

Evaluating SAM2 for Video Semantic Segmentation

Syed Ariff Syed Hesham^{1,2}

Yun Liu^{3*}

Guolei Sun^{3*}

Jing Yang⁴

Henghui Ding⁵

Xue Geng²

Xudong Jiang¹

¹School of Electrical and Electronic Engineering (EEE), Nanyang Technological University, Singapore 639798, Singapore

²Institute for Infocomm Research, A*STAR, Singapore 138632, Singapore

³College of Computer Science & Academy for Advanced Interdisciplinary Studies, Nankai University, Tianjin 300350, China

⁴State Key Laboratory of Public Big Data, Guizhou University, Guiyang 550025, China

⁵Fudan Vision and Learning Lab, Fudan University, Shanghai 200438, China

Abstract: The segment anything model 2 (SAM2) has proven to be a powerful foundation model for promptable visual object segmentation in both images and videos, capable of storing object-aware memories and transferring them temporally through memory blocks. While SAM2 excels in video object segmentation by providing dense segmentation masks based on prompts, extending it to dense video semantic segmentation (VSS) poses challenges due to the need for spatial accuracy, temporal consistency, and the ability to track multiple objects with complex boundaries and varying scales. This paper explores the extension of SAM2 for VSS, focusing on two primary approaches and highlighting firsthand observations and common challenges faced during this process. The first approach involves using SAM2 to extract unique objects as masks from a given image, with a segmentation network employed in parallel to generate and refine initial predictions. The second approach utilizes the predicted masks to extract unique feature vectors, which are then fed into a simple network for classification. The resulting classifications and masks are subsequently combined to produce the final segmentation. Our experiments suggest that leveraging SAM2 enhances overall performance in VSS, primarily due to its precise predictions of object boundaries.

Keywords: Video semantic segmentation (VSS), segment anything model 2 (SAM2), visual foundation models, boundary refinement, temporal consistency, mask-based classification.

Citation: S. A. S. Hesham, Y. Liu, G. Sun, J. Yang, H. Ding, X. Geng, X. Jiang. Evaluating SAM2 for video semantic segmentation. *Machine Intelligence Research*, vol.23, no.5, pp.998–1011, 2026. <http://doi.org/10.1007/s11633-026-1638-9>

1 Introduction

Segmentation is a fundamental task in computer vision that involves identifying and extracting the region of interest from a given image, representing an object of interest such as a person, barrier, or road. A specialized task in this field is video semantic segmentation (VSS)^[1], which focuses on classifying and segmenting each pixel in a video sequence into predefined categories over time. Unlike traditional image segmentation, which deals with static images, VSS accounts for the temporal dimension of video data, allowing for the identification and tracking of objects across multiple frames. This involves not only recognizing objects but also maintaining their identities and boundaries as they move, change shape, or interact

with other elements in the scene^[2, 3].

Recently, there has been an increasing interest in applying foundation models in numerous visual tasks like classification, detection, and VSS. The excitement in the usage of foundation models originates from their extensive pre-training on large web-scale datasets, which in turn enhances their ability to generalize effectively to various downstream tasks. The commercial success of ChatGPT, which uses the foundation model GPT^[4] for text generation, has accelerated the exploration of foundation models in other modalities, including the field of computer vision^[5–8]. The early developments of foundation models in computer vision include CLIP^[9], which makes use of a vast collection of image-text pairs from the web to align both modalities effectively. This was then applied to recognize new visual categories through text prompts and provided zero-shot generalization for vision concepts, ultimately inspiring the creation of foundation models for object detection called GLEE^[10]. Meta AI's introduction of the segment anything model (SAM)^[11] for promptable segmentation, further broadened the application of foundation models across multiple domains, including entertainment, medicine, and forensics.

Research Article

Special Issue on Multimodal Learning, Temporal Modeling, and Foundation Models for Video Understanding

Manuscript received on April 8, 2025; accepted on January 29, 2026; published online on August 24, 2026

Recommended by Associate Editor Chunhua Shen

Colored figures are available in the online version at <https://link.springer.com/journal/11633>

* Corresponding Author: Yun Liu (E-mail: liuyun@nankai.edu.cn); Guolei Sun (E-mail: guolei.sun@nankai.edu.cn)

© Institute of Automation, Chinese Academy of Sciences 2026

While SAM excels at segmenting static images, its effectiveness in real-world scenarios is restricted due to its inability to address the temporal information in streaming images/video data^[12]. To overcome this challenge, Meta AI developed SAM2^[13], an enhanced version of the original SAM that incorporates a memory module. This enhancement allows the storage of context memories for each generated masklet, allowing SAM2 to process video frames sequentially while referencing its prior memories of the objects of interest, which leads to a coherent segmentation of these objects over time. Together, SAM and SAM2 have sparked considerable interest among researchers, motivating investigations into SAM's potential for various applications and downstream tasks. Research has explored SAM's segmentation abilities across different contexts, including the medical domain and its role as a semantic segmentor^[14–20]. However, there has been limited focus on evaluating the capabilities of SAM2 in the context of VSS, particularly in challenging scenarios.

In this paper, we extend SAM2's application specifically to VSS in two main ways. First, we evaluate SAM2 as a post-processing refiner to extract different objects in a scene as masklets and then use them to enhance the segmentation quality of existing models, assessing the resulting improvements. Second, we explore SAM2's independent functionality as a semantic segmentor, where we utilize features from its image encoder to vectorize the representation of each detected object in an image and perform classification. Our experiments are conducted on the established VSPW dataset^[1], which includes a broad spectrum of short videos showcasing various motion levels, allowing for a thorough assessment of SAM2's capabilities in this specific task.

Overall, the contributions of this paper include:

- 1) We evaluate the effectiveness of SAM2 as a post-processing refiner to improve the quality of masks generated by VSS models across multiple datasets, identifying areas for potential enhancement.
- 2) We investigate the standalone functionality of SAM2 to work independently (I-SAM2) as a video semantic segmentor, focusing on its performance with minimal modifications to the base model.
- 3) We assess SAM2's performance using VSS benchmarks to identify its limitations and applicability in this domain, providing insights into the adaptation requirements of foundation models for dense prediction tasks.

2 Related work

Segmentation is a key challenge in computer vision, focusing on pixel-wise classification of images. It plays a crucial role in interpreting image information and is primarily explored through three main approaches: semantic segmentation, instance segmentation, and panoptic segmentation. Semantic segmentation^[21–23] involves assigning each pixel to a predefined category. Whereas in-

stance segmentation^[24, 25] takes this a step further by separating segments into individual instances within the same class. Panoptic segmentation, as introduced in ^[26], goes beyond both semantic and instance segmentation to provide a comprehensive understanding of the entire scene. While these techniques lay the foundation for image segmentation, the complexity of video data introduces additional challenges and opportunities for research, as VSS extends the principles of pixel-wise classification into the temporal domain to provide coherent segmentation across video frames.

2.1 Video semantic segmentation

The emergence of deep learning, particularly with the introduction of convolutional networks, has promoted extensive research focused on segmentation and its applications across various fields, including robotics^[27–31], autonomous driving^[32–36], and medical image analysis^[37–42]. While there has been a trend toward developing models specialized in segmentation, most approaches are developed to process images independently^[38, 43–47]. This independent processing presents limitations when applying image-based models to the real world, which is inherently continuous. Consequently, there is a growing interest in developing models specifically designed for videos, allowing them to incorporate temporal information and ensure smoother and more consistent predictions over time^[2, 48–55]. This increased focus on video segmentation has led to significant advancements that have positively impacted the overall performance of models in this domain^[3, 56–63], including recent modular architectures that leverage foundation models^[64–66]. However, the demand for further performance improvements remains crucial. Therefore, it becomes essential to explore the segmentation using the latest techniques to continue enhancing performance.

2.2 Segment anything

Recently, there has been a growing interest in generalized AI frameworks, commonly termed foundational models^[67]. These models make use of web-scale datasets coupled with a self-supervised learning approach to achieve powerful zero-shot generalization across a wide range of downstream tasks. Significant advancements in natural language processing (NLP) leading to models like BERT^[68] and GPT series^[4, 69–71] have inspired the computer vision field to create large visual models (LVMs)^[5–7]. Additionally, research integrating visual and textual data has produced models like CLIP^[9] and ALIGN^[72], making multimodal AI an exciting field. In line with the current trends, a significant contribution in this domain is Meta's SAM^[11], which aims to provide a versatile and promptable solution for image segmentation tasks.

While SAM excels in zero-shot generalization, it faces challenges when applied to video tasks as it needs to establish efficient frame-to-frame associations, process additional dimensions, handle the temporal information, and maintain temporal consistency. Several methods, such as track anything model (TAM)^[73] and SAM-Track^[74] for object tracking, and SAMText^[75] for video text spotting, have attempted to address these challenges. However, they often require adjustments to the model architecture for optimal performance. SAM2^[13] builds upon the original SAM^[11] to tackle these limitations. It achieves this by integrating a streaming memory, and further refines the segmentation by making use of interactive user prompts coupled with a memory attention mechanism, which allows the system to track target objects across frames, allowing for real-time video segmentation. As a result, SAM2 not only enhances the accuracy of segmentation but also improves overall efficiency with fewer user interactions, allowing for more effective segmentation in spatio-temporal contexts. The generalizability of SAM2 to various complex real-world scenarios is achieved through training on the SA-V dataset, a significantly larger and more diverse expansion of the dataset that was used to train SAM, comprising an impressive 50.9K videos and 652.6K masklets.

