serve visual assistance systems. Ma *et al.* [42] propose a robot system to achieve wayfinding tasks for the visually impaired. The work of “I am the follower, also the boss” [72] uses machine forms of a guiding robot and anatomy from different stages to achieve visually impaired assistance. Zheng *et al.* [74] focus on material recognition in wearable robotics. Apart from the existing works with user studies, there are still well-established works that only concentrate on segmentation without user study [6, 52, 57, 74] to showcase the feasibility of their concepts to assist visually impaired. Our work falls into the latter case.

## 3. Method: Open Scene Understanding

As shown in Fig. 4, our Open Scene Understanding (OpenSU) system consists of three parts: (1) the grounded situation recognition; (2) the segmentation map generation; (3) the information transformation. A rough image caption, including main activities and objects, and bounding boxes of the objects are produced by the grounded situation recognition model. Since the objects interact with each other, the overlapping bounding boxes can not classify the object in the region of interest properly. To address this, the Segment Anything Model (SAM) is adopted behind the GSR model to obtain the non-overlapping segmentation masks annotated with roles related to the activities and the object class. Finally, we propose several region indication methods to easily specify an object of interest by PVI.

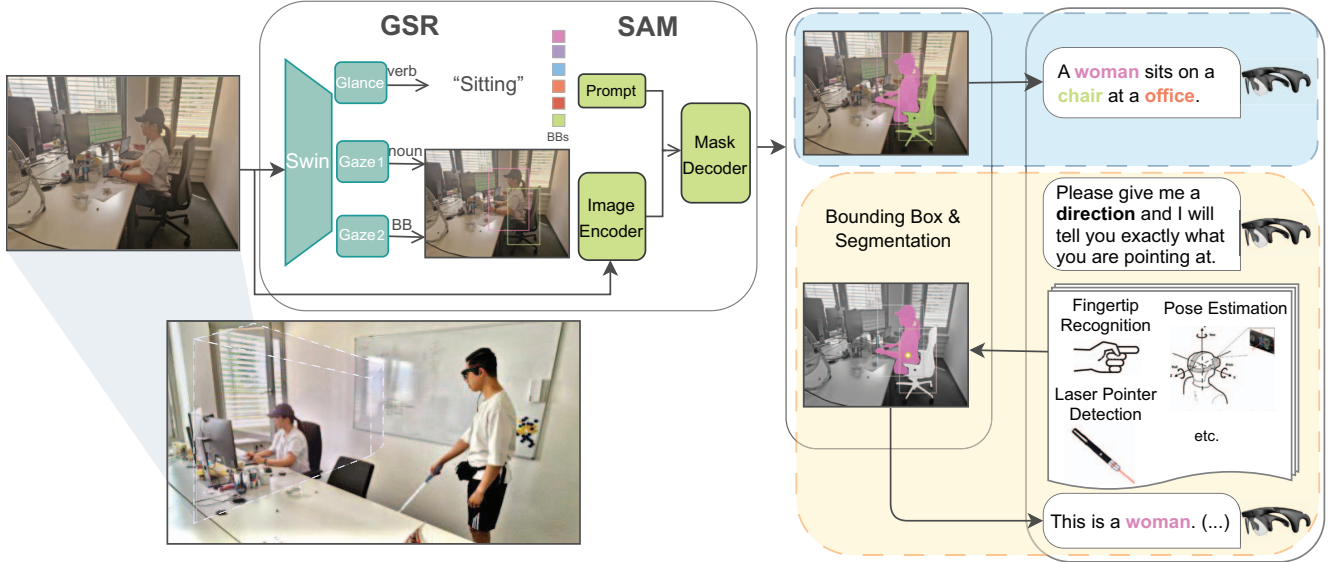


Figure 4: **The overview of Open Scene Understanding (OpenSU) system.** Grounded situation recognition captures the activity (e.g., *sitting*), nouns (e.g., *woman*, *chair*, *office*) related to the roles (Agent, Item, Place), and the bounding boxes of the objects. The caption template is “An Agent sits on an Item at a Place”, so the image caption of the verb *sitting* is “A woman sits on a chair at an office”. SAM uses bounding boxes as prompts to generate the segmentation masks. According to application interaction, potential region indication methods (e.g., fingertip, head pose, laser pointer) can be used to specify a region of interest, and the information can be reported to the user via bone-conducting earphones of wearable systems [71].

### 3.1. Grounded Situation Recognition

**Preliminary of CoFormer.** Our GSR Model is based on CoFormer [13], which is inspired by Kahneman’s cognitive theory [25]. The theory suggests that humans employ two comprehensive thinking systems to make decisions. In the case of the scene understanding task, humans rapidly glance at a scene to perceive the overall situation, then concentrate on specific details to analyze the objects involved and their relations. CoFormer builds upon the DETR [7] model, a significant milestone in end-to-end object detection, to simultaneously predict the activity, salient objects, and their corresponding bounding boxes. The image feature map extracted from CNN [21] backbone is processed with three transformer decoders. The first decoder, the Glance Transformer extracts holistic information and predicts the activity (verb) depicted in the scene. Based on the predicted verb, the set of roles is determined. The second decoder, Gaze-S1 captures the nouns and their relations using learnable role tokens as input. The third decoder, Gaze-S2 takes the aggregated image features and learnable tokens, predicted verb, and its associated roles as input. It extracts role features that are crucial for predicting grounded nouns.

