

# Assessing YOLACT++ for real time and robust instance segmentation of medical instruments in endoscopic procedures

Juan Carlos Ángeles Cerón<sup>1</sup>, Leonardo Chang<sup>1</sup>, Gilberto Ochoa Ruiz<sup>1</sup>, Sharib Ali<sup>2</sup>

**Abstract**—Image-based tracking of laparoscopic instruments plays a fundamental role in computer and robotic-assisted surgeries by aiding surgeons and increasing patient safety. Computer vision contests, such as the Robust Medical Instrument Segmentation (ROBUST-MIS) Challenge, seek to encourage the development of robust models for such purposes, providing large, diverse, and high-quality datasets. To date, most of the existing models for instance segmentation of medical instruments were based on two-stage detectors, which provide robust results but are nowhere near to the real-time, running at 5 frames-per-second (fps) at most. However, for the method to be clinically applicable, a real-time capability is utmost required along with high accuracy. In this paper, we propose the addition of attention mechanisms to the YOLACT architecture to allow real-time instance segmentation of instruments with improved accuracy on the ROBUST-MIS dataset. Our proposed approach achieves competitive performance compared to the winner of the 2019 ROBUST-MIS challenge in terms of robustness scores, obtaining 0.313 MI.DSC and 0.338 MI.NSD while reaching real-time performance at >45 fps.

**Keywords:** *Deep learning; laparoscopy; segmentation; surgical instrument, real-time.*

## I. INTRODUCTION

Computer-assisted minimally invasive surgery such as endoscopy has grown in popularity over the past years. However, due to the nature of these procedures, issues like limited field-of-view, extreme lighting conditions, lack of depth information, and difficulty in manipulating operating instruments demand strenuous amounts of effort from the surgeons [1]. Surgical data science applications could provide physicians with context-aware assistance during minimally invasive surgery to overcome these limitations and increase patient safety. One of the main forms of assistance is providing accurate tracking of medical instruments using instance segmentation methods. These systems are expected to be a crucial component in tasks ranging from surgical navigation, skill analysis, complication prediction, and other computer-integrated surgery (CIS) applications [2].

Nonetheless, instrument tracking methods are often deployed in difficult scenarios such as bleeding, over or under-exposure, smoke, and reflections [3]. The net effect of these issues increases the missed detection rates in endoscopic surveillance, hampering the adoption of AI-based tools in this context [4]. Therefore, the development of robust techniques that can be effectively deployed in real endoscopy interventions is very much necessary.

<sup>1</sup>Tecnológico de Monterrey, School of Engineering and Sciences, Av. Eugenio Garza Sada 2501, Monterrey, N.L., México 64849, Mexico

<sup>2</sup>Institute of Biomedical Engineering and Big Data Institute, Old Road Campus, University of Oxford, Oxford, UK

C. Authors: jcarlos.angelesc@gmail.com, gilberto.ochoa@tec.mx

Endoscopic computer vision contests, such as the Robust Medical Instrument Segmentation (ROBUST-MIS) Challenge [1] represent important and necessary efforts to encourage the development of robust models for surgical instrument segmentation. They integrate the developments in computer-assisted surgeries and benchmark the generalization capabilities of the developed methods on different clinical scenarios. Furthermore, they provide large-high-quality datasets to overcome one of the main bottlenecks of developing robust methodologies, which is the lack of annotated data.

Previous approaches for instance segmentation submitted to the 2019 ROBUST-MIS challenge, were exclusively based on two-stage detectors such as Mask R-CNN [5]. While these models exhibited good performances in terms of robustness, they all suffered from very high inference times averaging around 5 fps, preventing them from achieving real-time performances. Realistically, real-time performance is mandatory in order to fully exploit the capabilities of tracking applications in live surgeries.

In order to overcome these inference limitations while maintaining a robust performance in terms of tool segmentation results, we propose a new approach based on YOLACT++ [6] equipped with attention modules on the multi-scale outputs of the CNN backbone and Feature Pyramid Network (FPN). The increased representation power achieved by using attention allows the extraction of more discriminant features while suppressing the less effective ones.

In this work, we evaluate the Criss-cross Attention Module (CCAM) [7]. CCAM, which is depicted in Figure 1, recursively integrates global context across feature maps in a fast and clever criss-cross fashion. By integrating this attention mechanism, our proposed model outperforms previous approaches in the state-of-the-art by a slight margin, but it attains real-time performances, which makes our method clinically applicable both in inference time and robustness.

