From Generalization to Precision: Exploring SAM for Tool Segmentation in Surgical Environments

Kanyifeechukwu J. Oguine, Roger D. Soberanis-Muku, Nathan Drenkow, Mathias Unberath

Abstract

Purpose: Accurate tool segmentation is essential in computer-aided procedures. However, this task conveys challenges due to artifacts' presence and the limited training data in medical scenarios. Methods that generalize to unseen data represent an interesting venue, where zero-shot segmentation presents an option to account for data limitation. Initial exploratory works with the Segment Anything Model (SAM) show that bounding-box-based prompting presents notable zero-short generalization. However, point-based prompting leads to a degraded performance that further deteriorates under image corruption. We argue that SAM drastically over-segment images with high corruption levels, resulting in degraded performance when only a single segmentation mask is considered, while the combination of the masks overlapping the object of interest generates an accurate prediction. **Method:** We use SAM to generate the over-segmented prediction of endoscopic frames. Then, we employ the ground-truth tool mask to analyze the results of SAM when the best single mask is selected as prediction and when all the individual masks overlapping the object of interest are combined to obtain the final predicted mask. We analyze the Endovis18 and Endovis17 instrument segmentation datasets using synthetic corruptions of various strengths and an In-House dataset featuring counterfactually created real-world corruptions. **Results:** Combining the over-segmented masks contributes to improvements in the IoU. Furthermore, selecting the best single segmentation presents a competitive IoU score for clean images. **Conclusions:** Combined SAM predictions present improved results and robustness up to a certain corruption level. However, appropriate prompting strategies are fundamental for implementing these models in the medical domain.

1 Introduction

The automatic segmentation of instruments plays a crucial role in computer-assisted interventions, where applications such as surgical understanding, tool tracking, tool-tissue iteration analysis, and navigation can benefit from tool segmentation models. These applications require precise and robust tool segmentation algorithms capable of working under the challenging environment that surgical images impose, where smoke, motion blur, and brightness and intensity variations can be present during the procedure [1]. Over the years, numerous image segmentation models have been developed, gradually progressing from conventional methods like clustering, edge-based, and contour-based detection towards deep learning models such as U-Net [10] and SegNet [4]. Deep learning models have demonstrated exceptional performance in segmentation tasks, yet they require large datasets for good generalization and robustness to corruptions in medical imaging. The availability of adequate training medical datasets presents a limitation in the medical imaging analysis domain.

Recently, advancements in foundational models like Large Language Models (LLMs) [5, 9], and recently in the image domain with the Segment Anything Model (SAM) [8] have revolutionized our perception of model generalization. SAM has established a notable reputation for its exceptional performance and abstraction across diverse unseen datasets, leading to a promising avenue toward efficient surgical tool segmentation. However, it is known that the performance of segmentation algorithms can be compromised by factors commonly encountered in surgical scenarios, such as variations in lighting, environmental and acquisition conditions, and imaging hardware. While SAM has demonstrated exceptional performance in general real-world image segmentation, its applicability to surgical environments still needs to be explored.