2.2.1 Video applications

The introduction of SAM2^[13] has led to a wave of research focused on various applications, including video generation^[76–78], 3D reconstruction^[79, 80], deepfake detection^[81], video editing^[82, 83], and model compression^[84]. Its applications in video object segmentation (VOS) cover areas like video object tracking (VOT)^[73, 85], audio-visual segmentation^[86], and domain-specific applications, including medical video segmentation^[87], among others^[88, 89].

While SAM has been applied to different use cases, there has been limited work done on a comprehensive evaluation of SAM2 in real-world video scenarios. This evaluation is crucial for assessing its performance and identifying the factors that influence its effectiveness as both a post-processing method and a standalone semantic segmentor. Therefore, this work aims to evaluate SAM2's performance in real-world scenarios, conducting thorough tests on the various parameters to report its direct practical applicability to images captured in the wild.

3 Methods & experiments

This section provides an overview of the experiments conducted, outlining the datasets used, the evaluation metrics employed, and the specific implementation settings applied throughout the study.

3.1 Methods

To evaluate the effectiveness of SAM2^[13] as both a post-processing refiner and a standalone tool for generat-

ing segmentation masks, the automatic mask generator is used with a 16×16 point grid as input prompts. Furthermore, to evaluate the SAM2's ability to refine masks while ensuring new objects are tracked, a window of 32 frames is used to share and propagate masklets before they are recomputed. In the experiments, thresholds for the predicted intersection-over-union (IoU), stability, and stability offset are established at 0.6, 0.8, and 0.9, respectively, based on the ablation studies and visual assessments discussed in Section 4.1.2. A detailed discussion of the evaluation setup for SAM2 in both contexts is covered in this section.

3.1.1 SAM2 as refiner

Using SAM2 as a refiner serves as a qualitative study to evaluate its effectiveness as a post-processing algorithm for refining masks produced by a trained segmentation model. The aim is to assess whether SAM2's capabilities as a generic foundation model can effectively extract unique objects from a scene, thereby potentially enhancing the uniformity and accuracy of the predicted masks through improved detection and refinement of object-level details.

The approach for using SAM2 as a refiner is illustrated in Fig. 1. In this method, a semantic segmentation model generates a semantic mask S , while SAM2 operates in parallel as an object segmentation network that extracts individual objects from the scene as masklet $M' = \{m_1, m_2, \dots, m_n\}$. After passing through an aggregator, each masklet m_i is used to refine the predicted semantic mask by determining the predominant class c_i within the masklet. The refined segmentation mask R is then computed as follows:

$$R(x, y) = \begin{cases} c_i, & \text{if } (x, y) \in m_i \\ S(x, y), & \text{otherwise.} \end{cases} \quad (1)$$

Here, $S(x, y)$ is the original semantic mask at pixel location (x, y) and c_i represents the predominant class within the masklet m_i . To determine c_i , we compute the class label that appears most frequently within the mask m_i :

$$c_i = \arg \max_{c \in \mathcal{C}} \sum_{(x, y) \in m_i} \delta(S(x, y), c). \quad (2)$$

Here, $\delta(S(x, y), c)$ is the Kronecker delta function, which returns 1 if $S(x, y) = c$ and 0 otherwise. Applying these equations across all the pixels produces the final refined mask R . Evaluations are then performed on this refined mask R to evaluate how the performance changes after using SAM2 as a post-processing unit.

3.1.2 SAM2 as segmentor

To investigate SAM2's potential as an independent semantic segmentor (I-SAM2) with minimal architectural modifications, a framework incorporating only the essential changes is employed, as illustrated in Fig. 2. The

primary change involves the use of a simple classifier composed of a multi-layer perceptron (MLP) layer. The process of obtaining a segmentation mask begins with the images or live-stream frames as input, which are encoded by SAM2's image encoder. The encoded features F with feature dimension d are subsequently processed through a mask decoder, resulting in the generation of masklets $M = \{m_1, m_2, \dots, m_n\}$ that represent n detected object regions in the scene. A mean of the encoded features corresponding to the area covered by each masklet is computed to derive a representation for each masklet. This yields a single vector that characterizes each mask which can be expressed as

$$v_i = \frac{1}{|m_i|} \sum_{(x,y) \in m_i} F(x,y). \quad (3)$$

Here, $v_i \in \mathbb{R}^d$ represents a feature vector of the i -th mask ($i \in \{1, 2, \dots, n\}$), $F(x,y) \in \mathbb{R}^d$ is the encoded feature vector at pixel location (x,y) , and $|m_i|$ represents the cardinality, i.e., the number of pixels covered by the mask m_i .

The feature vectors $\{v_1, v_2, \dots, v_n\}$ are subsequently fed into the classifier for class prediction of each vector. The outputs from this classification process are combined with the masklets to create a unified segmentation mask, which is then used to assess the performance of SAM2 as a standalone semantic segmentor.

To further investigate the capabilities of SAM2 and compare its effectiveness with semantic information encoding, evaluations are conducted on other VSS methods by using their encoders to extract encoded features for training independent classifiers. Following the methodology outlined in (3), the encoded features from the trained segmentation models are extracted and used to train the classifiers for evaluation purposes.

Additionally, it is observed that the detection of masklets within the images occasionally results in low mask coverage, leaving significant areas without predicted labels. To address this issue, a base model is introduced, which fills areas not covered by the masklets with predictions from the base segmentor, as shown in Fig. 2. The incorporation of a base model aids in providing a thorough evaluation across diverse scenarios. Further in-




Fig. 1 Framework for using SAM2^[13] as a refiner to post-process the predicted masks from a segmentation model (Colored figures are available in the online version at <https://link.springer.com/journal/11633>)




Fig. 2 Framework for using SAM2^[13] as an independent semantic segmentation model (Colored figures are available in the online version at <https://link.springer.com/journal/11633>)

sights and discussions regarding the evaluations and their results can be found in Section 4.2.

3.2 Dataset

There are numerous datasets available for VOS and video instance segmentation (VIS), including one from SAM2^[16]. However, there is a limited availability of densely annotated datasets for evaluating semantic segmentation, with the VSPW dataset^[1] being a notable exception. The VSPW dataset stands out by providing dense annotations at a frame rate of 15FPS and includes 124 unique categories. It contains 2 806 training clips (198 244 frames), 343 validation clips (24 502 frames), and 387 test clips (28 887 frames), and it has been recognized by the community as a benchmarking dataset for VSS. We chose to use the VSPW dataset for our experiments as it covers a wide range of indoor and outdoor scenes, making it particularly suitable for performance evaluation.

3.3 Performance metrics

A well-established metric for semantic segmentation, mean IoU (mIoU) is used for quantitative assessment of the experiments. Considering the characteristics of SAM2 as a generic foundation model suitable for video object segmentation, the evaluation extends beyond mIoU to include additional metrics to assess visual consistency across frames and the precision of detected object boundaries. Such multidimensional evaluation provides deeper insights into the model's effectiveness, ensuring that the segmentation results maintain high fidelity over temporal video sequences. Firstly, to assess the visual consistency of the predictions, the metric mean video consistency (mVC)^[1] is used. The metric mVC specifically measures the temporal smoothness of the predicted segmentation maps across frames, assessing performance in the temporal domain. Formally, for a video clip represented as $\{\mathbf{I}_c\}_{c=0}^{C_v}$, with corresponding ground-truth masks $\{\mathbf{M}_c\}_{c=1}^{C_v}$ and predicted outputs $\{\mathbf{O}_c\}_{c=1}^{C_v}$, the calculation of video consistency VC_n is performed as follows:

$$\text{VC}_n = \frac{1}{C_v - n + 1} \sum_{i=1}^{C_v - n + 1} \frac{\left(\bigcap_{i=1}^{i+n-1} \mathbf{M}_i \right) \cap \left(\bigcap_{i=1}^{i+n-1} \mathbf{O}_i \right)}{\bigcap_{i=1}^{i+n-1} \mathbf{M}_i}. \quad (4)$$

The final mVC_n is computed by taking the mean of all the calculated VC_n values, given that $C_v \geq n$. As shown in (4), mVC_n captures the common areas of the predicted masks across frames, reflecting the consistency of these predictions over time. For further details on this metric, please refer to [1].

Secondly, to evaluate the accuracy of object boundaries, the ground-truth masks M , consisting of C classes,

are processed using morphological operations with a 5-pixel kernel K . This processing extracts the boundaries for each class $c \in C$, resulting in a boundary mask B . Subsequently, the IoU is computed between these boundary masks B and the predicted output masks O . An average of the IoU values across N images in the dataset provides the mean boundary IoU, referred to as mBIoU. This process is formulated as

$$B = \bigcup_{c \in C} (M_c \ominus K) \quad (5)$$

$$\text{mBIoU} = \frac{1}{N} \sum_{i=1}^N \frac{|(M_i \cap O_i) \cap B|}{|M_i \cap B| + |O_i \cap B| - |(M_i \cap O_i) \cap B|}. \quad (6)$$

The above (6) acts as a comprehensive measurement for the accuracy of borders for the predicted segmentation masks, which helps in evaluating the alignment of the predicted boundaries with the ground truth. The metric will be particularly beneficial in applications where the precise delineation of object edges is of critical importance, such as in medical imaging, autonomous driving, and various other computer vision tasks.