**Swin Transformer for GSR.** CoFormer performed state-of-the-art on the grounded scene understanding dataset, SWiG [48]. However, since the multi-stream transformer decoder makes predictions only based on the output fea-

ture map from the backbone, we argue that the capability of the backbone in extracting high-level image features greatly impacts the overall performance of CoFormer. According to Dosovitskiy *et al.* [16], CNNs exhibit inferior performance in modeling global context due to the restricted receptive fields of the convolution operation, while vision transformers extract both global and local information through the self-attention mechanism, delivering the ability to capture long-range dependencies effectively. To include the long-range modeling capacity for GSR, we utilize the lightweight version of Swin Transformer [39] pre-trained on ImageNet-1k [15] to replace the CNN backbone. This design is non-trivial, since Swin Transformer shows superiority to extract the local representation through local window attention and the interaction among regions through shifted window attention. Performing self-attention computations within multiple windows is significantly more computationally efficient compared to applying global self-attention. Current GSR models [13] apply multiple decoders to extract both global (activity) and local (entities) information from the feature map obtained by the CNN backbone. In contrast, Swin Transformer which produces feature maps containing both global and local information is better suited for the GSR task than CNN. More results are presented in Sec. 4.3.

**GELU for GSR.** Additionally, we accelerate the convergence by replacing the ReLU [45] activation function with GELU [22], a non-linear approximation of ReLU with the

adaptive dropout feature, in the decoder transformers and the classifiers. Since ReLU can guarantee the gradient on the activation function to be constant 1 or 0 during back-propagation, it solves the gradient vanishing issue and reduces the huge amount of operations caused by updating the gradient with tiny values in practical engineering problems. It simplifies the complexity of computation and brings the effect of sparse activation, which is competitive in deep neural networks. ReLU [45] is formulated as:

$$\text{ReLU}(x) = \max(0, x). \quad (1)$$

However, as a result of sparsity in ReLU, several neurons with ReLU activation function are in a DEAD state and the parameters cannot be updated. Instead, this leads to a decrease in the representation capability, resulting in a certain loss of performance as well as training speed. Moreover, the ReLU outputs remain unbounded on the positive side, leading to potential issues like gradient explosion during training. In contrast, GELU looks at the non-linearity problem from a probabilistic view and delivers a continuous differentiable output. It incorporates the benefits of both stochastic regularization and ReLU. On the basis of one-sided saturation, it maintains a certain negative correction capability by retaining certain values in the partially negative region. GELU circumvents the issue of dead ReLU to a certain extent but retains the sparsity caused by dead ReLU. Additionally, the nonlinearity of GELU in the positive region enhances the fitting of complex nonlinear functions compared to ReLU, allowing the GSR model to capture more complex and nuanced relationships in the data. This leads to an advanced understanding of the grounded situation. To realize GELU, the input  $x$  is multiplied with the standard Gaussian cumulative distribution function  $\Phi(x)$ :

$$\begin{aligned} \text{GELU}(x) &= xP(X \leq x) = x\Phi(x) \\ &= x \cdot \frac{1}{2} \left[ 1 + \text{erf}\left(\frac{x}{\sqrt{2}}\right) \right]. \end{aligned} \quad (2)$$

GELU activation function is widely used in vision transformers [16, 39, 60], but its performance is barely discussed. In our experiments (Sec. 4.4), we showcase the effectiveness of GELU in our pure transformer architecture for grounded situation recognition and scene understanding.

### 3.2. Segmentation Map Generation

To gain the accurate boundaries of the salient objects, we leverage the original image as input and the bounding boxes outputted by the GSR as prompts for Segment Anything Models (SAMs). SAM serves as a promptable foundation model for image segmentation, combining the benefits of both semantic and instance segmentation. It provides the capability to generate hundreds of masks per image without prompt or generate specific masks using sparse prompts

such as points, boxes, or text, as well as dense prompts like masks. Typically, SAM utilizes bounding boxes as prompts for achieving instance segmentation, as masks generated by SAM with point prompts might represent minute components of an instance.