Initial exploratory works with SAM [11] have shown that bounding-box-based prompting presents notable zero-short generalization. However, point-based prompting leads to a degraded performance that further deteriorates under image noise and corruption. In contrast, we argue that SAM can generate an accurate mask provided that all individual regions overlapping the object of interest are consolidated into a single mask. This can suggest that, while SAM drastically over-segment images with high corruption levels, resulting in degraded performance when only a single segmentation mask is considered, it still manages to accurately detect object boundaries while hallucinating additional ones. Under this perspective, this work evaluates the stability of SAM in tool segmentation under different levels of simulated and real image corruptions. We consider the scenario when the prediction is composed of a single segmentation mask and compare it with the case when predictions are composed of multiple sub-masks that overlap the region of interest (ROI).
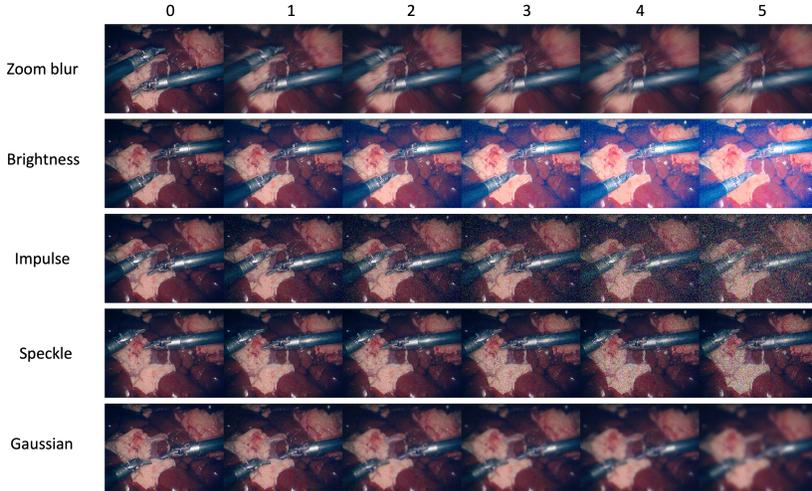
Figure 1: Examples of five types of corruptions applied to the images with their corresponding severity level. Images from our In-House dataset.

2 Method

To perform the evaluation, we employ SAM's automatic segmentation with default parameters to generate a grid-based over-segmented version of the image. We select the predictions that overlap the ROI defined by the tools' ground-truth mask to generate the tools' segmentation from the complete set of sub-masks. A sub-mask overlaps the ROI if its intersection ratio (i_r) with the ROI is more than 50 %. This ratio is defined by Eq. 1, where $M_{ROI}(x)$ represents the binary ground-truth mask, $M_s(x)$ a particular binary segmentation sub-mask obtained with SAM, and x is a pixel position.

$$i_r = \frac{\sum_x M_{ROI}(x) \cdot M_s(x)}{\sum_x M_s(x)} \tag{1}$$

This will lead to a set $S = M_1, M_2, \ldots, M_n$ of n overlapping sub-masks. The predicted segmentation is generated considering two scenarios. The first scenario evaluates the robustness and stability of the best single sub-mask as defined by Eq. 2.

$$M_{single} = \arg\max_i \sum_x M_i(x), \ M_i \in S, \ i = 1, 2, \ldots, n \tag{2}$$

The second scenario combines all the overlapping sub-masks to generate the final prediction, using Eq. 3 where "+" defines the pixel-wise addition between the binary masks.

$$M_{comb} = M_1 + M_2 + \cdots + M_n \tag{3}$$

3
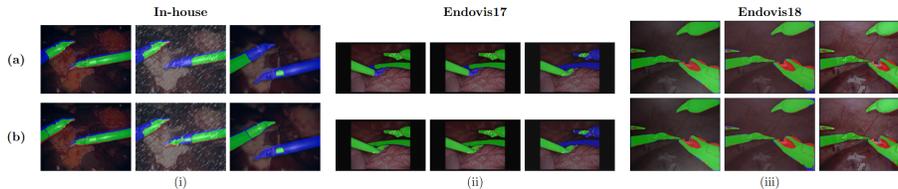
Figure 2: Overlay of (a) single and (b) combine prediction masks for different corruption types. True positives, false positives, and false negatives are indicated in green, red, and blue, respectively. Samples of the three datasets are presented.

Overall, the single mask segmentation can be interpreted as the segmentation obtained when a point-based prompt is placed in a representative section of the tool. Using a similar analogy, we can understand the combined segmentation resulting from placing multiple prompts in the tools.

