

A Benchmark for Endoluminal Scene Segmentation of Colonoscopy Images

David Vázquez^{1,2} · Jorge Bernal¹ ·
F. Javier Sánchez¹ · Gloria
Fernández-Esparrach⁴ · Antonio M.
López^{1,2} · Adriana Romero² · Michal
Drozdal^{3,5} · Aaron Courville²

Received: 20 May 2016 / Accepted: date

Abstract Colorectal cancer (CRC) is the third cause of cancer death worldwide. Currently, the standard approach to reduce CRC-related mortality is to perform regular screening in search for polyps and colonoscopy is the screening tool of choice. The main limitations of this screening procedure are polyp miss-rate and inability to perform visual assessment of polyp malignancy. These drawbacks can be reduced by designing Decision Support Systems (DSS) aiming to help clinicians in the different stages of the procedure by providing endoluminal scene segmentation. Thus, in this paper, we introduce an extended benchmark of colonoscopy image segmentation, with the hope of establishing a *new strong benchmark for colonoscopy image analysis research*. The proposed dataset consists of 4 relevant classes to inspect the endoluminal scene, targeting different clinical needs. Together with the dataset and taking advantage of advances in semantic segmentation literature, we provide new baselines by training standard fully convolutional networks (FCN). We perform a comparative study to show that FCN *significantly outperform, without any further post-processing*, prior results in endoluminal scene segmentation, especially with respect to polyp segmentation and localization.

Keywords Colonoscopy · Polyp · Semantic Segmentation · Deep Learning

David Vázquez

E-mail: dvazquez@cvc.uab.es

¹ Computer Vision Center, Computer Science Department, Universitat Autònoma de Barcelona, Spain

² Montreal Institute for Learning Algorithms, Université de Montréal, Canada

³ École Polytechnique de Montréal, Montréal, Canada

⁴ Endoscopy Unit, Gastroenterology Service, CIBERHED, IDIBAPS, Hospital Clinic, Universidad de Barcelona, Spain

⁵ Imagia Inc., Montréal, Canada

1 Introduction

Colorectal cancer (CRC) is the third cause of cancer death worldwide [33]. CRC arises from adenomatous polyps (adenomas) which are initially benign; however, with time some of them can become malignant. Currently, the standard approach to reduce CRC-related mortality is to perform regular screening in search for polyps and colonoscopy is the screening tool of choice. During the examination, clinicians visually inspect intestinal wall (see Figure 1(a) for an example of intestinal scene) in search of polyps. Once detected, they are resected and sent for histological analysis to determine their degree of malignancy and define the corresponding treatment the patient should undertake.

The main limitations of colonoscopy are its associated polyp miss-rate (small/flat polyps or the ones hidden behind intestine folds can be missed [20]) and the fact that polyp's malignancy degree is only known after histological analysis. These drawbacks can be reduced by developing new colonoscopy modalities to improve visualization (e.g. High Definition imaging, Narrow Band Imaging (NBI) [23] and magnification endoscopes [9]) and/or by developing Decision Support Systems (DSS) aiming to help clinicians in the different stages of the procedure. A clinically useful DSS should be able to detect, segment and assess the malignancy degree (e.g. by optical biopsy [31]) of polyps during the colonoscopy procedure, following a similar pipeline to the one shown in Figure 1(b).

The development of DSS for colonoscopy has been an active research topic during the last decades. The majority of available works on optical colonoscopy are focused on polyp detection (e.g. see [4, 32, 17, 26, 28, 36]) and only few works address the problems of endoluminal scene segmentation.

Endoluminal scene segmentation is of crucial relevance for clinical applications [3, 4, 25, 2]. *Polyp* segmentation is important to define the area covered by a potential lesion that should be carefully inspected and possibly removed by clinicians. Moreover, having a system for accurate in-vivo prediction of polyp histology might significantly improve clinical workflow. *Lumen* segmentation is relevant to help clinicians navigate through the colon during the procedure. Additionally, it can be used to establish quality metrics related to the degree of colon wall that has been explored, since a weak exploration can lead to polyp overlooking. Finally, *specular highlights* have proven to be useful in reducing polyp detection false positive ratio in the context of hand-crafted methods [6].