The rest of the paper is organized as follows. In Section II we discuss the ROBUST-MIS dataset, the context of the challenge, and the official metrics to assess robustness. Section III introduces our modifications to the YOLACT model to incorporate attention; also, we discuss our training and validation processes. In Section IV we present our results and we discuss areas for further improvements. Finally, Section V concludes the article.

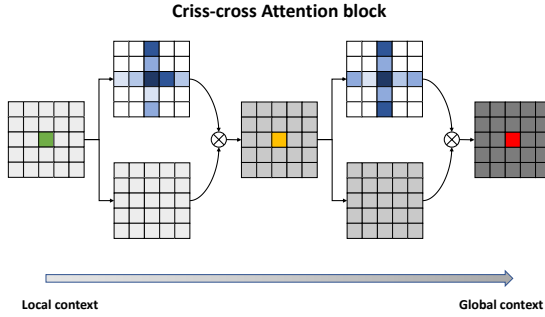


Fig. 1. Diagram of the criss-cross attention module. For each position in the input feature map, the attention module generates a sparse attention map by aggregating information in the horizontal and vertical axes. After the second operation, each pixel has extracted context from all the pixels.

TABLE I  
CASE DISTRIBUTION FOR EACH STAGE OF THE CHALLENGE [1].

Procedure	Training	Testing		
		Stage 1	Stage 2	Stage 3
Proctocolectomy	2,943	325	255	0
Rectal resection	3,040	338	289	0
Sigmoid resection*	0	0	0	2,880
TOTAL	5,983	663	514	2,880

## II. MATERIALS AND METHODS

### A. Materials

The Heidelberg Colorectal Data Set [2] served as a basis for the ROBUST-MIS challenge. It comprises 30 surgical procedures from three different types of surgeries and includes detailed segmentation maps for the surgical instruments in more than 10,000 laparoscopic video frames. The generalization and performance of the submitted algorithms are typically assessed in three stages, with increasing levels of difficulty:

- **Stage 1:** Test data is taken from the procedures from which the training data was extracted.
- **Stage 2:** Test data is taken from the same type of surgery as the training data but procedures not included in the training.
- **Stage 3:** Test data is taken from a different but similar type of surgery compared to the training data.

The detailed case distribution for each stage is presented in Table I.

### B. Metrics

The two multi-instance segmentation metrics were used to assess the performance of the models. Multiple Instance Dice Similarity Coefficient (MI\_DSC) and Multiple Instance Normalized Surface Dice (MI\_NSD). The DSC [8] is defined as the harmonic mean of precision and recall:

$$DSC(Y, \hat{Y}) := \frac{2 |Y \cap \hat{Y}|}{|Y| + |\hat{Y}|}, \quad (1)$$

Where  $Y$  indicates the ground truth annotation and  $\hat{Y}$  the corresponding prediction of an image frame.

Unlike DSC, which measures the overlap of volumes, the NSD measures the overlap of two mask borders [9]. The metric uses a threshold that is related to the inter-rater variability of the annotators. According to [1], their calculations resulted in a threshold of  $\tau := 13$  for the challenge's data set. To calculate the MI\_DSC and MI\_NSD, matches of instrument instances were computed. Then, the resulting metric scores per instrument instance per image were aggregated by the mean.

Note that the challenge reports robustness and accuracy rankings. However, to compute accuracy, it is mandatory to know the per image results per participant, which are not available due to privacy issues. For this reason, we will be reporting only robustness rankings.

The robustness rankings pay particular attention in stage 3 of the challenge since it was built to test generalization, and focus on the worst-case performance of methods. For this reason, MI\_DSC and MI\_NSD are aggregated by the 5% percentile instead of by the mean or median [1].

### C. Data preprocessing

A total of 996 frames with no visible instruments were removed from the training set, leaving 4,987 usable frames. An 85-15 percent split was made for training and validation purposes from this subset, respectively.

Data augmentation techniques were heavily applied to introduce as much variability as possible and increase the model's performance. The augmentation techniques used are random photometric distortions, random scaling, random sample cropping, and random mirroring.

## III. PROPOSED MODEL

### A. Architecture of the proposed model

In order to improve the robustness of the real-time YOLACT architecture used in our proposal, we introduce attention modules on the multi-scale outputs of the ResNet-101 backbone and the output features of the FPN (see Figure 2). Attention enables the network to focus on the most relevant features and avoid redundant use of information.