We evaluate the performance of SAM on three endoscopic datasets, including EndoVis17 [2], EndoVis18 [3], and an In-House dataset. Both EndoVis datasets contain binocular frames. In our experiments, we use the left frames of the EndoVis17 and EndoVis18 training split, leading to 225 and 149 images for each dataset, respectively. Our In-House stereo endoscopic dataset [6] was collected using the da Vinci Research Kit (dVRK) [7] and includes 400 frames. Additionally, the In-House contains three counterfactually created real-world corruptions. Note that the In-House dataset does not include ground-truth tool masks. We generated the annotations using a semi-automatic approach, employing SAM to generate the over-segmented image and then manually correcting the segmentation masks. We apply 18 image corruptions to the images of each with five levels of severity. Fig. 1 present examples for five of these corruptions for reference. We then obtain the single and combined segmentation considering corruption/severity groups. We report the average intersection over union (IoU) per group in our experiments. Note that to test the zero-shot generalizability of the network, the segment anything model was used out-of-the-box with its predefined weights, with no additional fine-tuning to the datasets.

3   Results

Results for the single (s) and combined (c) mask are presented in Figs. 3, 4, and 5 for the In-House, EndoVis17, and EndoVis18 datasets, respectively. Similarly, the performance of the segmentation masks in the real perturbations is presented in Table 1. Additionally, Fig. 2 shows visual examples of segmented images in the different datasets. Results show that the overall performance of the SAM degrades with the severity of the corruption but at different levels, depending on the type of perturbation. For example, it is suggested that zoom blur (Fig. 6) causes a high degradation in the performance consistently across the datasets and predicted mask, even at the first level of severity, in contrast to brightness
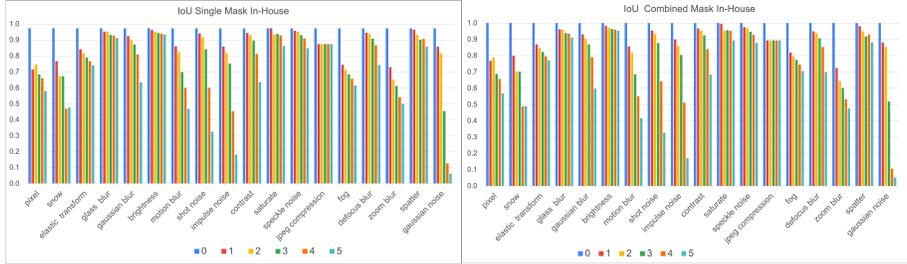
Figure 3: Average IoU for the 18 types of corruption, with five levels of severity for the Single v.s. Combined SAM segmentation mask in the In-House dataset.
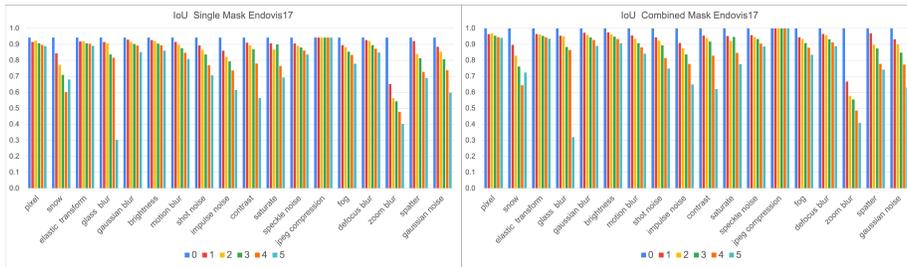
Figure 4: Average IoU for the 18 types of corruption, with five levels of severity for the Single v.s. Combined SAM segmentation mask in the EndoVis17 dataset.

(Fig. 7) or JPEG compression, where the performance is more stable for both single and combined masks.

When comparing the outcomes of the single prediction with the combined mask, we can observe a relative improvement in the IoU when employing the combined segmentations, especially in both EndoVis datasets. This suggests combining the combined masks captures additional details missed by the single mask prediction. For example, it can be noticed in Fig. 6 and Fig. 7 that the single mask fails to segment the tooltip in different cases, while by using a combined mask, it is possible to obtain these details of the tool.