In recent years, convolutional neural networks (CNNs) have become a *de facto* standard in computer vision, achieving state-of-the-art performance in tasks such as image classification, object detection and semantic segmentation; and making traditional methods based on hand-crafted features obsolete. Two major components in this groundbreaking progress were the availability of increased computational power (GPUs) and the introduction of large labeled datasets [13, 21]. Despite the additional difficulty of having limited amounts of labeled data, CNNs have successfully been applied to a variety of medical imaging tasks, by resorting to aggressive data augmentation techniques [30, 14]. More precisely, CNNs have excelled at semantic segmentation tasks in

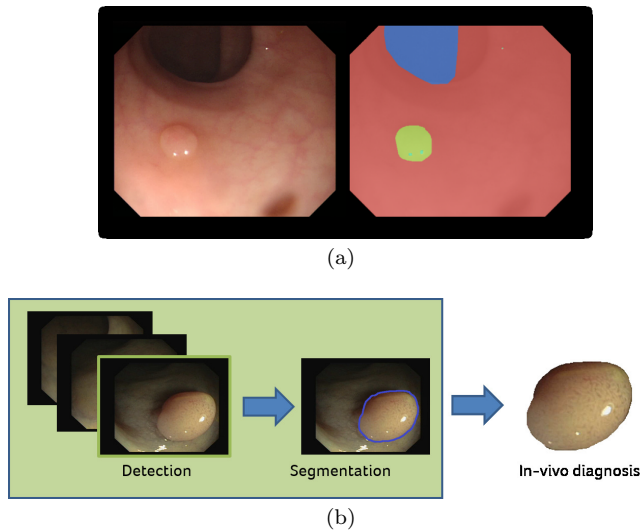


Fig. 1 (a) Colonoscopy image and corresponding labeling: blue for lumen, red for background (mucosa wall), and green for polyp, (b) Proposed pipeline of a decision support system for colonoscopy.

medical imaging, such as the EM ISBI 2012 dataset [1], BRATS [24] or MS lesions [35], where the top entries are built on CNNs [30, 11, 18, 8, 14]. Surprisingly, to the best of our knowledge, CNNs have not been applied to semantic segmentation of colonoscopy data. We associate this to the lack of large publicly available annotated databases, which is needed in order to train and validate such networks.

In this paper, we aim to overcome this limitation and introduce an extended benchmark of colonoscopy images by combining the two largest datasets of colonoscopy images [5, 4] and incorporating additional annotations to segment lumen and specular highlights, with the hope of establishing a *new strong benchmark for colonoscopy image analysis research*. We provide new baselines on this dataset by training standard fully convolutional networks (FCN) for semantic segmentation [22], and *significantly outperforming, without any further post-processing*, prior results in endoluminal scene segmentation.

Therefore, the contributions of this paper are two-fold:

1. Extended benchmark for colonoscopy image segmentation.
2. New state-of-the-art in colonoscopy image segmentation.

The rest of the paper is organized as follows. In Section 2, we present the new extended benchmark, including the introduction of datasets as well as the performance metrics. After that, in Section 3, we introduce the FCN architecture used as baseline for the new endoluminal scene segmentation benchmark. Then, in Section 4, we show qualitative and quantitative experimental results. Finally, Section 5 concludes the paper.

Table 1 Summary of prior database content. All frames show at least one polyp.

Database	# patients	# seq.	# frames	Resolution	Annotations
CVC-ColonDB	13	13	300	500x574	polyp lumen
CVC-ClinicDB	23	31	612	384x288	polyp
CVC-EndoSceneStill	36	44	912	500x574 & 384x288	polyp lumen background specularity border (void)

2 Endoluminal scene segmentation benchmark

In this section, we describe the endoluminal scene segmentation benchmark, including evaluation metrics.

2.1 Dataset

Inspired by already published benchmarks for polyp detection, proposed within a challenge held in conjunction with MICCAI 2015¹ [7], we introduce a benchmark for endoluminal scene object segmentation.

We combine **CVC-ColonDB** and **CVC-ClinicDB** into a new dataset (**CVC-EndoSceneStill**) composed of 912 images obtained from 44 video sequences acquired from 36 patients.

- **CVC-ColonDB** contains 300 images with associated polyp and background (here, mucosa and lumen) segmentation masks obtained from 13 polyp video sequences acquired from 13 patients.
- **CVC-ClinicDB** contains 612 images with associated polyp and background (here, mucosa and lumen) segmentation masks obtained from 31 polyp video sequences acquired from 23 patients.

