Polyp and Surgical Instrument Segmentation with Double Encoder-Decoder Networks

Adrian Galdran¹

1. BCN MedTech, Departament de Tecnologies de la Informació i les Comunicacions, Pompeu Fabra University, Barcelona, Spain. E-mail any correspondence to: adrian.galdran@upf.edu

Abstract

This paper describes a solution for the MedAI competition, in which participants were required to segment both polyps and surgical instruments from endoscopic images. Our approach relies on a double encoder-decoder neural network which we have previously applied for polyp segmentation, but with a series of enhancements: a more powerful encoder architecture, an improved optimization procedure, and the post-processing of segmentations based on tempered model ensembling. Experimental results show that our method produces segmentations that show a good agreement with manual delineations provided by medical experts.

Keywords: endoscopic image analysis; surgical instrument segmentation; polyp segmentation.

Introduction

Examination of the human gastrointestinal tract is key for the early detection and treatment of different diseases like Colorectal Cancer [1]. This is usually achieved by a colonoscopic screening, a procedure in which a flexible tube equipped with a camera is introduced through the rectum to look for, and potentially remove lesions throughout the colon. In this context, automated endoscopic image analysis and decision support systems based on artificial intelligence have shown great potential for improving examination effectiveness and facilitating medical intervention [2, 3].

This paper describes our solution to the MedAI medical image segmentation competition [4], a public contest in which participants were asked to solve two independent tasks, namely polyps and surgical instrument extraction from endoscopic images. For this, we employed a deep neural network architecture based upon previous work of us that won the Endotect 2021 challenge [5, 6].

Here we adopt several improvements over our previous polyp segmentation solution, and apply it also for surgical instrument segmentation, as described in the next section. We also experiment with temperature scaling and analyze its impact on the resulting performance metrics.

Materials and methods

Semantic segmentation is often approached with encoder-decoder networks [7] that produce pixel-wise probabilities. The encoder acts as a feature extractor, downsampling spatial resolutions and increasing the number of channels by learning convolutional filters. The decoder upsamples this compressed representation back to the original input size. Double encoder-decoders are a direct extension in which two encoder-decoder networks are sequentially combined [8]. With x an input RGB image, $E^{(1)}$ the first network, and $E^{(2)}$ the second network, in a double encoder-decoder the output $E^{(1)}(x)$ of the first network is fed to the second network together with x . This allows $E^{(2)}$ to focus on the most interesting parts of the image:

$$E(x) = E^{(2)}(x, E^{(1)}(x)), \quad (1)$$

where x and $E^{(1)}(x)$ are stacked so the input to $E^{(2)}$ has four channels, as illustrated in Fig. 1. Here we employ the same structure in both $E^{(1)}$ and $E^{(2)}$: a Feature Pyramid Network architecture as the decoder [9], and a Resnext101 as pretrained decoder [10]. We also utilize SAM [11] wrapping ADAM as the optimization procedure, and rotate the train/validation data in a 4-fold manner, resulting in four models E_1, E_2, E_3, E_4 . We then apply temperature sharpening on the resulting ensemble [12]:

$$p = \frac{1}{4} \sum_{i=1}^4 E_i(x)^t \quad (2)$$

with t a hyper-parameter, to obtain our final predictions.

Figure 1: Graphical description of the adopted double encoder-decoder architecture for polyp and instrument segmentation.

Temperature →	T=0.5			T=1			T=2		
	DICE	PRECISION	RECALL	DICE	PRECISION	RECALL	DICE	PRECISION	RECALL
Polyps	88.59	87.76	94.56	89.65	92.42	90.09	88.69	93.37	87.56
Instruments	94.94	92.08	98.84	96.18	95.06	97.88	96.35	96.32	96.92

Table 1: Performance analysis on the test set, provided by the organizers, in terms of Dice score, Precision, and Recall, for different Temperature sharpening values, for the tasks of polyp segmentation (top row) and instrument segmentation (bottom row). Best temperature value for each metric marked bold.

Experimental Results

The main metric of interest in medical image segmentation tasks is usually the Dice score, given by:

$$D(f_{\theta}(x), y) = \frac{2 \cdot |f_{\theta}(x) \cap y|}{|f_{\theta}(x)| + |y|}, \quad (3)$$

where $f_{\theta}(x)$ is the model's binary prediction for example x and y is its corresponding manual ground-truth. If we denote True Positives as TP, False Negatives as FN, etc., we see that in the numerator of eq. (3), $f_{\theta}(x) \cap y$ corresponds to TP. On the other hand, the denominator describes the amount of pixels that are predicted as foreground ($|f_{\theta}(x)| = TP + FP$) and the amount of pixels that actually belong to the foreground ($|y| = TP + FN$). Therefore we can re-write the Dice score as $D = 2TP / (2TP + FP + FN)$, and realize that True Negatives have no impact on this metric. This means that correctly predicting small objects will be rewarded much more than correctly producing good segmentations of large objects. For a more comprehensive evaluation we also report in Table 1 Precision and Recall, defined as:

$$Precision = \frac{TP}{TP + FP}, \quad Recall = \frac{TP}{TP + FN}. \quad (4)$$

Achieving a high precision can be considered akin to a good Dice score; recall also encompasses False Negatives and captures performance on segmenting large objects.

Table 1 shows final performance on the hidden test set, as provided by the organizers, for the two tasks of Polyp and surgical instrument segmentation after training on the Kvasir-SEG [13] and Kvasir-Instrument [14] datasets.