4 Results and ablations

This section provides a comprehensive analysis of the results obtained from various experiments conducted on SAM2, aimed at applying it to the semantic segmentation task. All experiments were performed using the VSPW dataset^[1]. The following discussion will present the findings and offer insights into the interpretation of the results obtained.

4.1 Results as refiner

The performance of SAM2 as a mask refiner is evaluated using state-of-the-art VSS models, following the framework illustrated in Fig. 1. The evaluation results presented in Table 1 indicate that SAM2 is effective as a refiner, enhancing the overall performance of the models across all evaluation metrics. Furthermore, SAM2 shows significant improvements in object segmentation as reflected by the notable increase in mBIoU scores. A consistent improvement in visual consistency across various VSS backbones and their variants (MiT-B0 to MiT-B5) can be attributed to SAM2's ability to treat objects as masklets and effectively track them across frames. The consistent enhancements across different evaluation metrics underscore SAM2's effectiveness as a post-processing tool, particularly in refining object boundaries and improving overall segmentation accuracy and consistency. Qualitative results in Fig. 3 illustrate these improvements, showing that SAM2 refines object boundaries and reduces seg-

Table 1 Quantitative comparison of our model with existing methods on the VSPW dataset^[1] with an input resolution of 480×853

Methods	Backbones	mIoU $\uparrow$ (%)	Weighted IoU $\uparrow$ (%)	mBIoU $\uparrow$ (%)	mVC $\uparrow$ (%)	mVC $\uparrow$ (%)
Segformer ^[90]	MiT-B0	32.9	56.8	–	82.7	77.3
CFFM ^[2]	MiT-B0	35.4	58.5	21.8	87.7	82.9
MRCFA ^[49]	MiT-B0	35.2	57.9	21.5	88.0	83.2
SAM2-CFFM	MiT-B0	36.7	60.0	25.8	88.8	84.4
SAM2-MRCFA	MiT-B0	36.8	59.2	25.8	89.1	84.6
Segformer ^[90]	MiT-B1	36.5	58.8	–	84.7	79.9
CFFM ^[2]	MiT-B1	38.5	60.0	24.7	88.6	84.1
MRCFA ^[49]	MiT-B1	38.9	60.0	24.7	88.8	84.4
TV3S ^[63]	MiT-B1	40.0	–	25.6	90.7	87.0
SAM2-CFFM	MiT-B1	40.3	61.1	29.3	89.4	85.1
SAM2-MRCFA	MiT-B1	40.6	60.9	29.4	89.5	85.2
SAM2-TV3S	MiT-B1	41.1	62.3	30.1	91.4	87.8
Segformer ^[90]	MiT-B2	43.9	63.7	–	86.0	81.2
CFFM ^[2]	MiT-B2	44.9	64.9	28.9	89.8	85.8
MRCFA ^[49]	MiT-B2	45.3	64.7	28.7	90.3	86.2
TV3S ^[63]	MiT-B2	46.3	–	29.8	91.5	88.4
SAM2-CFFM	MiT-B2	46.7	65.3	33.9	90.4	86.5
SAM2-MRCFA	MiT-B2	46.9	65.7	33.8	91.0	87.1
SAM2-TV3S	MiT-B2	47.8	65.2	34.8	92.1	89.1
Segformer ^[90]	MiT-B5	48.9	65.1	–	87.8	83.7
CFFM ^[2]	MiT-B5	49.3	65.8	32.1	90.8	87.1
MRCFA ^[49]	MiT-B5	49.9	66.0	32.0	90.9	87.4
TV3S ^[63]	MiT-B5	49.8	–	32.9	91.7	88.7
SAM2-CFFM	MiT-B5	50.8	66.3	37.0	91.2	87.7
SAM2-MRCFA	MiT-B5	51.3	66.6	37.1	91.5	88.2
SAM2-TV3S	MiT-B5	51.3	67.1	38.1	92.2	89.3




Fig. 3 Qualitative comparison of SAM2 refinement on VSPW validation set. SAM2 refinement improves boundary precision and reduces noise in predicted masks, particularly for objects with complex boundaries. (Colored figures are available in the online version at <https://link.springer.com/journal/11633>)

mentation noise across diverse scene categories.

Computational cost analysis. We evaluate the computational overhead of the SAM2-refiner on the

VSPW validation set using a single NVIDIA A100 GPU. The refiner operates at 1.1 to 3.5FPS depending on scene complexity and the number of objects tracked, compared

to 4FPS to 15FPS for base to large VSS models. This computational cost positions SAM2 as the most suitable for offline refinement scenarios where segmentation quality, particularly boundary accuracy and temporal consistency, is prioritized over real-time performance.

4.1.1 Cross-dataset generalization

To evaluate the generalization capability of SAM2 as a refiner across different datasets and scenes, we conduct experiments on VIPSeg^[91], a large-scale video panoptic segmentation dataset with 124 classes covering diverse in-the-wild scenarios. We apply the SAM2-refiner to predictions from recent state-of-the-art video panoptic segmentation methods by converting their panoptic outputs to semantic segmentation masks. Specifically, we evaluate CAVIS^[64], DVIS^[65], and DVIS++^[66] with different backbones without any fine-tuning or parameter adjustment to the SAM2-refiner. The results in Table 2 demonstrate that SAM2 consistently improves boundary accuracy as indicated by mBIoU gains across all methods, while maintaining competitive performance on other metrics. This validates that the refinement capability of SAM2 generalizes well to diverse datasets and modern architectures without dataset-specific tuning.

Table 2 Cross-dataset generalization of SAM2 as refiner on VIPSeg dataset

Methods	mIoU↑ (%)	mBIoU↑ (%)	mVC ₈ ↑ (%)	mVC ₁₆ ↑ (%)
CAVIS (R50) ^[64]	50.4	55.3	86.2	79.4
DVIS++ (R50) ^[66]	49.9	55.5	85.8	79.5
DVIS (Swin-L) ^[65]	55.2	57.2	89.5	84.2
DVIS (Swin-L) [†] ^[65]	57.0	60.9	91.4	87.7
SAM2-CAVIS	50.3	56.8	86.7	80.1
SAM2-DVIS++	50.0	57.0	86.3	80.1
SAM2-DVIS	55.0	59.1	89.8	84.6
SAM2-DVIS (offline)	57.2	62.3	91.8	88.3

4.1.2 Ablations on refiner

Ablation studies are conducted to analyze the impact of various parameters during evaluation and their effects on the final segmentation results. All ablation experiments are conducted by making use of the MiT-B1 backbone^[90] of the CFFM^[2] VSS model, following the default parameters outlined in Section 3. The ablations are conducted on four parameters that significantly influence model performance during inference: the temporal window size for mask propagation, the number of point grid inputs used as prompts, and the thresholds for prediction confidence and masklet stability. Understanding and testing these parameters is crucial for optimizing the model's effectiveness and ensuring reliable results.

Effect of temporal window size. The evaluation of SAM2's ability to propagate masklets across sequential

frames is performed by varying the temporal window size for masklet propagation. Using a temporal window allows for the detection of new objects while refreshing information about existing masklets, leading to more accurate segmentation. The results of the experiments are summarized in Table 3, where the evaluations begin with a single-frame temporal window and are incrementally increased up to a total of 32 frames. It is observed that, while increasing the number of frames has a negligible impact on the quantitative quality of the segmentation masks (mIoU), there is a consistent improvement in visual consistency. This improvement can be attributed to the tracking of masklets across the temporal dimension, which leads to more stable predictions of individual objects within the scene. Considering this observation, along with the advantage of faster inference times in SAM2 with increased temporal window, a temporal window of 32 frames is used for all subsequent experiments.

Table 3 Ablation study on varying temporal window sizes (i.e., the number of frames)

#Frames	mIoU (%)	FWIoU (%)	mVC ₈ (%)	mVC ₁₆ (%)
1	40.3	61.1	88.2	83.9
2	40.3	61.1	88.4	84.0
4	40.3	61.1	88.6	84.3
8	40.3	61.0	88.9	84.5
10	40.3	61.0	89.1	84.7
16	40.4	61.1	89.2	84.8
32	40.4	61.1	89.4	85.1

Influence of grid prompt size. To effectively segment all objects in a scene using an automatic prompt segmentor, selecting the appropriate size for the point grid to perform prompting is essential. Based on the evaluations in Table 4, there are minimal differences in performance across various prompt sizes. While the 8×8 option might seem appealing due to potentially faster inference times, this study focuses on identifying all objects in the scene with minimal omissions, ensuring an accurate assessment of semantic segmentation. Upon further visual inspection, the 16×16 grid appears to be the optimal choice, as it does not significantly impact evaluation time. However, in scenarios where rapid inference is critical and some object omission is acceptable, smaller prompts may provide advantages based on the observed results.