We extend the old annotations to account for lumen, specular highlights with new hand-made pixel-wise annotations as well as we define a void class for black borders present in each frame. In the new annotations, background only contains mucosa (intestinal wall). Please refer to Table 1 for datasets details and to Figure 2 for a dataset sample.

We split the resulting dataset into three sets: training, validation and test containing 60%, 20% and 20% images respectively. We impose the constraint that one patient can not be in different sets. As a result, the final training set contains 20 patients and 547 frames, the validation set contains 8 patients and 183 frames; and test set contains 8 patients and 182 frames. The dataset will be available online upon paper acceptance².

¹ <http://endovis.grand-challenge.org>

² <http://www.cvc.uab.es/CVC-Colon/index.php/cvc-endoscenestill/>




Fig. 2 Example of a colonoscopy image and its corresponding ground truth: (a) Original image, (b) Polyp mask, (c) Specular highlights mask and (d) lumen mask.

2.2 Metrics

We use Intersection over Union (IoU), also known as *Jaccard index*, and per pixel accuracy as segmentation metrics. These metrics are commonly used in medical image segmentation tasks [10, 27].

We compute the mean of per-class IoU. Each per-class IoU is computed over a validation/test set according to following formula:

$$IoU(\text{PR}(\text{class}), \text{GT}(\text{class})) = \frac{|\text{PR}(\text{class}) \cap \text{GT}(\text{class})|}{|\text{PR}(\text{class}) \cup \text{GT}(\text{class})|}, \quad (1)$$

where PR represents the binary mask produced by the segmentation method, GT represents the ground truth mask, \cap represents set intersection and \cup represents set union.

We compute the mean global accuracy for each set as follows:

$$Acc(\text{PR}, \text{GT}) = \frac{\#TP}{\#\text{pixels}}, \quad (2)$$

where TP represents the number of true positives.

Notably, this new benchmark might as well be used for the relevant task of polyp localization. In that case, we follow Pascal VOC challenge metrics [16] and determine that a polyp is localized if it has a high overlap degree with its associated ground truth, namely:

$$IoU(\text{PR}(\text{polyp}), \text{GT}(\text{polyp})) > 0.5, \quad (3)$$

where the metric is computed for each polyp independently, and averaged per set to give a final score.

3 Baseline

CNN are a standard architecture used for tasks, where a single prediction per input is expected (e.g. image classification). Such architectures capture hierarchical representations of the input data by stacking blocks of convolutional, non-linearity and pooling layers on top of each other. Convolutional layers

extract local features. Non-linearity layers allow deep networks to learn non-linear mappings of the input data. Pooling layers reduce the spatial resolution of the representation maps by aggregating local statistics.

FCNs [22, 30] were introduced in the computer vision and medical imaging communities in the context of semantic segmentation. FCNs naturally extend CNNs to tackle per pixel prediction problems, by adding upsampling layers to recover the spatial resolution of the input at the output layer. As a consequence, FCNs can process images of arbitrary size. In order to compensate for the resolution loss induced by pooling layers, FCNs introduce skip connections between their downsampling and upsampling paths. Skip connections help the upsampling path recover fine-grained information from the downsampling layers.

We implemented FCN8 architecture from [22] and trained the network by means of stochastic gradient descent with rmsprop adaptive learning rate [38]. The validation split is used to early-stop the training, we monitor mean IoU for validation set and use patience of 50. We used a mini-batch size of 10 images. The input image is normalized in the range 0-1. We randomly crop the training images to 224×224 pixels. As regularization, we use dropout[34] of 0.5, as mentioned in the paper [22]. We do not use any weight decay.

As described in Section 2.1, colonoscopy images have a black border that we consider as a void class. Void classes do not influence the computation of the loss nor the metrics of any set, since the pixels marked as void class are ignored. As the number of pixels per class is unbalanced, in some experiments, we apply the median frequency balancing of [15].

During training, we experiment with data augmentation techniques such as random cropping, rotations, zooming, shearing and elastic transformations.

4 Experimental results

In this section, we report semantic segmentation and polyp localization results on the new benchmark.

4.1 Endoluminal Scene Semantic Segmentation

In this section, we first analyze the influence of different data augmentation techniques. Second, we evaluate the effect of having different number of endoluminal classes on polyp segmentation results. Finally, we compare our results with previously published methods.