The solid but heavyweight image encoder ViT-H [16] with 611M parameters empowers Segment Anything to perform superior. Recent efforts [70, 73] dedicate to compress Segment Anything meanwhile remaining its robustness. MobileSAM [70] shares the same architecture with Segment Anything and employs TinyViT [62] with 5M parameters as an image encoder. The knowledge extracted from SAM is transferred into MobileSAM via knowledge distillation. MobileSAM improves not only the inference efficiency but also the training efficiency. It is trained on a single GPU with 1% of the original images for less than one day. Because of the limited computation resources of mobile assistance systems, it is crucial to perform efficient inference to enhance the user experience. Therefore, in this work, we spend a large effort to test and compare both vanilla SAM and MobileSAM in our OpenSU framework.

### 3.3. Region of Interest Indication

Multiple sensors, like an RGB-D camera and Inertial Measurement Units (IMUs), are installed on smart glasses [71]. They can be used to capture the indicator and infer the specified direction of interest. Here, we introduce several potential methods to indicate the region of interest. Please note that integrating different interaction methods in our system remains in our future work.

Firstly, we can capture the gesture through cameras and use gesture recognition algorithms [1, 9] to determine where the user is pointing. Gesture recognition is already involved in some VR/AR development kits, like Hololens. Furthermore, the work of Miksik *et al.* [43] on laser pointer detection coupled with optical see-through glasses (EPSON MOVERIO BT-200) serves as an inspiration for our direction indication. The utilization of IMUs enables the outside-in 6 DOF pose tracking and pose estimation in AR glasses, as analyzed by Frintepe *et al.* [20]. The system output is focused on identifying the central object within the scene.

## 4. Experiments

### 4.1. Dataset and Metrics

**Dataset.** The GSR dataset, SWiG [48] is extended from the situation recognition dataset imSitu [69] by adding 278,336 bounding boxes. Each sample of imSitu is labeled with a *verb*, which is associated with up to 6 *roles*. Three workers annotate the *entities* and their *bounding boxes* for the images using Amazon’s Mechanical Turk framework. The *bounding boxes* are described with format  $[x1, y1, x2, y2]$ . SWiG dataset contains 75K/25K/25K images for train-

Table 1. **Results (%) on the SWiG dataset**, including three settings and five metrics evaluated on the val and the test set.

Models	Top-1-Verb					Top-5-Verb					Ground-Truth-Verb			
	verb	value	val-all	grnd	grnd-all	verb	value	val-all	grnd	grnd-all	value	val-all	grnd	grnd-all
Methods for Grounded Situation Recognition (val set)														
ISL [49]	38.83	30.47	18.23	22.47	7.64	65.74	50.29	28.59	36.90	11.66	72.77	37.49	52.92	15.00
JSL [49]	39.60	31.18	18.85	25.03	10.16	67.71	52.06	29.73	41.25	15.07	73.53	38.32	57.50	19.29
GSRTR [14]	41.06	32.52	19.63	26.04	10.44	69.46	53.69	30.66	42.61	15.98	74.27	39.24	58.33	20.19
SituFormer [61]	44.32	35.35	22.10	29.17	13.33	71.01	55.85	33.38	45.78	19.77	76.08	42.15	61.82	24.65
CoFormer [13]	44.41	35.87	22.47	29.37	12.94	72.98	57.58	34.09	46.70	19.06	76.17	42.11	61.15	23.09
GSRFormer [12]	46.64	37.69	23.58	31.61	14.42	73.43	58.75	35.82	48.42	21.67	<b>78.76</b>	44.71	63.95	25.85
<b>OpenSU</b>	<b>49.96</b>	<b>41.12</b>	<b>26.69</b>	<b>34.69</b>	<b>16.27</b>	<b>77.77</b>	<b>62.81</b>	<b>39.23</b>	<b>52.56</b>	<b>23.50</b>	78.68	<b>46.66</b>	<b>65.33</b>	<b>27.66</b>
Methods for Grounded Situation Recognition (test set)														
ISL [49]	39.36	30.09	18.62	22.73	7.72	65.51	50.16	28.47	36.60	11.56	72.42	37.10	52.19	14.58
JSL [49]	39.94	31.44	18.87	24.86	9.66	67.60	51.88	29.39	40.60	14.72	73.21	37.82	56.57	18.45
GSRTR [14]	40.63	32.15	19.28	25.49	10.10	69.81	54.13	31.01	42.50	15.88	74.11	39.00	57.45	19.67
SituFormer [61]	44.20	35.24	21.86	29.22	13.41	71.21	55.75	33.27	46.00	20.10	75.85	42.13	61.89	24.89
CoFormer [13]	44.66	35.98	22.22	29.05	12.21	73.31	57.76	33.98	46.25	18.37	75.95	41.87	60.11	22.12
GSRFormer [12]	46.53	37.48	23.32	31.53	14.23	73.44	58.84	35.82	48.43	21.41	<b>78.81</b>	44.68	63.87	25.35
<b>OpenSU</b>	<b>50.10</b>	<b>41.20</b>	<b>26.56</b>	<b>34.27</b>	<b>15.70</b>	<b>77.91</b>	<b>62.95</b>	<b>39.00</b>	<b>51.90</b>	<b>22.74</b>	78.67	<b>46.31</b>	<b>64.22</b>	<b>26.57</b>

Table 2. **Efficiency analysis** of scene understanding methods and Segment Anything Model variants.