Impact on the choice of thresholds. The impact of the confidence threshold has been thoroughly evaluated and documented in Table 5, while assessments of stability are provided in Table 6. The selection of optimal thresholds for confidence and stability should be guided by the specific requirements of the downstream tasks that utilize the segmentation capabilities of SAM2. A high confidence threshold results in a limited number

Table 4 Ablation study on the grid prompt size

Prompt	mIoU (%)	FWIoU (%)	mVC ₈ (%)	mVC ₁₆ (%)
8 × 8	39.9	61.0	89.3	85.1
16 × 16	40.4	61.1	89.4	85.1
32 × 32	40.3	61.1	89.3	85.0
64 × 64	40.2	61.0	89.2	84.9

Table 5 Ablation study on the confidence threshold

IoU Conf.	mIoU (%)	FWIoU (%)	mVC ₈ (%)	mVC ₁₆ (%)
0.4	40.1	61.1	89.3	84.9
0.5	40.2	61.1	89.3	85.0
0.6	40.4	61.1	89.4	85.1
0.7	40.3	61.1	89.4	85.1
0.8	40.2	61.0	89.5	85.2
0.9	39.8	60.8	89.3	85.1

of masklets that align well with the objects, often leading to the exclusion of many masklets due to stringent thresholding, while a low confidence threshold can produce an overwhelming number of masklets, even if they are poorly aligned. In terms of stability, a high threshold ensures consistency among the masklets, whereas a low threshold may permit unstable masklets to be included. For the purpose of this study, aimed at achieving effective semantic segmentation, it is essential to obtain a sufficient number of masklets while maintaining the stability of the predictions. Thus, the threshold values are implemented with 0.6 for confidence and 0.8 for stability to establish an optimal balance.

4.2 Results as segmentor

Using SAM2 for semantic segmentation shifts the focus from classifying each pixel into specified categories to

Table 6 Ablation study on the stability

Stability	mIoU (%)	FWIoU (%)	mVC ₈ (%)	mVC ₁₆ (%)
0.4	40.5	61.2	89.3	85.0
0.5	40.5	61.2	89.3	85.0
0.6	40.4	61.2	89.3	85.0
0.7	40.4	61.1	89.3	85.0
0.8	40.4	61.1	89.4	85.1
0.9	40.0	60.9	89.3	85.0

detecting every individual item in the scene – both “thing” and “stuff” – before classifying these detected objects with a classifier. As discussed in Section 3.1.2, an evaluation is conducted to compare the independent SAM2 (I-SAM2) with other VSS models, CFFM^[2], in a similar setting (I-CFFM). Additionally, to verify the implication of the missed masks, a baseline mode is employed with optimal thresholding for the current task, as discussed in Section 4.1.2, where the regions not covered by the masklets are filled with the segmentation predictions. For the I-SAM2 base mode (I-SAM2_{base}), the segmentation results from CFFM variant with the MiT-B1 backbone^[90] are used.

The evaluation, detailed in Table 7, indicates that SAM2 performs poorly in this task, even with assistance from the segmentors in the base mode, which has the worst performance compared to other methods. This trend holds true even for the trained VSS models, as there is a noticeable decline in performance, with a drop of at least 10% in the quantitative accuracy of predictions, measured by mIoU values. Furthermore, there is a significant decrease in the stability of the results, evidenced by a drastic decline in visual consistency. The drastic drop in performance in SAM2 strongly suggests that it is not suitable for independent semantic segmentation when the core SAM2 model is not retrained for the specific task.

Table 7 Quantitative comparison of our model with existing methods on the VSPW dataset^[1] with an input resolution of 480 × 853

Methods	Backbones	mIoU↑ (%)	Weighted IoU↑ (%)	mBIoU↑ (%)	mVC ₈ ↑ (%)	mVC ₁₆ ↑ (%)
CFFM ^[2]	MiT-B0	35.4	58.5	21.5	87.7	82.9
I-CFFM	MiT-B0	29.9	49.4	21.1	66.9	58.1
I-CFFM _{base}	MiT-B0	33.5	56.9	21.4	73.9	66.8
CFFM	MiT-B1	38.5	60.0	24.7	88.6	84.1
I-CFFM	MiT-B1	30.3	49.9	21.3	66.0	57.3
I-CFFM _{base}	MiT-B1	34.1	57.6	22.0	73.2	66.2
CFFM	MiT-B2	44.9	64.9	28.9	89.8	85.8
I-CFFM	MiT-B2	37.5	55.1	27.1	70.8	62.7
I-CFFM _{base}	MiT-B2	42.3	63.2	28.1	79.1	73.0
I-SAM2	–	18.0	36.5	13.3	52.8	43.7
I-SAM2 _{base}	–	20.9	41.8	15.7	55.5	46.6

Analysis of I-SAM2 limitations. The performance gap observed in I-SAM2 can be understood by examining the fundamental differences between object segmentation and semantic segmentation objectives. SAM2 is trained for promptable object segmentation, where the primary goal is to delineate individual objects regardless of their semantic category. Consequently, the features learned by SAM2's image encoder are optimized for capturing general object properties such as boundaries, textures, and spatial coherence rather than class-discriminative information. This contrasts with models like CLIP^[9], which explicitly align visual and textual representations to encode semantic meaning, or task-specific VSS encoders that are supervised with dense categorical labels. The limitations are further compounded by the feature aggregation strategy in (3), which computes a mean feature vector over each masklet assuming semantic homogeneity within detected regions. While this assumption holds for class-aware encoders, it becomes problematic for SAM2's class-agnostic representations where averaging may dilute any latent discriminative signals. These observations suggest that leveraging SAM2 for semantic segmentation benefits more from utilizing its precise boundary predictions as a refinement tool rather than relying on its feature representations for direct classification.

5 Challenges and future directions

The limitations observed in our standalone I-SAM2 experiments reveal fundamental challenges in adapting foundation models for dense VSS. As discussed in Section 4.2, SAM2's class-agnostic features optimized for objectness struggle with semantic discrimination required for pixel-level classification. This finding highlights the need for targeted adaptations that inject semantic awareness while preserving SAM2's strong boundary prediction capabilities. To address this, specialized modules or adapters can enhance SAM2's ability to understand and track multiple objects with complex boundaries and varying scales. For instance, SAM2-Adapter^[92] has been developed to adapt SAM2 for downstream tasks, achieving state-of-the-art results in areas like camouflaged object detection and shadow detection.

Building on these insights, future work could explore lightweight parameter-efficient strategies to bridge the semantic gap we have identified. Low-rank adaptation (LoRA) could selectively update attention weights to enable semantic discrimination without full fine-tuning of SAM2's encoder^[92, 93]. Alternatively, learnable fusion modules could dynamically weight SAM2's object-aware features against task-specific semantic embeddings^[94], replacing our fixed majority voting approach with adaptive integration that learns optimal combination strategies during training^[95, 96]. Such approaches would directly address the feature-level limitations revealed in our experiments while maintaining SAM2's zero-shot generalization

for novel object categories.

Beyond semantic adaptation, evaluating SAM2 in open-vocabulary segmentation tasks highlights additional opportunities for enhancement. While SAM2 shows strong zero-shot generalization, its performance in open-vocabulary segmentation can be further improved by incorporating composable prompts. The SAM-CP^[97] approach pairs SAM2 with composable prompts and has achieved superior results in panoptic segmentation. This suggests that augmenting SAM2 with language-driven prompting could enhance its ability to segment a wider range of object categories in diverse and dynamic environments, complementing the semantic adaptations discussed above.

6 Conclusions

This study investigates the application of SAM2 in VSS by evaluating it across multiple datasets and architectures. Our experiments on VSPW and VIPSeg datasets demonstrate that SAM2, as a mask refiner, significantly enhances semantic segmentation performance by improving the precision of object boundaries and enhancing visual consistency across various existing segmentation models. The consistent improvements in mBIOU across different methods and datasets validate SAM2's effectiveness as a post-processing refinement tool. However, when evaluated as a standalone segmentor, SAM2 performs subpar compared to models explicitly trained for semantic segmentation. This limitation stems from SAM2's class-agnostic feature representations optimized for objectness rather than semantic discrimination. These findings underscore the critical need to adapt foundation models to specific tasks to fully leverage their advanced capabilities, while highlighting SAM2's practical value as a quality-focused refinement module for VSS applications.

Acknowledgements

This work was supported in part by the A*STAR Computing and Information Science (ACIS) Scholarship, Singapore, in part by the National Natural Science Foundation of China (No. 62576176), and in part by the Fundamental Research Funds for the Central Universities, China (Nankai University, No. 070-63253235). The computational devices is supported by the Supercomputing Center of Nankai University (NKSC), China.

Declarations of conflict of interest

The authors declared that they have no conflicts of interest to this work.

References

- [1] J. Miao, Y. Wei, Y. Wu, C. Liang, G. Li, Y. Yang. VSPW: A large-scale dataset for video scene parsing in the wild. In *Proceedings of IEEE/CVF Conference on*