Table 1: Average IoU for the three counterfactually generated corruption for the Single (s) v.s. Combined (c) SAM segmentation mask in the In-House dataset.

level	Mask	High Brightness	Low Brightness	Smoke
0	s	0.97	0.97	0.97
	c	0.99	0.99	0.99
1	s	0.94	0.61	0.72
	c	0.95	0.67	0.75

Regarding stability to the presence of perturbations, both masks present a degree of robustness to the severity of the disturbances. However, this is more
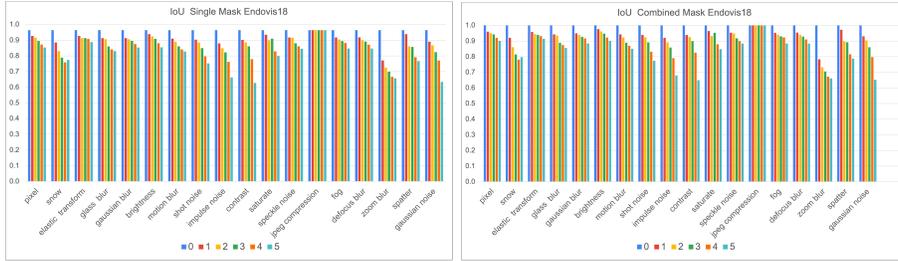
Figure 5: Average IoU for the 18 types of corruption, with five levels of severity for the Single v.s. Combined SAM segmentation mask in the EndoVis18 dataset.
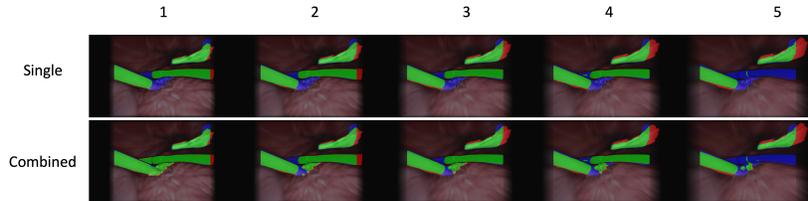
Figure 6: Frame segmentation visual performance in the presence of different levels of zoom blur. The figure presents the segmentation of the single and combined masks. True positives, false positives, and false negatives are indicated in green, red, and blue, respectively.

evident when evaluating the combined segmentation. In some cases, both masks present a high level of degradation in their performance. It is possible that, up to a specific type and level, the perturbations are strong enough to heavily modify the scene's visibility, leading to loss of information and the degradation of the results. For example, blur can reduce the contrast between the tool contour and the background, creating incorrect segmentations that contain a large portion of the tissue. However, it can be noticed in Fig. 6 that the use of a combined mask can reduce the impact of the noise and still recover sections of the tooltip.

4 Conclusions

The performance of SAM in the tool segmentation tasks can be affected by prompt selection. While a single mask can obtain adequate generalization and performance in this task, a prediction composed of multiple sub-masks can recover additional details of the tool, as suggested by an improvement in the IoU. This can indicate a tendency of SAM to over-segment the objects of interest and the necessity of defining adequate prompting strategies that lead to optimal results.

6
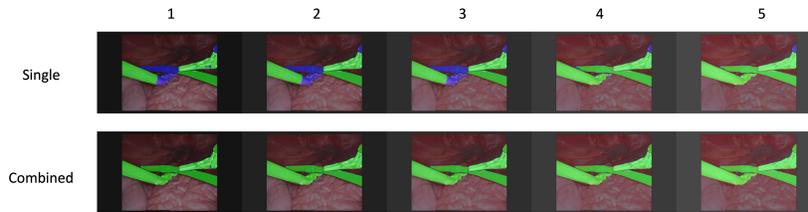
Figure 7: Frame segmentation visual performance in the presence of different levels of brightness. The figure presents the segmentation of the single and combined masks. True positives, false positives, and false negatives are indicated in green, red, and blue, respectively.