4.1.1 Influence of data augmentation

Table 2 presents an analysis on the influence of different data augmentation techniques and their impact on the validation performance. We evaluate random zoom from 0.9 to 1.1, rotations from 0 to 180 degrees, shearing from

Data Augmentation	IoU background	IoU polyp	IoU lumen	IoU spec.	IoU mean	Acc mean
None	88.93	44.45	54.02	25.54	57.88	92.48
Zoom	89.89	52.73	51.15	37.10	57.72	90.72
Warp	90.00	54.00	49.69	37.27	58.97	90.93
Shear	89.60	46.61	54.27	36.86	56.83	90.49
Rotation	90.52	52.83	56.39	35.81	58.89	91.38
Combination	92.62	54.82	55.08	35.75	59.57	93.02

Table 2 FCN8 endoluminal scene semantic segmentation results for different data augmentation techniques. The results are reported on validation set.

0 to 0.4 and warping with σ ranging from 0 to 10. Finally, we evaluate the combination of all the data augmentation techniques.

As shown in the table, polyps significantly benefit from all data augmentation methods, in particular, from warping. Note that warping applies small elastic deformation locally, accounting for many realistic variations in the polyp shape. Rotation and zoom also have a strong positive impact on the polyp segmentation performance. It goes without saying that such transformations are the least aggressive ones, since they do not alter the polyp appearance. Shearing is most likely the most aggressive transformation, since it changes the polyp appearance and might, in some cases, result in unrealistic deformations.

While for lumen it is difficult to draw any strong conclusions, it looks like zooming and warping slightly deteriorate the performance, whereas shearing and rotation slightly improve it. As for specular highlights, all the data augmentation techniques that we tested significantly boost the segmentation results. Finally, background (mucosa) shows only slight improvement when incorporating data augmentations. This is not surprising, given its predominance throughout the data it could be even considered background.

Overall, combining all the discussed data augmentation techniques, leads to better results in terms of mean IoU and mean global accuracy. More precisely, we increase the mean IoU by 4.51% and the global mean accuracy by 1.52%

4.1.2 Influence of number of classes

Table 3 presents endoluminal scene semantic segmentation results for different number of classes. As shown in the table, using more underrepresented classes such as lumen or specular highlights makes the optimization problem more difficult. As expected and contrary to hand-crafted segmentation methods, when considering polyp segmentation, deep learning based approaches do not suffer from specular highlights, showing the robustness of the learnt features towards saturation zones in colonoscopy images.

Best results for polyp segmentation are obtained in the 2 classes scenario (polyp vs background). However, segmenting lumen is a relevant clinical problem as mentioned in Section 1. Results achieved in the 3 classes scenario are

# classes	IoU background	IoU polyp	IoU lumen	IoU spec.	IoU mean	Acc mean
4	92.07	39.37	59.55	40.52	57.88	92.48
3	92.19	50.70	56.48	–	66.46	92.82
2	96.63	56.07	–	–	76.35	96.77

Table 3 FCN8 endoluminal scene semantic segmentation results for different number of classes. The results are reported on validation set. In all cases, we selected the model that provided best validation results (with or without class balancing).

	Data Augmentation	IoU background	IoU polyp	IoU lumen	IoU spec.	IoU mean	Acc mean
<i>FCN8 performance</i>							
4 classes	None	86.36	38.51	43.97	32.98	50.46	87.40
3 classes	None	84.66	47.55	36.93	–	56.38	86.08
2 classes	None	94.62	50.85	–	–	72.74	94.91
4 classes	Combination	88.81	51.60	41.21	38.87	55.13	89.69
<i>State-of-the-art methods</i>							
[3, 2, 6]	-	73.93	22.13	23.82	44.86	41.19	75.58

Table 4 Results on the tests set: FCN8 with respect to previously published methods.

very encouraging, with a IoU higher than 50% for both polyp and lumen classes.

4.1.3 Comparison to state-of-the-art

Finally, we evaluate the FCN model on the test set. We compare our results to the combination of previously published hand-crafted methods: 1) [3] an energy-map based method for polyp segmentation, 2) [2] a watershed-based method for lumen segmentation and 3) [6] for specular highlights segmentation.