System	Runtime(s)	#Params(M)
Grounding DINO [38]+SAM+BLIP [28]	74.04	1310.66
OpenSU (SAM)	1.34	738.09
OpenSU (MobileSAM)	<b>0.70</b>	<b>107.13</b>

ing/validation/testing. The total numbers of *verb*, *roles* and *entities* classes are 504, 190 and 11,538. The activity-role pairs are derived from FrameNet [4] and the entity classes source from ImageNet [15].

**Metrics.** Five metrics are reported to evaluate the GSR models’ performance: (1) **verb** to measure verb prediction accuracy, (2) **value** to assess the accuracy of a predicted *noun* for a given *role*, (3) **value-all** to measure the correctness of all *nouns* within a frame, (4) **grounded-value** to determine if the *noun* is accurately predicted and its bounding box has an IoU over 0.5, and (5) **grounded-value-all** to measure the frequency that all the *noun-bounding box* pairs are correct. Because of the structured annotation, all metrics are highly related to the *verb* accuracy. And several activities could take place in a scene. Three settings are proposed to evaluate the predictions comprehensively. (1) **Top-1 Verb**: once the top-1 predicted *verb* is not the ground truth *verb* of the sample, all the *verb*, *noun*, and *bounding boxes* are regarded as wrong. (2) **Top-5 Verb**: the *verb* is correct, when the ground truth *verb* is contained in the *verbs* with top-5 probabilities. (3) **Ground-Truth Verb**: the ground truth *verb* is assumed to be given in this case.

## 4.2. Implementation Details

Our model is trained with a batch size of 4 for 40 epochs on four Tesla V100-SXM2-32GB GPUs. Each input image is padded to size  $700 \times 700$ . The transformer backbone Swin-Tiny [39] is pre-trained on ImageNet-1k [15]. We adopt AdamW optimizer with  $10^{-4}$  weight decay,  $\beta_1=0.9$




Figure 5: **Training curves** of the GSR model, CoFormer [13], and our OpenSU systems, considering both ReLU and GELU as the decoder activation functions.

and  $\beta_2=0.999$ . The learning rates are different for backbone and transformer decoders, *i.e.*,  $10^{-5}$  and  $10^{-4}$ . After epoch 30, the learning rates decrease with factor 10. We use the same loss functions as CoFormer [13] to align logits and targets, which are verb classification loss, noun classification loss, box existence loss, and box regression loss. To compare the inference time, 10 images with size  $1042 \times 1042$  go through our OpenSUs and other frameworks. The average runtime per image and the model parameters are calculated. The evaluation process is conducted on one Tesla V100-SXM2-32GB GPU.

## 4.3. Results on SWiG

In this experiment, we focus on comparing the performance of our proposed method with other state-of-the-art methods. The quantitative results are illustrated in Tab. 1, which indicates the proposed OpenSU outperforms other counterparts among all evaluation metrics except for the value metric in the Ground-Truth Verb setting. Compared with the second best method [12], OpenSU achieves a maximum 4.34% performance gain with the verb metric