Our attention mechanism of choice was Criss-cross Attention Module (CCAM) [7], specifically because of its fast, computationally efficient ( $N\sqrt{N}$ ), and low GPU memory usage. These characteristics are crucial in order to introduce as little time-processing overhead as possible into the model and preserve real-time performance.

CCAM captures global contextual information in a similar fashion to the non-local module [10] by generating an attention map for each pixel in the input feature map and then transforming it into a refined feature map. However, instead of generating attention maps for each pixel-pair which is computationally expensive, for each pixel in a feature map, CCAM aggregates contextual information only in its horizontal and vertical directions, as shown in Figure 1. By consecutively stacking two criss-cross attention modules, each pixel can collect contextual information from

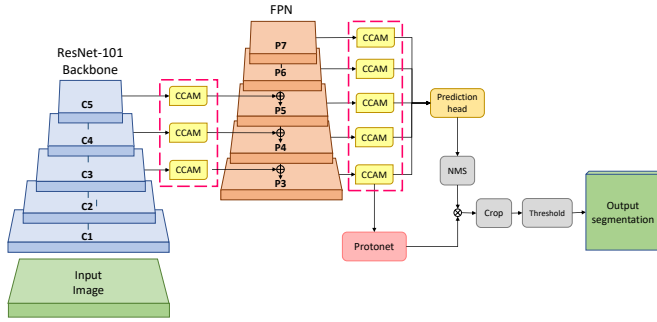


Fig. 2. Proposed YOLACT++ architecture with criss-cross attention modules in ResNet-101 backbone + FPN. Note that certain modules are removed in some experiments. See Section III for further details.

all pixels in a given feature map. Next, the contextually rich feature is concatenated and convolved with the original feature maps for feature fusion. Our experiments consisted of systematically adding the attention modules in two strategic locations in the network: first, in the backbone’s output features, next in the FPN output features, and ultimately on both locations. As a result, we came up with three models, *CCAM-Backbone*, *CCAM-FPN*, and *CCAM-Full*, plus the attentionless baseline *Base YOLACT++*.

### B. Training and validation

We trained the models for up to 100,000 iterations with a learning rate of 0.001, momentum of 0.9, weight decay of  $5 \times 10^{-4}$ , and batch size 16 in an NVIDIA DGX-1 system. The performance was assessed using the evaluation code for the challenge [11]. Moreover, the rankings were computed using the R package *challenger* [12].

## IV. RESULTS AND DISCUSSION

Figure 3 shows dot-and-boxplots of the metric values for each algorithm over all test cases in stage 3 of the challenge. Among the three model variations to which we added attention modules, *CCAM-Backbone* achieved the best results in terms of robustness. This result indicates that the contextually enriched feature maps from the ResNet-101 backbone are powerful enough to generate more accurate mask prototypes and coefficients in the YOLACT architecture and ultimately better segmentation outputs.

From our results, we can determine that adding attention mechanisms on the FPN outputs only increases the model performance slightly when compared to the baseline. However, this slight improvement becomes non-significant when considering the strict aggregated challenge metrics.

Considering the good performance of *CCAM-Backbone* and the minor improvements of *CCAM-FPN*, one might believe that combining both configurations in *CCAM-Full* would lead to better results than having only one of them. However, as we can observe in their respective boxplots, this is not the case. An explanation of this behavior is that adding too many attention modules that integrate global context

TABLE II

AGGREGATED EVALUATION PERFORMANCE FOR STAGE 3 OF THE CHALLENGE. THE TOP SECTION OF THE TABLE INCLUDES THE TEAMS FROM THE 2019 CHALLENGE. TEAM SCORES WERE TAKEN FROM [1]. THE BOTTOM SECTION INCLUDES OUR THREE ATTENTION MODELS PLUS THE BASELINE WITHOUT ATTENTION.

Team/Algorithm	MI_DSC	MI_NSD	FPS
<i>www</i> (Mask R-CNN)	0.31	0.35	5*
<i>Uniandes</i> (Mask R-CNN)	0.26	0.29	5*
<i>SQUASH</i> (Mask R-CNN)	0.22	0.26	5*
<i>CASIA_SRL</i> (Dense Pyramid Network)	0.19	0.27	5*
<i>fisensee</i> (2D U-Net)	0.17	0.16	18*
<i>caresyntax</i> (Mask R-CNN)	0.00	0.00	5*
<i>VIE</i> (Mask R-CNN)	0.00	0.00	5*
CCAM-Backbone	0.313	0.338	49
CCAM-Full	0.308	0.333	45
CCAM-FPN	0.000	0.000	60
Base YOLACT++	0.000	0.000	75

\*Approximated from base method.