- Computer Vision and Pattern Recognition*, Nashville, USA, pp. 4131–4141, 2021. DOI: [10.1109/CVPR46437.2021.00412](https://doi.org/10.1109/CVPR46437.2021.00412).
- [2] G. Sun, Y. Liu, H. Ding, T. Probst, L. Van Gool. Coarse-to-fine feature mining for video semantic segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, New Orleans, USA, pp. 3116–3127, 2022. DOI: [10.1109/CVPR52688.2022.00313](https://doi.org/10.1109/CVPR52688.2022.00313).
 - [3] G. Sun, Y. Liu, H. Ding, M. Wu, L. Van Gool. Learning local and global temporal contexts for video semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 10, pp. 6919–6934, 2024. DOI: [10.1109/TPAMI.2024.3387326](https://doi.org/10.1109/TPAMI.2024.3387326).
 - [4] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Vancouver, Canada, Article number 159, 2020.
 - [5] X. Zhai, A. Kolesnikov, N. Houlsby, L. Beyer. Scaling vision transformers. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, New Orleans, USA, pp. 1204–1213, 2022. DOI: [10.1109/CVPR52688.2022.01179](https://doi.org/10.1109/CVPR52688.2022.01179).
 - [6] M. Dehghani, J. Djolonga, B. Mustafa, P. Padlewski, J. Heek, J. Gilmer, A. P. Steiner, M. Caron, R. Geirhos, I. Alabdulmohsin, R. Jenatton, L. Beyer, M. Tschannen, A. Arnab, X. Wang, C. R. Ruiz, M. Minderer, J. Puigcerver, U. Evci, M. Kumar, S. Van Steenkiste, G. F. Elsayed, A. Mahendran, F. Yu, A. Oliver, F. Huot, J. Bastings, M. Collier, A. A. Gritsenko, V. Birodkar, C. N. Vasconcelos, Y. Tay, T. Mensink, A. Kolesnikov, F. Pavetic, D. Tran, T. Kipf, M. Lucic, X. Zhai, D. Keyzers, J. J. Harmsen, N. Houlsby. Scaling vision transformers to 22 billion parameters. In *Proceedings of the 40th International Conference on Machine Learning*, Honolulu, USA, pp. 7480–7512, 2023.
 - [7] Z. Liu, H. Hu, Y. Lin, Z. Yao, Z. Xie, Y. Wei, J. Ning, Y. Cao, Z. Zhang, L. Dong, F. Wei, B. Guo. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, New Orleans, USA, pp. 11999–12009, 2022. DOI: [10.1109/CVPR52688.2022.01170](https://doi.org/10.1109/CVPR52688.2022.01170).
 - [8] L. Wang, B. Huang, Z. Zhao, Z. Tong, Y. He, Y. Wang, Y. Wang, Y. Qiao. VideoMAE V2: Scaling video masked autoencoders with dual masking. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Vancouver, Canada, pp. 14549–14560, 2023. DOI: [10.1109/CVPR52729.2023.01398](https://doi.org/10.1109/CVPR52729.2023.01398).
 - [9] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, I. Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748–8763, 2021.
 - [10] J. Wu, Y. Jiang, Q. Liu, Z. Yuan, X. Bai, S. Bai. General object foundation model for images and videos at scale. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, USA, pp. 3783–3795, 2024. DOI: [10.1109/CVPR52733.2024.00363](https://doi.org/10.1109/CVPR52733.2024.00363).
 - [11] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W. Y. Lo, P. Dollár, R. Girshick. Segment anything. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, Paris, France, pp. 3992–4003, 2023. DOI: [10.1109/ICCV51070.2023.00371](https://doi.org/10.1109/ICCV51070.2023.00371).
 - [12] W. Ji, J. Li, Q. Bi, T. Liu, W. Li, L. Cheng. Segment anything is not always perfect: An investigation of SAM on different real-world applications. *Machine Intelligence Research*, vol. 21, no. 4, pp. 617–630, 2024. DOI: [10.1007/s11633-023-1385-0](https://doi.org/10.1007/s11633-023-1385-0).
 - [13] N. Ravi, V. Gabeur, Y. T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C. Y. Wu, R. B. Girshick, P. Dollár, C. Feichtenhofer. SAM 2: Segment anything in images and videos. In *Proceedings of the 13th International Conference on Learning Representations*, Singapore, 2025.
 - [14] L. Yang, Z. Zhao, H. Zhao. UniMatch V2: Pushing the limit of semi-supervised semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 47, no. 4, pp. 3031–3048, 2025. DOI: [10.1109/TPAMI.2025.3528453](https://doi.org/10.1109/TPAMI.2025.3528453).
 - [15] Y. Benigimim, S. Roy, S. Essid, V. Kalogeiton, S. Lathuilière. Collaborating foundation models for domain generalized semantic segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, USA, pp. 3108–3119, 2024. DOI: [10.1109/CVPR52733.2024.00300](https://doi.org/10.1109/CVPR52733.2024.00300).
 - [16] X. Xiong, Z. Wu, S. Tan, W. Li, F. Tang, Y. Chen, S. Li, J. Ma, G. Li. SAM2-UNet: Segment anything 2 makes strong encoder for natural and medical image segmentation, [Online], Available: <https://arxiv.org/abs/2408.08870>, 2024.
 - [17] J. Wu, W. Ji, Y. Liu, H. Fu, M. Xu, Y. Xu, Y. Jin. Medical SAM adapter: Adapting segment anything model for medical image segmentation, [Online], Available: <https://arxiv.org/abs/2304.12620>, 2023.
 - [18] Y. He, Z. Gao, Y. Li, Z. Wang. A lightweight multi-modality medical image semantic segmentation network base on the novel UNeXt and Wave-MLP. *Computerized Medical Imaging and Graphics*, vol. 111, Article number 102311, 2024. DOI: [10.1016/j.compmedimag.2023.102311](https://doi.org/10.1016/j.compmedimag.2023.102311).
 - [19] X. J. Wu, C. L. Liu. Hierarchical open-vocabulary part-object segmentation with knowledge-guided SAM. *Machine Intelligence Research*, vol. 23, no. 1, pp. 214–226, 2026. DOI: [10.1007/s11633-025-1610-0](https://doi.org/10.1007/s11633-025-1610-0).
 - [20] L. Chen, Y. Zhu, Q. Yang, X. Niu, K. Ding, S. Xiang. SAM-MI: A mask-injected framework for enhancing open-vocabulary semantic segmentation with SAM. *Machine Intelligence Research*, vol. 23, no. 3, pp. 616–632, 2026. DOI: [10.1007/s11633-025-1615-8](https://doi.org/10.1007/s11633-025-1615-8).
 - [21] L. C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected CRFs. In *Proceedings of the 3rd International Conference on Learning Representations*, San Diego, USA, 2015.
 - [22] S. A. S. Hesham, J. Z. Sui, G. E. S. Xin, L. C. Ping, L. Jiang. Pix2Pt map for transfer-based few-shot learning. In *Proceedings of IEEE International Conference on Multimedia & Expo Workshops*, Shenzhen, China, 2021. DOI: [10.1109/ICMEW53276.2021.9456002](https://doi.org/10.1109/ICMEW53276.2021.9456002).
 - [23] L. C. Chen, G. Papandreou, F. Schroff, H. Adam. Rethinking atrous convolution for semantic image segmentation, [Online], Available: <https://arxiv.org/abs/1706.05587>, 2017.