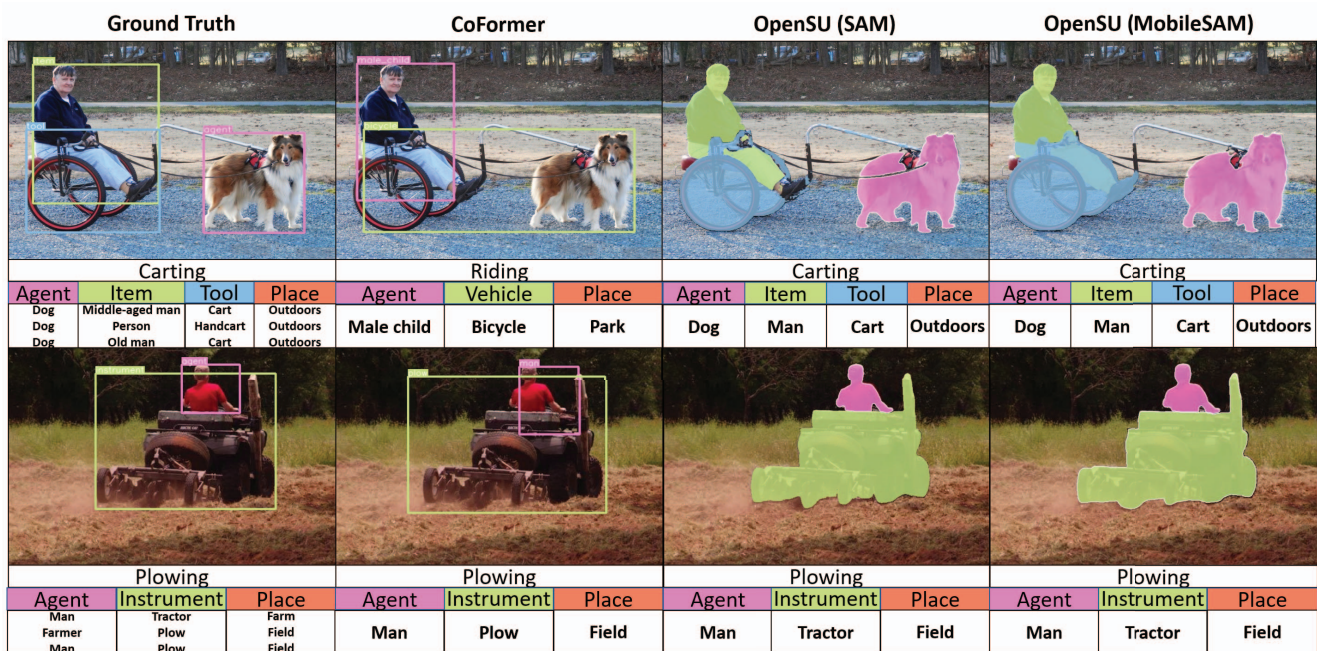


Figure 6: **Qualitative comparison** between GSR ground truth, predictions obtained from the GSR baseline CoFormer [13], and fine-grained predictions of two SAM-based variants of our OpenSU system.

in Top-5-Verb setting on the validation set, and minimum 1.39% boost with the grounded-value metric in Ground-Truth-Verb setting. Similar improvements can be found on the SWiG test set with maximum 4.47% and minimum 0.36% referring to the verb metric in Top-5-Verb and the grounded-value metric in the Ground-Truth-Verb setting, respectively. As for the value metric in the Ground-Truth-Verb setting, OpenSU achieves very close performance to GSRFormer [12], the most competitive counterpart listed in Tab. 1. It is worth noting that the performance gap between the validation and test set is very small, which showcases our OpenSU has an excellent generalization ability.

#### 4.4. Efficiency Analysis

We further conduct experiments to analyze the efficiency. Tab. 2 showcases the runtime in seconds and the number of parameters of all systems. It can be easily found that the MobileSAM-based OpenSU has the lowest number of parameters, which is less than one-tenth ( $-1203.53$  MParams) compared with the combination of Grounding DINO [38], SAM and BLIP [28]. Besides, OpenSU with MobileSAM can achieve the smallest runtime, *i.e.*,  $0.70s$ , reducing  $73.34s$  compared to the combination baseline. Furthermore, Fig. 5 presents the convergence speed of models using different activation functions. Our OpenSU with GELU converges faster than OpenSU with RELU and CoFormer [13] while keeping the highest verb accuracy under the Top-1-Verb setting during training. This evidence proves that the smoothness and continuity of GELU facilitate more stable and efficient optimization during training.

#### 4.5. Qualitative Results

Qualitative results of CoFormer [13] and our OpenSUs are presented in Fig. 6. For each image, there are three frames that originate from three annotators, each of which can be regarded as ground truth. In the first example, since CoFormer wrongly predicts the activity, the related information including roles, entity classes, and bounding boxes is incorrect as well. Our OpenSU systems with the state-of-the-art GSR model obtain the correct text and spatial information. Compared to OpenSU with MobileSAM, where the upper body and the legs are delineated as separate entities, OpenSU with original SAM can segment the details of the overlapping objects. The second row showcases another situation. All three models predict the correct verb and nouns. However, the positions of the bounding boxes predicted by CoFormer have a serious shift and the large area of overlap leads to confusion. Both variations of the proposed OpenSU output the corresponding segmentation masks instead of the bounding boxes resulting in more precise prediction and better understanding.

#### 4.6. Field Test

Apart from the evaluation on the SWiG dataset, the OpenSU system underwent a field test on two visually impaired assistance datasets, *i.e.*, Obstacle Dataset (OD) [55] and Walk On The Road (WOTR) [64]. The field test encompassed comprehensive daily scenes, allowing for a thorough evaluation of the system performance in assisting PVI. Fig. 7 visualizes the results of OpenSU, including ① con-

struction site, ② dining hall, ③ campus, ④ garage, ⑤ supermarket, ⑥ airport, ⑦ road, and ⑧ farm. Accurate segmentation of entities can be achieved even when the samples (① and ②) contain five interactive roles. The detection of the woman and garbage in ② and ⑤ is affected by the objects with the same class and same color respectively, leading to bounding boxes with incorrectly predicted sizes. Overlapping bounding boxes do not degrade segmentation performance, but inaccurate positioning and sizing can reduce accuracy. The absence of human agents presented in the samples ④, ⑥, and ⑧ signifies that GSR differs from Human-Object Interaction (HOI) and action recognition tasks. These tasks are generalized by introducing verbs, such as “glowing”, “taxiing”, and “pawing”, to specifically describe objects’ activities instead of people’s actions. We observe that the descriptions of roadblocks are incorrect in ① and ④. As part of future work, we suggest involving additional information about roadblocks in the dataset to improve accuracy. Based on the field test, the image captions provided in our analysis are considered rough, as they consist of predefined frames for each verb with entities filled in. However, despite their roughness, people can comprehend them due to the structured image description.