Original measurement was not reported.

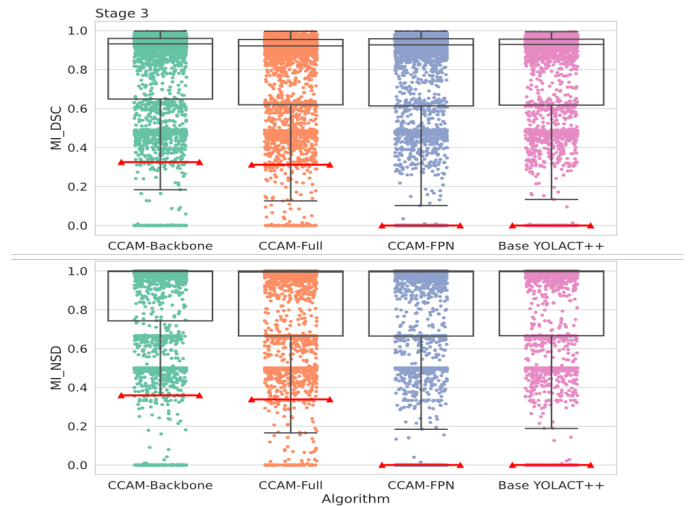


Fig. 3. Dot-and-boxplots showing the individual performance of algorithms in stage 3 of the challenge. Red lines indicate the aggregated metric by 5% percentile also reported in Table II.

might lead to an over-mixing of information and noise, commonly known as over-smoothing, a common problem in graph neural networks from which CCAM takes inspiration.

Next, we compare our proposed models and baseline to the top participants of the 2019 challenge (note that the 2020 edition did not take place). Table II shows the aggregated MI\_DSC and MI\_NSD values achieved for each participant/model, as well as the network architecture they were based on and their approximated/measured frame rate. Regarding the robustness of our method, *CCAM-Backbone* achieves competitive results in both metrics with respect to the top-performing contestant, reaching 0.313 MI\_DSC and 0.338 MI\_NSD compared to 0.31 and 0.35 respectively by team *www*.

It is important to note, though, that this team used data from the EndoVis 2017 challenge to improve their generalization performance on additional data. However, our

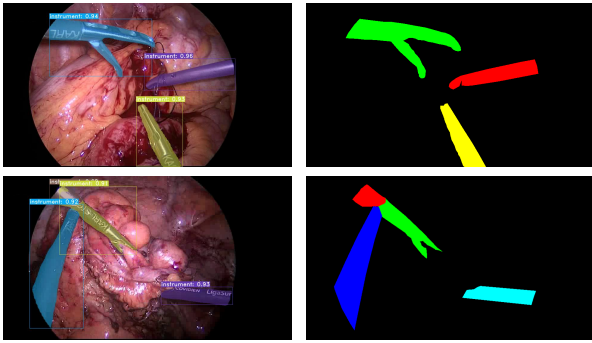


Fig. 4. Left: CCAM-Backbone evaluation results on stage 3 frames. Right: Ground truth annotations. All images have the confidence threshold set to 0.3.

best model outperforms the second-best contestant by a considerable margin: 0.053 MI<sub>DSC</sub> and 0.048 MI<sub>NSD</sub>.

An important contribution of our method is its ability to run in real-time. Inference speed performance was tested on a 10 second video snippet from the ROBUST-MIS dataset a total of ten times per model. The reported frame rates were then aggregated by the mean. Inference was tested on a single Tesla P100 GPU from the DGX-1 cluster with video multi-frame enabled.

As seen in Table II, the vast majority of the submitted models were based on Mask R-CNN, which is inevitably slow as it relies on a two-stage detector that limits its performance to 5 fps at most. In contrast, our models comfortably fulfill real-time operation requirements for clinically usable models, running at  $\geq 45$  fps.

Notably, our top model produces high-quality and temporally consistent masks. Figure 4 shows some examples with varying types and number of instruments together with their respective ground truth annotations. The model is robust to occluded instruments and various harsh conditions, like blood, smoke, and poor lighting. Nevertheless, it struggles with transparent instruments and small instruments on the edge of the field of view. Figure 5 illustrates some examples of challenging frames for our proposed algorithm, which we will seek to address in future work.