The segmentation results on the test set are reported in Table 4 and show a clear improvement of FCN8 over previously published methods. The following improvements can be observed when comparing previously published methods to 4-class FCN8 model trained with data augmentation: 15% in IoU for background (mucosa), 29% in IoU for polyps, 18% in IoU for lumen, 14% in mean IoU and 14% in mean accuracy. FCN8 is still outperformed by traditional methods when it comes to specular highlights class. However, it is important to note that specular highlights class is used by hand-crafted methods to reduce false positive ratio of polyp detection and from our analysis it looks like FCN model is able to segment well polyps even when ignoring this class. For example, the best mean IoU of 72.74 % and mean accuracy of 94.91% are obtained by 2-class model without additional data augmentation.

Figure 3 shows qualitative results of 4-class FCN8 model trained with data augmentation. From left to right, each row shows a colonoscopy frame, followed by the corresponding ground truth annotation and FCN8 prediction. Rows 1 to 4 show correct segmentation masks, with very clean polyp segmentation.

Rows 5 and 6 show failure modes of the model, where polyps have been missed or under-segmented. In row 5, the small polyp is missed by our segmentation method while, in row 6, the polyp is under-segmented. All cases exhibit decent lumen segmentation and good background (mucosa) segmentation.

4.2 Polyp localization

Endoluminal scene segmentation can be seen as a proxy to proper polyp detection in colonoscopy video. In order to understand how well suited FCN is to localize polyps, we perform a last experiment. In this experiment, we compute the polyp localization rate as a function of IoU between the model prediction and the ground truth. We can compute this IoU per frame, since our dataset contains a maximum of one polyp per image. This analysis describes the ability of a given method to cope with polyp appearance variability and stability on polyp localization.

The localization results are presented in Figure 4 and show a significant improvement when comparing FCN8 variants to the previously published method [3]. For example, when considering a correct polyp localization to have at least 50% IoU, we observe an increase of 40% in polyp localization rate. As a general trend, we observe that architectures trained using fewer number of classes achieve a higher IoU, though the polyp localization difference starts to be more visible when really high overlapping degrees are imposed. Finally, as one would expect, we observe that the architectures that show better results in polyp segmentation are the ones that show better results in polyp localization.

4.3 Towards clinical applicability

Subsections 4.1.3 and 4.2 presented results of a comparative study between FCNs and previous state-of-the-art of endoluminal scene objects segmentation in colonoscopy images. As mentioned in Section 1, we foresee several clinical applications, which can be built from the results of endoluminal scene segmentation. However, in order to be deployed in the exploration room, they must comply with real-time constraints apart from offering a good segmentation performance. In this case, and considering videos recorded at 25 frames per second, a DSS should not take more than 40 ms to process an image in order not to delay the procedure.

Considering this, we have computed processing times for each of the approaches studied in this paper. Results are presented in Table 5.

As shown in the table, none of the presented approaches currently meet real-time constraints. Running the FCN-8 inference on an NVIDIA Titan X GPU takes 88ms per frame. Note that this could easily be addressed by taking advantage of recent research on model compression [29] by applying fancier FCN architectures that encourage feature reuse [19]. Alternatively, we could