#### 4.7. Discussion

**(1) What can AGI promote assistive technology?** Artificial General Intelligence (AGI) has the potential to promote the development of assistive technology, including scene perception and understanding, scene reasoning, adaptive learning for specific needs and preferences of users, high-level decision-making in mobility and navigation, etc. These capabilities can be leveraged to develop more intelligent and adaptive assistive technologies, enabling them to better respond to the special needs of users, particularly for people with disabilities. However, the recent AGI technology is resource-intensive. Thus, making AGI-based assistance systems accessible and affordable to a wide range of users, poses a significant challenge. Therefore, improving efficiency is crucial to enhance the overall user experience.

**(2) How to leverage big vision models to help PVI?** By leveraging the advantages of big vision models, *i.e.*, Segment Anything Model (SAM), we enhance the object recognition and positioning of GSR by delivering fine-grained segmentation masks instead of ambiguous bounding boxes. This approach enables more precise interaction experiences when navigating and interacting with known or unknown environments. The proposed OpenSU system is the first attempt to combine SAM with GSR for helping PVI. Moving forward, the integration of SAM in various assistive applications will be further explored to unlock its full potential, such as customized object recognition, prompt-based object finding, path finding, social distancing, *etc.*




Figure 7: **Field test with the OpenSU system.** The samples are derived from datasets specifically proposed for assisting the visually impaired, *i.e.*, OD [55] and WOTR [64].

#### 5. Conclusion

People with Visual Impairments (PVI) often encounter challenges in understanding scenes in their daily lives. In this work, we propose an open scene understanding system (OpenSU) to assist PVI to obtain more comprehensive scene descriptions. Specifically, we utilize the Swin transformer as the backbone, enabling more accurate Grounded Situation Recognition (GSR). Besides, we for the first time, adopt the Segment Anything Model (SAM) to generate finer-grained masks for more precise object positioning compared to traditional bounding boxes of GSR. Through field testing, OpenSU showcases the effectiveness and the advantage of using SAM and GSR to perform open scene understanding in daily scenarios.



## References

- [1] Mohammad Mahmudul Alam, Mohammad Tariqul Islam, and S. M. Mahbubur Rahman. Unified learning approach for egocentric hand gesture recognition and fingertip detection. *PR*, 2022. 5
- [2] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. SPICE: Semantic propositional image caption evaluation. In *ECCV*, 2016. 2
- [3] Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. ScanQA: 3D question answering for spatial scene understanding. In *CVPR*, 2022. 3
- [4] Collin F. Baker, Charles J. Fillmore, and John B. Lowe. The berkeley framenet project. In *COLING*, 1998. 2, 6
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020. 1
- [6] Zhengcai Cao, Xiaowen Xu, Biao Hu, and MengChu Zhou. Rapid detection of blind roads and crosswalks by using a lightweight semantic segmentation network. *T-ITS*, 2021. 3
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 4
- [8] Po-Yi Chen, Alexander H. Liu, Yen-Cheng Liu, and Yu-Chiang Frank Wang. Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In *CVPR*, 2019. 3
- [9] Yung-Han Chen, Chi-Hsuan Huang, Sin-Wun Syu, Tien-Ying Kuo, and Po-Chyi Su. Egocentric-view fingertip detection for air writing based on convolutional neural networks. *Sensors*, 2021. 5
- [10] Zhuo Chen, Xiaoming Liu, Masaru Kojima, Qiang Huang, and Tatsuo Arai. A wearable navigation device for visually impaired people based on the real-time semantic visual SLAM system. *Sensors*, 2021. 3
- [11] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *NeurIPS*, 2021. 3
- [12] Zhi-Qi Cheng, Qi Dai, Siyao Li, Teruko Mitamura, and Alexander Hauptmann. GSRFormer: Grounded situation recognition transformer with alternate semantic attention refinement. In *MM*, 2022. 2, 3, 6, 7
- [13] Junhyeong Cho, Youngseok Yoon, and Suha Kwak. Collaborative transformers for grounded situation recognition. In *CVPR*, 2022. 2, 3, 4, 6, 7
- [14] Junhyeong Cho, Youngseok Yoon, Hyeonjun Lee, and Suha Kwak. Grounded situation recognition with transformers. *BMVC*, 2021. 6
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *ICCV*, 2009. 4, 6
- [16] 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. 4, 5
- [17] Yu Du, Fangyun Wei, Zihe Zhang, Miaoqing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *CVPR*, 2022. 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. *TMM*, 2021. 3
- [19] Mingyuan Fan, Shenqi Lai, Junshi Huang, Xiaoming Wei, Zhenhua Chai, Junfeng Luo, and Xiaolin Wei. Rethinking BiSeNet for real-time semantic segmentation. In *CVPR*, 2021. 3
- [20] Ahmet Firintepe, Alain Pagani, and Didier Stricker. A comparison of single and multi-view IR image-based ar glasses pose estimation approaches. In *VRM*, 2021. 5
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [22] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016. 4
- [23] 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. 3
- [24] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, 2021. 3
- [25] Daniel Kahneman. Maps of bounded rationality: Psychology for behavioral economics. *AER*, 2003. 1, 4
- [26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloé Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross B. Girshick. Segment anything. In *ICCV*, 2023. 1, 2, 3
- [27] Guoxin Li, Jiaqi Xu, Zhijun Li, Chao Chen, and Zhen Kan. Sensing and navigation of wearable assistance cognitive systems for the visually impaired. *TCDS*, 2023. 3
- [28] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 3, 6, 7
- [29] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 3
- [30] Ruiyu Li, Makarand Tapaswi, Renjie Liao, Jiaya Jia, Raquel Urtasun, and Sanja Fidler. Situation recognition with graph neural networks. In *ICCV*, 2017. 3
- [31] Shuailin Li, Chuyu Zhang, and Xuming He. Shape-aware semi-supervised 3D semantic segmentation for medical images. In *MICCAI*, 2020. 3
- [32] Xiang Li, Hanzhang Cui, John-Ross Rizzo, Edward Wong, and Yi Fang. Cross-Safe: A computer vision-based approach to make all intersection-related pedestrian signals accessible for the visually impaired. In *CVC*, 2020. 3