- [24] S. Liu, L. Qi, H. Qin, J. Shi, J. Jia. Path aggregation network for instance segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, USA, pp. 8759–8768, 2018. DOI: [10.1109/CVPR.2018.00913](https://doi.org/10.1109/CVPR.2018.00913).
- [25] D. Bolya, C. Zhou, F. Xiao, Y. J. Lee. YOLACT: Real-time instance segmentation. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, Seoul, Republic of Korea, pp. 9156–9165, 2019. DOI: [10.1109/ICCV.2019.00925](https://doi.org/10.1109/ICCV.2019.00925).
- [26] A. Kirillov, K. He, R. Girshick, C. Rother, P. Dollár. Panoptic segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Long Beach, USA, pp. 9396–9405, 2019. DOI: [10.1109/CVPR.2019.00963](https://doi.org/10.1109/CVPR.2019.00963).
- [27] S. Ainetter, F. Fraundorfer. End-to-end trainable deep neural network for robotic grasp detection and semantic segmentation from RGB. In *Proceedings of IEEE International Conference on Robotics and Automation*, Xi'an, China, pp. 13452–13458, 2021. DOI: [10.1109/ICRA48506.2021.9561398](https://doi.org/10.1109/ICRA48506.2021.9561398).
- [28] K. Li, W. Tao, L. Liu. Online semantic object segmentation for vision robot collected video. *IEEE Access*, vol. 7, pp. 107602–107615, 2019. DOI: [10.1109/ACCESS.2019.2933479](https://doi.org/10.1109/ACCESS.2019.2933479).
- [29] M. Siam, C. Jiang, S. Lu, L. Petrich, M. Gamal, M. El-hoseiny, M. Jagersand. Video object segmentation using teacher-student adaptation in a human robot interaction (HRI) setting. In *Proceedings of International Conference on Robotics and Automation*, Montreal, Canada, pp. 50–56, 2019. DOI: [10.1109/ICRA.2019.8794254](https://doi.org/10.1109/ICRA.2019.8794254).
- [30] S. Song, A. Zeng, J. Lee, T. Funkhouser. Grasping in the wild: Learning 6DoF closed-loop grasping from low-cost demonstrations. *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4978–4985, 2020. DOI: [10.1109/LRA.2020.3004787](https://doi.org/10.1109/LRA.2020.3004787).
- [31] V. Pradeep, R. Khemmar, L. Lecrosnier, Y. Duchemin, R. Rossi, B. Decoux. Self-supervised sidewalk perception using fast video semantic segmentation for robotic wheelchairs in smart mobility. *Sensors*, vol. 22, no. 14, Article number 5241, 2022. DOI: [10.3390/s22145241](https://doi.org/10.3390/s22145241).
- [32] F. Codevilla, M. Müller, A. López, V. Koltun, A. Dosovitskiy. End-to-end driving via conditional imitation learning. In *Proceedings of IEEE International Conference on Robotics and Automation*, Brisbane, Australia, pp. 4693–4700, 2018. DOI: [10.1109/ICRA.2018.8460487](https://doi.org/10.1109/ICRA.2018.8460487).
- [33] M. Siam, M. Gamal, M. Abdel-Razek, S. Yogamani, M. Jagersand, H. Zhang. A comparative study of real-time semantic segmentation for autonomous driving. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, Salt Lake City, USA, 2018.
- [34] K. Muhammad, T. Hussain, H. Ullah, J. Del Ser, M. Rezaei, N. Kumar, M. Hijji, P. Bellavista, V. H. C. de Albuquerque. Vision-based semantic segmentation in scene understanding for autonomous driving: Recent achievements, challenges, and outlooks. *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 12, pp. 22694–22715, 2022. DOI: [10.1109/TITS.2022.3207665](https://doi.org/10.1109/TITS.2022.3207665).
- [35] X. Xiao, Y. Zhao, F. Zhang, B. Luo, L. Yu, B. Chen, C. Yang. BASeg: Boundary aware semantic segmentation for autonomous driving. *Neural Networks*, vol. 157, pp. 460–470, 2023. DOI: [10.1016/j.neunet.2022.10.034](https://doi.org/10.1016/j.neunet.2022.10.034).
- [36] H. Guan, C. Song, Z. Zhang. LiDAR-camera cooperative semantic segmentation. *Machine Intelligence Research*, vol. 22, no. 5, pp. 956–968, 2025. DOI: [10.1007/s11633-024-1508-2](https://doi.org/10.1007/s11633-024-1508-2).
- [37] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, O. Ronneberger. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In *Proceedings of the 19th International Conference on Medical Image Computing and Computer-Assisted Intervention*, Athens, Greece, pp. 424–432, 2016. DOI: [10.1007/978-3-319-46723-8_49](https://doi.org/10.1007/978-3-319-46723-8_49).
- [38] O. Ronneberger, P. Fischer, T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Proceedings of the 18th International Conference on Medical Image Computing and Computer-Assisted Intervention*, Munich, Germany, pp. 234–241, 2015. DOI: [10.1007/978-3-319-24574-4_28](https://doi.org/10.1007/978-3-319-24574-4_28).
- [39] M. Grammatikopoulou, R. Sanchez-Matilla, F. Bragman, D. Owen, L. Culshaw, K. Kerr, D. Stoyanov, I. Luengo. A spatio-temporal network for video semantic segmentation in surgical videos. *International Journal of Computer Assisted Radiology and Surgery*, vol. 19, no. 2, pp. 375–382, 2024. DOI: [10.1007/s11548-023-02971-6](https://doi.org/10.1007/s11548-023-02971-6).
- [40] H. Wu, J. Liu, F. Xiao, Z. Wen, L. Cheng, J. Qin. Semi-supervised segmentation of echocardiography videos via noise-resilient spatiotemporal semantic calibration and fusion. *Medical Image Analysis*, vol. 78, Article number 102397, 2022. DOI: [10.1016/j.media.2022.102397](https://doi.org/10.1016/j.media.2022.102397).
- [41] X. Li, C. Cui, S. Shi, H. Fei, Y. Hu. Semi-supervised echocardiography video segmentation via adaptive spatio-temporal tensor semantic awareness and memory flow. *IEEE Transactions on Medical Imaging*, vol. 44, no. 5, pp. 2182–2193, 2025. DOI: [10.1109/TMI.2025.3526955](https://doi.org/10.1109/TMI.2025.3526955).
- [42] B. Wang, L. Li, Y. Nakashima, R. Kawasaki, H. Nagahara, Y. Yagi. Noisy-LSTM: Improving temporal awareness for video semantic segmentation. *IEEE Access*, vol. 9, pp. 46810–46820, 2021. DOI: [10.1109/ACCESS.2021.3067928](https://doi.org/10.1109/ACCESS.2021.3067928).
- [43] J. Long, E. Shelhamer, T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Boston, USA, pp. 3431–3440, 2015. DOI: [10.1109/CVPR.2015.7298965](https://doi.org/10.1109/CVPR.2015.7298965).
- [44] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, L. C. Chen. MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, USA, pp. 4510–4520, 2018. DOI: [10.1109/CVPR.2018.00474](https://doi.org/10.1109/CVPR.2018.00474).
- [45] L. C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A. L. Yuille. DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, 2018. DOI: [10.1109/TPAMI.2017.2699184](https://doi.org/10.1109/TPAMI.2017.2699184).
- [46] H. Zhao, J. Shi, X. Qi, X. Wang, J. Jia. Pyramid scene parsing network. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, USA, pp. 6230–6239, 2017. DOI: [10.1109/CVPR.2017.660](https://doi.org/10.1109/CVPR.2017.660).
- [47] Q. Yan, S. Li, C. Liu, M. Liu, Q. Chen. RoboSeg: Real-time semantic segmentation on computationally constrained robots. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 52, no. 3, pp. 1567–1577, 2022. DOI: [10.1109/TSMC.2020.3032437](https://doi.org/10.1109/TSMC.2020.3032437).
- [48] Z. Qiu, T. Yao, T. Mei. Learning deep spatio-temporal

- dependence for semantic video segmentation. *IEEE Transactions on Multimedia*, vol. 20, no. 4, pp. 939–949, 2018. DOI: [10.1109/TMM.2017.2759504](https://doi.org/10.1109/TMM.2017.2759504).
- [49] G. Sun, Y. Liu, H. Tang, A. Chhatkuli, L. Zhang, L. Van Gool. Mining relations among cross-frame affinities for video semantic segmentation. In *Proceedings of the 17th European Conference on Computer Vision*, Tel Aviv, Israel, pp. 522–539, 2022. DOI: [10.1007/978-3-031-19830-4_30](https://doi.org/10.1007/978-3-031-19830-4_30).
- [50] Y. Li, J. Shi, D. Lin. Low-latency video semantic segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, USA, pp. 5997–6005, 2018. DOI: [10.1109/CVPR.2018.00628](https://doi.org/10.1109/CVPR.2018.00628).
- [51] M. Paul, C. Mayer, L. Van Gool, R. Timofte. Efficient video semantic segmentation with labels propagation and refinement. In *Proceedings of IEEE Winter Conference on Applications of Computer Vision*, Snowmass, USA, pp. 2862–2871, 2020. DOI: [10.1109/WACV45572.2020.9093520](https://doi.org/10.1109/WACV45572.2020.9093520).
- [52] P. Hu, F. Caba, O. Wang, Z. Lin, S. Sclaroff, F. Perazzi. Temporally distributed networks for fast video semantic segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, USA, pp. 8815–8824, 2020. DOI: [10.1109/CVPR42600.2020.00884](https://doi.org/10.1109/CVPR42600.2020.00884).
- [53] E. Shelhamer, K. Rakelly, J. Hoffman, T. Darrell. Clockwork convnets for video semantic segmentation. In *Proceedings of European Conference on Computer Vision*, Amsterdam, The Netherlands, pp. 852–868, 2016. DOI: [10.1007/978-3-319-49409-8_69](https://doi.org/10.1007/978-3-319-49409-8_69).
- [54] Y. Zhu, K. Sapra, F. A. Reda, K. J. Shih, S. Newsam, A. Tao, B. Catanzaro. Improving semantic segmentation via video propagation and label relaxation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Long Beach, USA, pp. 8848–8857, 2019. DOI: [10.1109/CVPR.2019.00906](https://doi.org/10.1109/CVPR.2019.00906).
- [55] Y. X. Li, C. L. Z. Chen, S. Li, A. M. Hao, H. Qin. A novel divide and conquer solution for long-term video salient object detection. *Machine Intelligence Research*, vol. 21, no. 4, pp. 684–703, 2024. DOI: [10.1007/s11633-023-1388-x](https://doi.org/10.1007/s11633-023-1388-x).
- [56] T. Zhou, F. Porikli, D. J. Crandall, L. Van Gool, W. Wang. A survey on deep learning technique for video segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 6, pp. 7099–7122, 2023. DOI: [10.1109/TPAMI.2022.3225573](https://doi.org/10.1109/TPAMI.2022.3225573).
- [57] M. Gao, F. Zheng, J. J. Q. Yu, C. Shan, G. Ding, J. Han. Deep learning for video object segmentation: A review. *Artificial Intelligence Review*, vol. 56, no. 1, pp. 457–531, 2023. DOI: [10.1007/s10462-022-10176-7](https://doi.org/10.1007/s10462-022-10176-7).
- [58] F. Pan, X. Yin, S. Lee, A. Niu, S. Yoon, I. S. Kweon. MoDA: Leveraging motion priors from videos for advancing unsupervised domain adaptation in semantic segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, Seattle, USA, pp. 2649–2658, 2024. DOI: [10.1109/CVPRW63382.2024.00271](https://doi.org/10.1109/CVPRW63382.2024.00271).
- [59] D. Guo, D. P. Fan, T. Lu, C. Sakaridis, L. Van Gool. Vanishing-point-guided video semantic segmentation of driving scenes. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, USA, pp. 3544–3553, 2024. DOI: [10.1109/CVPR52733.2024.00340](https://doi.org/10.1109/CVPR52733.2024.00340).
- [60] H. Mai, R. Sun, Y. Wang, T. Zhang, F. Wu. Pay attention to target: Relation-aware temporal consistency for domain adaptive video semantic segmentation. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, Vancouver, Canada, pp. 4162–4170, 2024. DOI: [10.1609/aaai.v38i5.28211](https://doi.org/10.1609/aaai.v38i5.28211).
- [61] Y. Weng, M. Han, H. He, M. Li, L. Yao, X. Chang, B. Zhuang. Mask propagation for efficient video semantic segmentation. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, New Orleans, USA, Article number 314, 2024.
- [62] Y. Gao, Z. Wang, J. Zhuang, Y. Zhang, J. Li. Exploit domain-robust optical flow in domain adaptive video semantic segmentation. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, Washington DC, USA, pp. 641–649, 2023. DOI: [10.1609/aaai.v37i1.25140](https://doi.org/10.1609/aaai.v37i1.25140).
- [63] S. A. S. Hesham, Y. Liu, G. Sun, H. Ding, J. Yang, E. Konukoglu, X. Geng, X. Jiang. Exploiting temporal state space sharing for video semantic segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Nashville, USA, pp. 24211–24221, 2025. DOI: [10.1109/CVPR52734.2025.02255](https://doi.org/10.1109/CVPR52734.2025.02255).
- [64] S. Lee, J. Seo, K. Han, M. Choi, S. Im. CAVIS: Context-aware video instance segmentation. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, Honolulu, USA, pp. 4507–4517, 2025.
- [65] T. Zhang, X. Tian, Y. Wu, S. Ji, X. Wang, Y. Zhang, P. Wan. DVIS: Decoupled video instance segmentation framework. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, Paris, France, pp. 1282–1291, 2023. DOI: [10.1109/ICCV51070.2023.00124](https://doi.org/10.1109/ICCV51070.2023.00124).
- [66] T. Zhang, X. Tian, Y. Zhou, S. Ji, X. Wang, X. Tao, Y. Zhang, P. Wan, Z. Wang, Y. Wu. DVIS++: Improved decoupled framework for universal video segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 47, no. 7, pp. 5918–5929, 2025. DOI: [10.1109/TPAMI.2025.3552694](https://doi.org/10.1109/TPAMI.2025.3552694).
- [67] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. Chatterji, A. Chen, K. Creel, J. Q. Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon, J. Etchemendy, K. Ethayarajh, L. Fei-Fei, C. Finn, T. Gale, L. Gillespie, K. Goel, N. Goodman, S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D. E. Ho, J. Hong, K. Hsu, J. Huang, T. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani, O. Khattab, P. W. Koh, M. Krass, R. Krishna, R. Kudipudi, A. Kumar, F. Ladhak, M. Lee, T. Lee, J. Leskovec, I. Levent, X. L. Li, X. Li, T. Ma, A. Malik, C. D. Manning, S. Mirchandani, E. Mitchell, Z. Munyikwa, S. Nair, A. Narayan, D. Narayanan, B. Newman, A. Nie, J. C. Niebles, H. Nilforoshan, J. Nyarko, G. Ogut, L. Orr, I. Papadimitriou, J. S. Park, C. Piech, E. Portelance, C. Potts, A. Raghunathan, R. Reich, H. Ren, F. Rong, Y. Roohani, C. Ruiz, J. Ryan, C. Ré, D. Sadigh, S. Sagawa, K. Santhanam, A. Shih, K. Srinivasan, A. Tamkin, R. Taori, A. W. Thomas, F. Tramèr, R. E. Wang, W. Wang, B. Wu, J. Wu, Y. Wu, S. M. Xie, M. Yasunaga, J. You, M. Zaharia, M. Zhang, T. Zhang, X. Zhang, Y. Zhang, L. Zheng, K. Zhou, P. Liang. On the opportunities and risks of foundation models, [Online], Available: <https://arxiv.org/abs/2108.07258>, 2022.
- [68] J. Devlin, M. W. Chang, K. Lee, K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*,