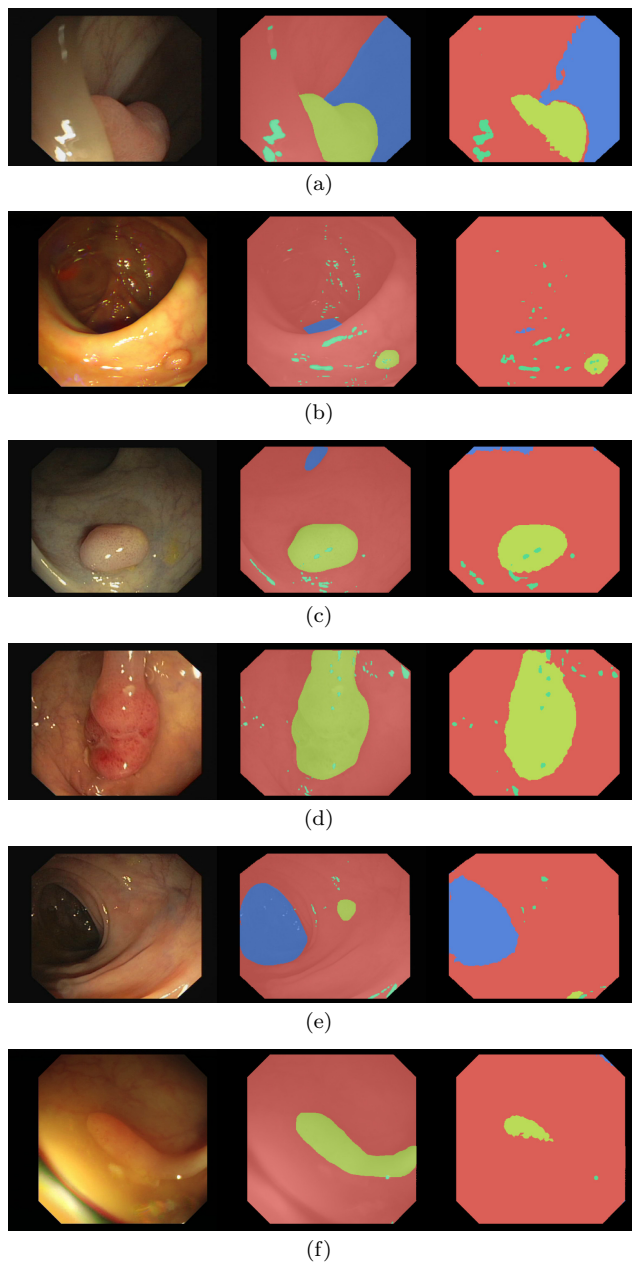


Fig. 3 Examples of predictions for 4-class FCN8 model. Each sub-figure represents a single frame, a ground truth annotation and a prediction image. We use the following color coding in the annotations: red for background (mucosa), blue for lumen, yellow for polyp and green specularity. (a, b, c, d) show correct polyp segmentation, whereas (e, d) show incorrect polyp segmentation.

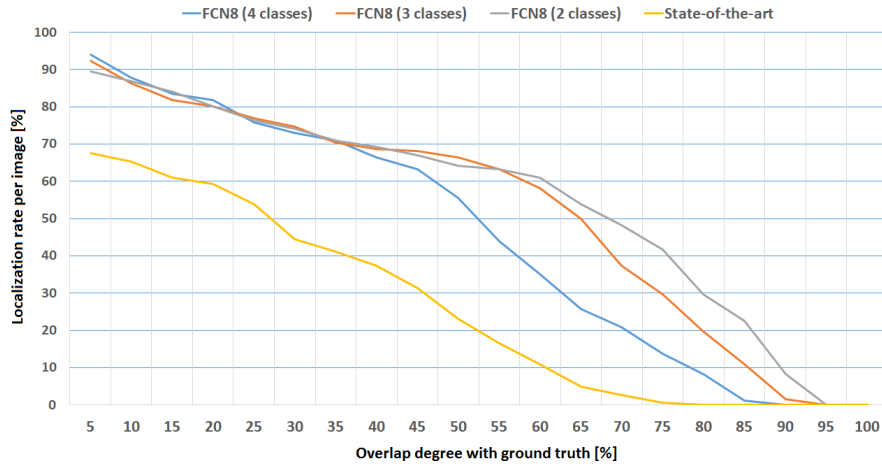


Fig. 4 Localization rate of polyps as a function of IoU. The x-axis represents the degree of overlap between ground truth and model prediction. The y-axis represents the percentage of correctly localized polyps. Different color plots represent different models: FCN8 with 4 classes, FCN8 with 3 classes, FCN8 with 2 classes and previously published method [3] (referred to as state-of-the-art in the plot)

Method	Polyp	Lumen	Specular highlights	Background
FCN	88ms*	88ms*	88ms*	88ms*
State-of-the-art	10000ms	8000ms	5000ms	23000ms

Table 5 Summary of processing times achieved by the different methods studied in the paper. FCN results are the same for all four classes considered as segmentation of the four classes is done at the same time.

exploit the temporal component and build more sophisticated architectures that would take advantage of the similarities among consecutive frames.

Clearly, hand-crafted methods take much longer to process one image. Moreover, they need to apply different methods to segment each class of interest, making them less clinically useful. Note that this is not the case for FCN-like architectures.

Despite computational constraints, FCNs’ superior performance could lead to more reliable and impactful computer-assisted clinical applications, since they offer both a better performance and computational efficiency.