References

[1] ALI, S., ZHOU, F., BRADEN, B., BAILEY, A., YANG, S., CHENG, G., ZHANG, P., LI, X., KAYSER, M., SOBERANIS-MUKUL, R. D., ALBARQOUNI, S., WANG, X., WANG, C., WATANABE, S., OKSUZ, I., NING, Q., YANG, S., KHAN, M. A., GAO, X., REALDON, S., LOSHCHENOV, M., SCHNABEL, J., EAST, J., WAGNIERES, G., LOSCHENOV, V., GRISAN, E., DAUL, C., BLONDEL, W., AND RITTSCHER, J. An objective comparison of detection and segmentation algorithms for artefacts in clinical endoscopy. *Scientific Reports* (2020).

[2] ALLAN, M., KONDO, S., BODENSTEDT, S., LEGER, S., KADKHODAMOHAMMADI, R., LUENGO, I., FUENTES, F., FLOUTY, E., MOHAMMED, A., PEDERSEN, M., KORI, A., ALEX, V., KRISHNAMURTHI, G., RAUBER, D., MENDEL, R., PALM, C., BANO, S., SAIBRO, G., SHIH, C.-S., CHIANG, H.-A., ZHUANG, J., YANG, J., IGLOVIKOV, V., DOBRENKII, A., REDDIBOINA, M., REDDY, A., LIU, X., GAO, C., UNBERATH, M., KIM, M., KIM, C., KIM, C., KIM, H., LEE, G., ULLAH, I., LUNA, M., PARK, S. H., AZIZIAN, M., STOYANOV, D., MAIER-HEIN, L., AND SPEIDEL, S. 2018 Robotic Scene Segmentation Challenge. *arXiv* (Jan. 2020).

[3] ALLAN, M., SHVETS, A., KURMANN, T., ZHANG, Z., DUGGAL, R., SU, Y.-H., RIEKE, N., LAINA, I., KALAVAKONDA, N., BODENSTEDT, S., HERRERA, L., LI, W., IGLOVIKOV, V., LUO, H., YANG, J., STOYANOV, D., MAIER-HEIN, L., SPEIDEL, S., AND AZIZIAN, M. 2017 Robotic Instrument Segmentation Challenge. *arXiv* (Feb. 2019).

[4] BADRINARAYANAN, V., KENDALL, A., AND CIPOLLA, R. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell. 39*, 12 (Jan. 2017), 2481–2495.

7

[5] DEVLIN, J., CHANG, M.-W., LEE, K., AND TOUTANOVA, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *ACL Anthology* (June 2019), 4171–4186.

[6] DING, H., ZHANG, J., KAZANZIDES, P., WU, J. Y., AND UNBERATH, M. CaRTS: Causality-Driven Robot Tool Segmentation from Vision and Kinematics Data. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*. Springer, Cham, Switzerland, Sept. 2022, pp. 387–398.

[7] KAZANZIDES, P., CHEN, Z., DEGUET, A., FISCHER, G. S., TAYLOR, R. H., AND DIMAIO, S. P. An open-source research kit for the da Vinci® Surgical System. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2014, pp. 6434–6439.

[8] KIRILLOV, A., MINTUN, E., RAVI, N., MAO, H., ROLLAND, C., GUSTAFSON, L., XIAO, T., WHITEHEAD, S., BERG, A. C., LO, W.-Y., DOLLÁR, P., AND GIRSHICK, R. Segment Anything. *arXiv* (Apr. 2023).

[9] OPENAI. GPT-4 Technical Report. *arXiv* (Mar. 2023).

[10] RONNEBERGER, O., FISCHER, P., AND BROX, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Springer, Cham, Switzerland, Nov. 2015, pp. 234–241.

[11] WANG, A., ISLAM, M., XU, M., ZHANG, Y., AND REN, H. SAM Meets Robotic Surgery: An Empirical Study in Robustness Perspective. *arXiv* (Apr. 2023).