- Minneapolis, USA, pp. 4171–4186, 2019. DOI: [10.18653/v1/N19-1423](https://doi.org/10.18653/v1/N19-1423).
- [69] OpenAI. GPT-4 technical report, [Online], Available: <https://arxiv.org/abs/2303.08774>, 2024.
- [70] OpenAI. GPT-4o system card, [Online], Available: <https://arxiv.org/abs/2410.21276>, 2024.
- [71] OpenAI. OpenAI o1 system card, [Online], Available: <https://arxiv.org/abs/2412.16720>, 2024.
- [72] C. Jia, Y. Yang, Y. Xia, Y. T. Chen, Z. Parekh, H. Pham, Q. Le, Y. H. Sung, Z. Li, T. Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 4904–4916, 2021.
- [73] J. Yang, M. Gao, Z. Li, S. Gao, F. Wang, F. Zheng. Track anything: Segment anything meets videos, [Online], Available: <https://arxiv.org/abs/2304.11968>, 2023.
- [74] Y. Cheng, L. Li, Y. Xu, X. Li, Z. Yang, W. Wang, Y. Yang. Segment and track anything, [Online], Available: <https://arxiv.org/abs/2305.06558>, 2023.
- [75] H. He, J. Zhang, M. Xu, J. Liu, B. Du, D. Tao. Scalable mask annotation for video text spotting, [Online], Available: <https://arxiv.org/abs/2305.01443>, 2023.
- [76] Z. Lu, Z. Xiao, J. Bai, Z. Xiong, X. Wang. Can SAM boost video super-resolution? [Online], Available: <https://arxiv.org/abs/2305.06524>, 2023.
- [77] T. Wang, L. Li, K. Lin, Y. Zhai, C. C. Lin, Z. Yang, H. Zhang, Z. Liu, L. Wang. Disco: Disentangled control for realistic human dance generation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, USA, pp. 9326–9336, 2024. DOI: [10.1109/CVPR52733.2024.00891](https://doi.org/10.1109/CVPR52733.2024.00891).
- [78] B. Qin, W. Ye, Q. Yu, S. Tang, Y. Zhuang. Dancing Avatar: Pose and text-guided human motion videos synthesis with image diffusion model, [Online], Available: <https://arxiv.org/abs/2308.07749>, 2023.
- [79] Y. Yang, X. Wu, T. He, H. Zhao, X. Liu. SAM3D: Segment anything in 3D scenes, [Online], Available: <https://arxiv.org/abs/2306.03908>, 2023.
- [80] S. Dong, F. Liu, G. Lin. Leveraging large-scale pre-trained vision foundation models for label-efficient 3D point cloud segmentation. In *Proceedings of the 13th International Conference on Image and Graphics*, Xuzhou, China, pp. 304–315, 2025. DOI: [10.1007/978-981-95-3393-0_25](https://doi.org/10.1007/978-981-95-3393-0_25).
- [81] Y. Lai, Z. Luo, Z. Yu. Detect any deepfakes: Segment anything meets face forgery detection and localization. In *Proceedings of the 17th Chinese Conference on Biometric Recognition*, Xuzhou, China, pp. 180–190, 2023. DOI: [10.1007/978-981-99-8565-4_18](https://doi.org/10.1007/978-981-99-8565-4_18).
- [82] J. Z. Wu, X. Li, D. Gao, Z. Dong, J. Bai, A. Singh, X. Xiang, Y. Li, Z. Huang, Y. Sun, R. He, F. Hu, J. Hu, H. Huang, H. Zhu, X. Cheng, J. Tang, M. Z. Shou, K. Keutzer, F. Iandola. CVPR 2023 text guided video editing competition, [Online], Available: <https://arxiv.org/abs/2310.16003>, 2023.
- [83] Y. Yin, Z. Fu, F. Yang, G. Lin. OR-NeRF: Object removing from 3D scenes guided by multiview segmentation with neural radiance fields, [Online], Available: <https://arxiv.org/abs/2305.10503>, 2023.
- [84] X. Zhao, W. Ding, Y. An, Y. Du, T. Yu, M. Li, M. Tang, J. Wang. Fast segment anything, [Online], Available: <https://arxiv.org/abs/2306.12156>, 2023.
- [85] W. H. Chu, A. W. Harley, P. Tokmakov, A. Dave, L. Guibas, K. Fragkiadaki. Zero-shot open-vocabulary tracking with large pre-trained models. In *Proceedings of IEEE International Conference on Robotics and Automation*, Yokohama, Japan, pp. 4916–4923, 2024. DOI: [10.1109/ICRA57147.2024.10611726](https://doi.org/10.1109/ICRA57147.2024.10611726).
- [86] S. Mo, Y. Tian. AV-SAM: Segment anything model meets audio-visual localization and segmentation, [Online], Available: <https://arxiv.org/abs/2305.01836>, 2023.
- [87] W. Yue, J. Zhang, K. Hu, Y. Xia, J. Luo, Z. Wang. SurgicalSAM: Efficient class promptable surgical instrument segmentation. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, Vancouver, Canada, pp. 6890–6898, 2024. DOI: [10.1609/aaai.v38i7.28514](https://doi.org/10.1609/aaai.v38i7.28514).
- [88] B. Benjdira, A. Koubaa, A. M. Ali. ROSGPT_Vision: Commanding robots using only language models' prompts, [Online], Available: <https://arxiv.org/abs/2308.11236>, 2023.
- [89] C. Fu, L. Yao, H. Zuo, G. Zheng, J. Pan. SAM-DA: UAV tracks anything at night with SAM-powered domain adaptation. In *Proceedings of International Conference on Advanced Robotics and Mechatronics*, Tokyo, Japan, pp. 31–38, 2024. DOI: [10.1109/ICARM62033.2024.10715901](https://doi.org/10.1109/ICARM62033.2024.10715901).
- [90] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, P. Luo. SegFormer: Simple and efficient design for semantic segmentation with transformers. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, Article number 924, 2021.
- [91] J. Miao, X. Wang, Y. Wu, W. Li, X. Zhang, Y. Wei, Y. Yang. Large-scale video panoptic segmentation in the wild: A benchmark. In *proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, New Orleans, USA, pp. 21001–21011, 2022. DOI: [10.1109/CVPR52688.2022.02036](https://doi.org/10.1109/CVPR52688.2022.02036).
- [92] T. Chen, A. Lu, L. Zhu, C. Ding, C. Yu, D. Ji, Z. Li, L. Sun, P. Mao, Y. Zang. SAM2-Adapter: Evaluating & adapting segment anything 2 in downstream tasks: Camouflage, shadow, medical image segmentation, and more, [Online], Available: <https://arxiv.org/abs/2408.04579>, 2024.
- [93] H. Zhong, S. Shen, K. Cai, Z. Wu, J. Yao, Y. Cheng, X. Li, X. Zhang, L. Song, Q. Hu. Serial low-rank adaptation of vision transformer. In *Proceedings of IEEE International Conference on Multimedia and Expo*, Nantes, France, 2025. DOI: [10.1109/ICME59968.2025.11209194](https://doi.org/10.1109/ICME59968.2025.11209194).
- [94] I. Curti, P. Z. Ramirez, A. Petrelli, L. Di Stefano. Multimodal SAM-adapter for semantic segmentation. *IEEE Access*, vol. 13, pp. 160438–160455, 2025. DOI: [10.1109/ACCESS.2025.3609640](https://doi.org/10.1109/ACCESS.2025.3609640).
- [95] K. Wang, D. Lin, C. Li, Z. Tu, B. Luo. Adapting segment anything model to multi-modal salient object detection with semantic feature fusion guidance, [Online], Available: <https://arxiv.org/abs/2408.15063>, 2024.
- [96] R. Chen, H. Su, S. Tang. NAS-LoRA: Empowering parameter-efficient fine-tuning for visual foundation models with searchable adaptation, [Online], Available: <https://arxiv.org/abs/2512.03499>, 2025.
- [97] P. Chen, L. Xie, X. Huo, X. Yu, X. Zhang, Y. Sun, Z. Han, Q. Tian. SAM-CP: Marrying SAM with composable prompts for versatile segmentation. In *Proceedings of the 13th International Conference on Learning Representations*, Singapore, 2025.