Discussion of the Results and Future Work

We conclude several observations from the above results. First, instrument segmentation is relatively easier than polyp extraction, due to the greater variability of the latter regarding color and appearance. Also, in both cases we achieve a Dice score close to or above 90%, a good overall scoring. Last, temperature has a noticeable impact in performance: when a model attains greater Dice score, it also achieves higher Precision but lower Recall. It can be seen that lower temperature values lead to higher recall at the expense of a decrease in Dice and Precision, whereas higher temperatures result in greater Dice/Precision, but lower Recall. Therefore, Temperature sharpening is a reasonable mechanism to deal with the Precision/Recall trade-off in segmentation tasks.

Although the training data contained images in which there was always a polyp or a surgical instrument present, the test set had frames with no object of interest to be segmented, which is an example of out-of-distribution data. Future research may involve the investigation of different approaches to better handling this kind of problem for the task of endoscopic image segmentation.

Experimental Reproducibility and Code Availability

An implementation of the above method can be found at https://github.com/agaldran/medai_challenge, together with details like hyper-parameter specifications.

Acknowledgments

The author is grateful to *datacrunch.io* for donating GPU time for this project.

References

1. Hagggar FA and Boushey RP. Colorectal Cancer Epidemiology: Incidence, Mortality, Survival, and Risk Factors. *Clinics in Colon and Rectal Surgery* 2009 Nov; 22:191–7. [Accessed on: 2020 Oct 9]
2. Lui TK, Hui CK, Tsui VW, Cheung KS, Ko MK, Foo D aCC, Mak LY, Yeung CK, Lui TH, Wong SY, and Leung WK. New insights on missed colonic lesions during colonoscopy through artificial intelligence–assisted real-time detection (with video). en. *Gastrointestinal Endoscopy* 2020 May. [Accessed on: 2020 Sep 24]
3. Jha D, Ali S, Tomar NK, Riegler MA, Johansen D, Johansen HD, and Halvorsen P. Exploring Deep Learning Methods for Real-Time Surgical Instrument Segmentation in Laparoscopy. *2021 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI)*. IEEE. 2021 :1–4
4. Hicks S, Jha D, Thambawita V, Riegler M, Halvorsen P, Singstad B, Gaur S, Pettersen K, Goodwin M, Parasa S, and Lange T de. MedAI: Transparency in Medical Image Segmentation. *Nordic Machine Intelligence* 2021
5. Galdran A, Carneiro G, and Ballester MAG. Double Encoder-Decoder Networks for Gastrointestinal Polyp Segmentation. *International Conference on Pattern Recognition*. Springer. 2021 :293–307
6. Hicks SA, Jha D, Thambawita V, Halvorsen P, Hammer HL, and Riegler MA. The EndoTect 2020 Challenge: Evaluation and Comparison of Classification, Segmentation and Inference Time for Endoscopy. en. *Pattern Recognition. ICPR International Workshops and Challenges*. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2021 :263–74. DOI: 10.1007/978-3-030-68793-9_18
7. Ronneberger O, Fischer P, and Brox T. U-Net: Convolutional Networks for Biomedical Image Segmentation. en. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2015 :234–41
8. Galdran A, Anjos A, Dolz J, Chakor H, Lombaert H, and Ayed IB. The Little W-Net That Could: State-of-the-Art Retinal Vessel Segmentation with Minimalistic Models. *arXiv:2009.01907* 2020 Sep. [Accessed on: 2020 Sep 24]
9. Lin TY, Dollár P, Girshick R, He K, Hariharan B, and Belongie S. Feature Pyramid Networks for Object Detection. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. ISSN: 1063-6919. 2017 Jul :936–44
10. Kolesnikov A, Beyer L, Zhai X, Puigcerver J, Yung J, Gelly S, and Houlsby N. Big transfer (bit): General visual representation learning. *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*. Springer. 2020 :491–507
11. Foret P, Kleiner A, Mobahi H, and Neyshabur B. Sharpness-aware Minimization for Efficiently Improving Generalization. *International Conference on Learning Representations*. 2021. Available from: <https://openreview.net/forum?id=6Tm1mposlRM>
12. Berthelot D, Carlini N, Goodfellow I, Papernot N, Oliver A, and Raffel CA. MixMatch: A Holistic Approach to Semi-Supervised Learning. *Advances in Neural Information Processing Systems*. Ed. by Wallach H, Larochelle H, Beygelzimer A, d’Alché-Buc F, Fox E, and Garnett R. Vol. 32. Curran Associates, Inc., 2019. Available from: <https://proceedings.neurips.cc/paper/2019/file/1cd138d0499a68f4bb72bee04bbec2d7-Paper.pdf>
13. Jha D, Smedsrud PH, Riegler MA, Halvorsen P, Lange T de, Johansen D, and Johansen HD. Kvasir-seg: A segmented polyp dataset. *International Conference on Multimedia Modeling*. Springer. 2020 :451–62
14. Jha D, Ali S, Emanuelsen K, Hicks SA, Thambawita V, Garcia-Ceja E, Riegler MA, Lange T de, Schmidt PT, Johansen HD, Johansen D, and Halvorsen P. Kvasir-Instrument: Diagnostic and Therapeutic Tool Segmentation Dataset in Gastrointestinal Endoscopy. *MultiMedia Modeling*. 2021 :218–29