- [33] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted CLIP. In *CVPR*, 2023. 3
- [34] Liang Liao, Wenyi Chen, Jing Xiao, Zheng Wang, Chia-Wen Lin, and Shin'ichi Satoh. Unsupervised foggy scene understanding via self spatial-temporal label diffusion. *TIP*, 2022. 3
- [35] Yimin Lin, Kai Wang, Wanxin Yi, and Shiguo Lian. Deep learning based wearable assistive system for visually impaired people. In *ICCVW*, 2019. 3
- [36] Huayao Liu, Ruiping Liu, Kailun Yang, Jiaming Zhang, Kunyu Peng, and Rainer Stiefelwagen. HIDA: Towards holistic indoor understanding for the visually impaired via semantic instance segmentation with a wearable solid-state LiDAR sensor. In *ICCVW*, 2021. 3
- [37] Ruiping Liu, Kailun Yang, Huayao Liu, Jiaming Zhang, Kunyu Peng, and Rainer Stiefelwagen. TransKD: Transformer knowledge distillation for efficient semantic segmentation. *arXiv preprint arXiv:2202.13393*, 2022. 1
- [38] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding DINO: marrying DINO with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 6, 7
- [39] 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. 4, 5, 6
- [40] Chaofan Ma, Yuhuan Yang, Yanfeng Wang, Ya Zhang, and Weidi Xie. Open-vocabulary semantic segmentation with frozen vision-language models. *arXiv preprint arXiv:2210.15138*, 2022. 3
- [41] Xianzheng Ma, Zhixiang Wang, Yacheng Zhan, Yinqiang Zheng, Zheng Wang, Dengxin Dai, and Chia-Wen Lin. Both style and fog matter: Cumulative domain adaptation for semantic foggy scene understanding. In *CVPR*, 2022. 3
- [42] Yanan Ma, Qi Xu, Yue Wang, Jing Wu, Chengnian Long, and Yi-Bing Lin. EOS: An efficient obstacle segmentation for blind guiding. *FGCS*, 2023. 3
- [43] 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. 5
- [44] Weidong Min, Ruikang Liu, Daojing He, Qing Han, Qingting Wei, and Qi Wang. Traffic sign recognition based on semantic scene understanding and structural traffic sign location. *T-ITS*, 2022. 3
- [45] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010. 4, 5
- [46] Ian Nunes, Matheus B. Pereira, Hugo Oliveira, Jefersson A. dos Santos, and Marcus Poggi. Conditional reconstruction for open-set semantic segmentation. In *ICIP*, 2022. 3
- [47] Kunyu Peng, Juncong Fei, Kailun Yang, Alina Roitberg, Jiaming Zhang, Frank Bieder, Philipp Heidenreich, Christoph Stiller, and Rainer Stiefelwagen. MASS: Multi-attentional semantic segmentation of LiDAR data for dense top-view understanding. *T-ITS*, 2022. 3
- [48] Sarah Pratt, Mark Yatskar, Luca Weihs, Ali Farhadi, and Aniruddha Kembhavi. Grounded situation recognition. In *IV*, 2020. 1, 2, 4, 5
- [49] Sarah Pratt, Mark Yatskar, Luca Weihs, Ali Farhadi, and Aniruddha Kembhavi. Grounded situation recognition. In *ECCV*, 2020. 6
- [50] Yuhang Qiu, Chia Shuo Chang, Jiun Lin Yan, Li Ko, and Tian Sheuan Chang. Semantic segmentation of intracranial hemorrhages in head ct scans. In *ICSESS*, 2019. 3
- [51] Alec Radford, Jong Wook Kim, Chris Hallacy, A. Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 3
- [52] Jarek Reynolds, Chandra Kanth Nagesh, and Danna Gurari. Salient object detection for images taken by people with vision impairments. *arXiv preprint arXiv:2301.05323*, 2023. 3
- [53] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Acdc: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *CVPR*, 2021. 3
- [54] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *CVPR*, 2021. 3
- [55] Wu Tang, De-er Liu, Xiaoli Zhao, Zenghui Chen, and Chen Zhao. A dataset for the recognition of obstacles on blind sidewalk. *Universal Access in the Information Society*, 2021. 7, 8
- [56] Shishun Tian, Minghuo Zheng, Wenbin Zou, Xia Li, and Lu Zhang. Dynamic crosswalk scene understanding for the visually impaired. *TNSRE*, 2021. 3
- [57] Yu-Yun Tseng, Alexander Bell, and Danna Gurari. VizWiz-FewShot: Locating objects in images taken by people with visual impairments. In *ECCV*, 2022. 3
- [58] Svenja Uhlemeyer, Matthias Rottmann, and Hanno Gottschalk. Towards unsupervised open world semantic segmentation. In *UAI*, 2022. 3
- [59] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. CIDEr: Consensus-based image description evaluation. In *CVPR*, 2015. 2
- [60] 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. 5
- [61] Meng Wei, Long Chen, Wei Ji, Xiaoyu Yue, and Tat-Seng Chua. Rethinking the two-stage framework for grounded situation recognition. In *AAAI*, 2022. 3, 6
- [62] Kan Wu, Jinnian Zhang, Houwen Peng, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. TinyViT: Fast pretraining distillation for small vision transformers. In *ECCV*, 2022. 5
- [63] Yu-Huan Wu, Yun Liu, Xin Zhan, and Ming-Ming Cheng. P2T: Pyramid pooling transformer for scene understanding. *TPAMI*, 2022. 3