5 Conclusions

In this paper, we have introduced an extended benchmark for endoluminal scene semantic segmentation. The benchmark includes extended annotations of polyps, background (mucosa), lumen and specular highlights. The dataset provides the standard training, validation and test splits for machine learning practitioners and will be publicly available upon paper acceptance. Moreover, standard metrics for the comparison haven been defined; with the hope to speed-up the research in the endoluminal scene segmentation area.

Together with the dataset, we provided new baselines based on fully convolutional networks, which outperformed by a large margin previously published results, without any further post-processing. We extended the proposed pipeline and used it as proxy to perform polyp detection. Due to the lack of non-polyp frames in the dataset, we reformulated the task as polyp localization. Once again, we highlighted the superiority of deep learning based models over traditional hand-crafted approaches. As expected and contrary to hand-crafted segmentation methods, when considering polyp segmentation, deep learning based approaches do not suffer from specular highlights, showing the robustness of the learnt features towards saturation zones in colonoscopy images. Moreover, given that FCN not only excels in terms of performance but also allows for nearly real-time processing, it has a great potential to be included in future DSS for colonoscopy.

Knowing the potential of deep learning techniques, efforts in the medical imaging community should be devoted to gather larger labeled datasets as well as designing deep learning architectures that would be better suited to deal with colonoscopy data. This paper pretends to make a first step towards novel and more accurate DSS by making all code and data publicly available, paving the road for more researcher to contribute to the endoluminal scene segmentation domain.

Acknowledgements The authors would like to thank the developers of Theano [37] and Keras [12]. We acknowledge the support of the following agencies for research funding and computing support: Imagia Inc., Spanish government through funded project AC/DC TRA2014-57088-C2-1-R and iVENDIS (DPI2015-65286-R), SGR projects 2014-SGR-1506, 2014-SGR-1470 and 2014-SGR-135 and TECNIOspring-FP7-ACCI grant, by the FSEED and NVIDIA Corporation for the generous support in the form of different GPU hardware units.

References

1. Arganda-Carreras, I., Turaga, S.C., Berger, D.R., et al.: Crowdsourcing the creation of image segmentation algorithms for connectomics. *Frontiers in Neuroanatomy* **9**(142) (2015). DOI 10.3389/fnana.2015.00142
2. Bernal, J., Gil, D., Sánchez, C., Sánchez, F.J.: Discarding non informative regions for efficient colonoscopy image analysis. In: *International Workshop on Computer-Assisted and Robotic Endoscopy*, pp. 1–10. Springer (2014)
3. Bernal, J., Núñez, J.M., Sánchez, F.J., Vilariño, F.: Polyp Segmentation Method in Colonoscopy Videos by Means of MSA-DOVA Energy Maps Calculation. In: *MICCAI 2014 Workshop on Clinical Image-Based Procedures*, pp. 41–49. Springer (2014)
4. Bernal, J., Sánchez, F.J., et al.: WM-DOVA maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized Medical Imaging and Graphics* **43**, 99–111 (2015)
5. Bernal, J., Sánchez, J., Vilarino, F.: Towards automatic polyp detection with a polyp appearance model. *Pattern Recognition* **45**(9), 3166–3182 (2012)
6. Bernal, J., Sánchez, J., Vilarino, F.: Impact of image preprocessing methods on polyp localization in colonoscopy frames. In: *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 7350–7354. IEEE (2013)
7. Bernal, J., Tajbakhsh, N., Sanchez, F.J., Matuszewski, B.J., Chen, H., Yu, L., Angermann, Q., Romain, O., Rustad, B., Balasingham, I., Pogorelov, K., Choi, S., Debar, Q.,