## V. CONCLUSION

We presented a novel approach for multi-instance segmentation of medical instruments based on the YOLACT architecture extended with embedded criss-cross attention modules. The addition of attention made it possible to extract better global context and exploit salient features leading to improved segmentation results. Our best model yielded competitive results in terms of robustness compared to the state-of-the-art, reaching 0.313 on area-based metric MI<sub>DSC</sub> and 0.338 on distance-based metric MI<sub>NSD</sub> while attaining real-time performance. Our experiments showed that adding attention modules to YOLACT boosts the performance of the model and increases robustness. However, since CCAM is based on graph neural networks, it can potentially create feature clusters that can cause over-smoothing of the learned

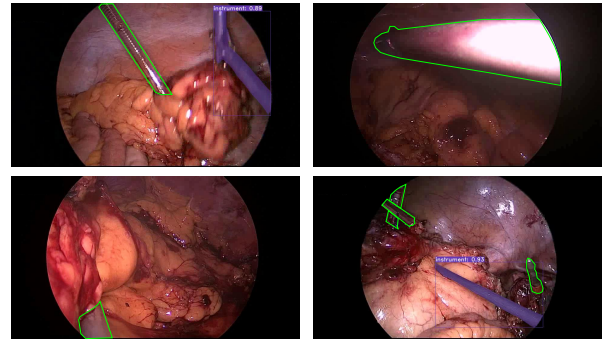


Fig. 5. Examples of challenging frames, green contours indicate instrument instances not detected by our model. From left to right and top to bottom: 1. transparent instrument, 2. large reflection on instrument, 3. small instrument on the edge of the field of view, and 4. multiple instruments, some partially occluded and on the edge of the field of view.

features. Thus, embedding attention modules at each layer may hinder the performance resulting from the over-mixing of global information and noise.

In future work, we plan to experiment with different types of attention mechanisms besides CCAM. Moreover, to increase the robustness of future models, we believe that stronger data augmentation aimed towards challenging instances could improve the model performances.

## ACKNOWLEDGMENTS

The authors wish to thank the AI Hub and the CIOT at ITESM for their support for carrying the experiments reported in this paper on their NVIDIA's DGX computer.

## REFERENCES

- [1] Tobias Ross *et al.*, Comparative validation of multi-instance instrument segmentation in endoscopy: Results of the ROBUST-MIS 2019 challenge, 2021; Medical Image Analysis, Vol. 70, May, 2021
- [2] Lena Maier-Hein *et al.*, Heidelberg Colorectal Data Set for Surgical Data Science in the Sensor Operating Room, 2020; arXiv:2005.03501.
- [3] Sebastian Bodenstedt *et al.*, Comparative evaluation of instrument segmentation and tracking methods in minimally invasive surgery, 2018; arXiv:1805.02475.
- [4] Sharib Ali *et al.*, Deep learning for detection and segmentation of artefact and disease instances in gastrointestinal endoscopy. Challenge Report Medical Image Analysis, Vol. 70, May, 2021
- [5] Kaiming He, Georgia Gkioxari, Piotr Dollár and Ross Girshick. Mask R-CNN, 2017; arXiv:1703.06870.
- [6] Bolya, Daniel and Zhou, Chong and Xiao, Fanyi and Lee, Yong Jae. YOLACT++: Better Real-time Instance Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1,(2020). Institute of Electrical and Electronics Engineers (IEEE).
- [7] Zilong Huang, Xinggang Wang, Yunchao Wei, Lichao Huang, Humphrey Shi, Wenyu Liu and Thomas S. Huang. CCNet: Criss-Cross Attention for Semantic Segmentation, 2018; arXiv:1811.11721.
- [8] Lee R. Dice. Measures of the Amount of Ecologic Association Between Species. *Ecology*, 26(3), pp. 297–302,(1945). Ecological Society of America.
- [9] Stanislav Nikolov *et al.*, Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy, 2018; arXiv:1809.04430.
- [10] Xiaolong Wang, Ross Girshick, Abhinav Gupta and Kaiming He. Non-local Neural Networks, 2017; arXiv:1711.07971.
- [11] Tobias Roß and Annika Reinke. RobustMIS2019, 2019. <https://phabricator.mitk.org/source/rmis2019/>.
- [12] Manuel Wiesenfarth *et al.*, Methods and open-source toolkit for analyzing and visualizing challenge results, 2021; Scientific Reports, 11(1), pp. 1–15,(2021). Nature Publishing Group.