Syed Ariff Syed Hesham received the B.Sc. degree in electrical and electronic engineering from Nanyang Technological University (NTU), Singapore in 2023. He is currently a Ph.D. degree candidate in electrical and electronic engineering at NTU, Singapore, and is also an A*STAR Computing and Information Science (ACIS) research scholar working under

the Institute for Infocomm Research (I2R), Singapore.

His research interest is the intersection of deep learning and computer vision, with a particular focus on advancing effective and label-efficient visual perception models.

E-mail: syedhesh002@e.ntu.edu.sg

ORCID iD: 0009-0000-9589-3056



Yun Liu received the B. Eng. and Ph. D. degrees in computer science from Nankai University, China in 2016 and 2020, respectively. Then, he worked with Prof. Luc Van Gool as a postdoctoral scholar at the Computer Vision Lab, ETH Zurich, Switzerland. After that, he worked as a senior scientist at the Institute for Infocomm Research (I2R),

A*STAR, Singapore. Now, he is a professor at the College of Computer Science and the Academy for Advanced Interdisciplinary Studies, Nankai University, China.

His research interests include computer vision and machine learning (especially deep learning).

E-mail: liuyun@nankai.edu.cn (Corresponding author)

ORCID iD: 0000-0001-6143-0264



Guolei Sun received the M. Sc. degree in computer science from the King Abdullah University of Science and Technology (KAUST), Saudi Arabia in 2018, and the Ph.D. degree at ETH Zurich, Switzerland in 2024, in Prof. Luc Van Gool's Computer Vision Lab. He was a postdoctoral researcher at the Computer Vision Lab, ETH Zurich. From 2018 to

2019, he worked as a research engineer with the Inception Institute of Artificial Intelligence, UAE. He is currently a professor at the College of Computer Science and the Academy for Advanced Interdisciplinary Studies, Nankai University, China.

His research interests include deep learning for video understanding, semantic/instance segmentation, object counting, and weakly supervised learning.

E-mail: guolei.sun@nankai.edu.cn (Corresponding author)

ORCID iD: 0000-0001-8667-9656



Jing Yang received the Ph. D. degree in mechanical and electronic engineering from Guizhou University, China in 2020. From August 2018 to September 2019, he was awarded a scholarship by the China Scholarship Council (CSC) under the State Scholarship Fund to pursue his study with Oklahoma State University, USA, as a Joint Ph.D. student with the

School of Computer Science and Technology, where he joined the Guoliang Fan's Group. From October 2022 to October 2023, he was a visiting scholar studying in the team of Professor Minyi Guo (IEEE Fellow) from Shanghai Jiao Tong University, China. He is currently an assistant professor with the

State Key Laboratory of Public Big Data, Guizhou University, China.

His research interests include high-performance computing, task scheduling in various architectures, and open-domain visual learning.

E-mail: jyang23@gzu.edu.cn

ORCID iD: 0000-0003-1915-9487



Henghui Ding received the B. Eng. degree from Xi'an Jiaotong University, China in 2016, and the Ph.D. degree from Nanyang Technological University (NTU), Singapore in 2020. He was a research scientist at ByteDance, China, and a post-doctoral researcher at ETH and NTU. He is currently a tenure-track professor at Fudan University, China. He

serves as an Associate Editor for *IET Computer Vision and Visual Intelligence* and serves/has served as the Area Chair of CVPR'24, NeurIPS'24, ICLR'25, ACM MM'24, and BMVC'24 and an SPC Member of AAAI 2022–2025 and IJCAI 2023–2024.

His research interests include computer vision and machine learning.

E-mail: hhding@fudan.edu.cn

ORCID iD: 0000-0003-4868-6526



Xue Geng received the B. Eng. degree in computer science from Northeastern University, China in 2012, and the Ph. D. degree in computer science from National University of Singapore (NUS), Singapore in 2017. Currently, she is a senior scientist with the Institute for Infocomm Research (I2R), A*STAR, Singapore.

Her research interests include model compression and efficient machine learning.

E-mail: snownus@gmail.com

ORCID iD: 0000-0002-2594-9648



Xudong Jiang received the B. Eng. and M. Eng. degrees from the University of Electronic Science and Technology of China (UESTC), China in 1983 and 1986, respectively, and the Ph. D. degree from Helmut Schmidt University, Germany in 1997. From 1998 to 2004, he was with the Institute for Infocomm Research, A*STAR, Singapore, as a lead scientist,

and the head of the Biometrics Laboratory. He joined Nanyang Technological University (NTU), Singapore, as a faculty member in 2004, where he was the director of the Center for Information Security from 2005 to 2011. He is currently a professor with the School of Electrical and Electronics Engineering (EEE), NTU, and the director of the Centre for Information Sciences and Systems, Singapore. He has authored more than 200 articles, with over 60 articles in IEEE journals and 30 papers in top conferences such as CVPR/ICCV/ECCV/AAAI/ICLR. He has served as an Associate Editor for *IEEE Signal Processing Letters* and *IEEE Transactions on Image Processing*. Currently, he serves as the Senior Area Editor for *IEEE Transactions on Image Processing* and the Editor-in-Chief of *IET Biometrics*.

His research interests include computer vision, machine learning, pattern recognition, image processing, and biometrics.

E-mail: exdjiang@ntu.edu.sg

ORCID iD: 0000-0002-9104-2315

Citation: S. A. S. Hesham, Y. Liu, G. Sun, J. Yang, H. Ding, X. Geng, X. Jiang. Evaluating sam2 for video semantic segmentation. *Machine Intelligence Research*, vol.23, no.5, pp.998–1011, 2026. <https://doi.org/10.1007/s11633-026-1638-9>

Articles may interest you

Deep video harmonization by improving spatial-temporal consistency. *Machine Intelligence Research*, vol.21, no.1, pp.46-54, 2024.
DOI: [10.1007/s11633-023-1447-3](https://doi.org/10.1007/s11633-023-1447-3)

Sam-mi: a mask-injected framework for enhancing open-vocabulary semantic segmentation with sam. *Machine Intelligence Research*, vol.23, no.3, pp.616-632, 2026.
DOI: [10.1007/s11633-025-1615-8](https://doi.org/10.1007/s11633-025-1615-8)

Lidar-camera cooperative semantic segmentation. *Machine Intelligence Research*, vol.22, no.5, pp.956-968, 2025.
DOI: [10.1007/s11633-024-1508-2](https://doi.org/10.1007/s11633-024-1508-2)

Deep learning-based face video restoration technique: a survey. *Machine Intelligence Research*, vol.23, no.3, pp.545-564, 2026.
DOI: [10.1007/s11633-025-1623-x](https://doi.org/10.1007/s11633-025-1623-x)

Hierarchical open-vocabulary part-object segmentation with knowledge-guided sam. *Machine Intelligence Research*, vol.23, no.1, pp.214-226, 2026.
DOI: [10.1007/s11633-025-1610-0](https://doi.org/10.1007/s11633-025-1610-0)

Segment anything is not always perfect: an investigation of sam on different real-world applications. *Machine Intelligence Research*, vol.21, no.4, pp.617-630, 2024.
DOI: [10.1007/s11633-023-1385-0](https://doi.org/10.1007/s11633-023-1385-0)

Interpretable visual neural decoding with unsupervised semantic disentanglement. *Machine Intelligence Research*, vol.22, no.3, pp.553-570, 2025.
DOI: [10.1007/s11633-023-1484-y](https://doi.org/10.1007/s11633-023-1484-y)



WeChat: MIR



Twitter: MIR_Journal