- [64] Haiying Xia, Cong Yao, Yumei Tan, and Shuxiang Song. A dataset for the visually impaired walk on the road. *Displays*, 2023. 7, 8
- [65] Yingda Xia, Yi Zhang, Fengze Liu, Wei Shen, and Alan L Yuille. Synthesize then compare: Detecting failures and anomalies for semantic segmentation. In *ECCV*, 2020. 3
- [66] 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. 3
- [67] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. GroupViT: Semantic segmentation emerges from text supervision. In *CVPR*, 2022. 3
- [68] Jilan Xu, Junlin Hou, Yuejie Zhang, Rui Feng, Yi Wang, Yu Qiao, and Weidi Xie. Learning open-vocabulary semantic segmentation models from natural language supervision. In *CVPR*, 2023. 3
- [69] Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. Situation recognition: Visual semantic role labeling for image understanding. In *CVPR*, 2016. 5
- [70] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023. 2, 5
- [71] Jiaming Zhang, Kailun Yang, Angela Constantinescu, Kunyu Peng, Karin Müller, and Rainer Stiefelhagen. Trans4Trans: Efficient transformer for transparent object segmentation to help visually impaired people navigate in the real world. In *ICCVW*, 2021. 1, 3, 4, 5
- [72] Yan Zhang, Ziang Li, Haole Guo, Luyao Wang, Qihe Chen, Wenjie Jiang, Mingming Fan, Guyue Zhou, and Jiangtao Gong. “I am the follower, also the boss”: Exploring different levels of autonomy and machine forms of guiding robots for the visually impaired. In *CHI*, 2023. 3
- [73] Xu Zhao, Wenchao Ding, Yongqi An, Yinglong Du, Tao Yu, Min Li, Ming Tang, and Jinqiao Wang. Fast segment anything. *arXiv preprint arXiv:2306.12156*, 2023. 5
- [74] Junwei Zheng, Jiaming Zhang, Kailun Yang, Kunyu Peng, and Rainer Stiefelhagen. MateRobot: Material recognition in wearable robotics for people with visual impairments. *arXiv preprint arXiv:2302.14595*, 2023. 1, 3
- [75] 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
- [76] Wujie Zhou, Shaohua Dong, Jingsheng Lei, and Lu Yu. MTANet: Multitask-aware network with hierarchical multimodal fusion for RGB-T urban scene understanding. *T-IV*, 2022. 3
- [77] Wenbin Zou, Guoguang Hua, Yue Zhuang, and Shishun Tian. Real-time passable area segmentation with consumer RGB-D cameras for the visually impaired. *TIM*, 2023. 3