- Maier-Hein, L., Speidel, S., Stoyanov, D., Brandao, P., Cordova, H., Sanchez-Montes, C., Gurudu, S.R., Fernandez-Esparrach, G., Dray, X., Liang, J., Histace, A.: Comparative validation of polyp detection methods in video colonoscopy: Results from the miccai 2015 endoscopic vision challenge. *IEEE Transactions on Medical Imaging* **PP**(99), 1–1 (2017). DOI 10.1109/TMI.2017.2664042
8. Brosch, T., Tang, L.Y.W., Yoo, Y., et al.: Deep 3d convolutional encoder networks with shortcuts for multiscale feature integration applied to multiple sclerosis lesion segmentation. *IEEE TMI* **35**(5), 1229–1239 (2016). DOI 10.1109/TMI.2016.2528821
 9. Bruno, M.: Magnification endoscopy, high resolution endoscopy, and chromoscopy; towards a better optical diagnosis. *Gut* **52**(suppl 4), iv7–iv11 (2003)
 10. Cha, K.H., Hadjiiski, L., Samala, R.K., Chan, H.P., Caoili, E.M., Cohan, R.H.: Urinary bladder segmentation in ct urography using deep-learning convolutional neural network and level sets. *Medical physics* **43**(4), 1882–1896 (2016)
 11. Chen, H., Qi, X., Cheng, J., Heng, P.A.: Deep contextual networks for neuronal structure segmentation. In: Proceedings of the 13th AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA., pp. 1167–1173 (2016)
 12. Chollet, F.: Keras. <https://github.com/fchollet/keras> (2015)
 13. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A Large-Scale Hierarchical Image Database. In: CVPR09 (2009)
 14. Drozdal, M., Vorontsov, E., Chartrand, G., Kadoury, S., Pal, C.: The importance of skip connections in biomedical image segmentation. *CoRR* **abs/1608.04117** (2016). URL <http://arxiv.org/abs/1608.04117>
 15. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. *CoRR* **abs/1411.4734** (2014). URL <http://arxiv.org/abs/1411.4734>
 16. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* **88**(2), 303–338 (2010)
 17. Gross, S., Stehle, T., et al.: A comparison of blood vessel features and local binary patterns for colorectal polyp classification. In: SPIE Medical Imaging, pp. 72,602Q–72,602Q. International Society for Optics and Photonics (2009)
 18. Havaei, M., Davy, A., Warde-Farley, D., et al.: Brain tumor segmentation with deep neural networks. *CoRR* **abs/1505.03540** (2015)
 19. Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L.: Densely connected convolutional networks. *CoRR* **abs/1608.06993** (2016). URL <http://arxiv.org/abs/1608.06993>
 20. Leufkens, A., Van Oijen, M., et al.: Factors influencing the miss rate of polyps in a back-to-back colonoscopy study. *Endoscopy* **44**(5), 470–475 (2012)
 21. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollr, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European Conference on Computer Vision (ECCV). Zrich (2014)
 22. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3431–3440 (2015)
 23. Machida, H., Sano, Y., et al.: Narrow-band imaging in the diagnosis of colorectal mucosal lesions: a pilot study. *Endoscopy* **36**(12), 1094–1098 (2004)
 24. Menze, B., Jakab, A., Bauer, S., et al.: The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). *IEEE TMI* p. 33 (2014). DOI 10.1109/TMI.2014.2377694
 25. Núñez, J.M., Bernal, J., Ferrer, M., Vilariño, F.: Impact of keypoint detection on graph-based characterization of blood vessels in colonoscopy videos. In: International Workshop on Computer-Assisted and Robotic Endoscopy, pp. 22–33. Springer (2014)
 26. Park, S.Y., Sargent, D.: Colonoscopic polyp detection using convolutional neural networks. In: SPIE Medical Imaging, pp. 978,528–978,528. International Society for Optics and Photonics (2016)
 27. Prastawa, M., Bullitt, E., Ho, S., Gerig, G.: A brain tumor segmentation framework based on outlier detection. *Medical image analysis* **8**(3), 275–283 (2004)
 28. Ribeiro, E., Uhl, A., Häfner, M.: Colonic polyp classification with convolutional neural networks. In: 2016 IEEE 29th International Symposium on Computer-Based Medical Systems (CBMS), pp. 253–258. IEEE (2016)

29. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. In: In Proceedings of ICLR (2015)
30. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 234–241. Springer (2015)
31. Roy, H.K., Goldberg, M.J., Bajaj, S., Backman, V.: Colonoscopic optical biopsy: Bridging technological advances to clinical practice. *Gastroenterology* **140**(7), 1863 (2011)
32. Silva, J., Histace, A., Romain, O., Dray, X., Granado, B.: Toward embedded detection of polyps in wce images for early diagnosis of colorectal cancer. *International Journal of Computer Assisted Radiology and Surgery* **9**(2), 283–293 (2014). DOI 10.1007/s11548-013-0926-3. URL <http://dx.doi.org/10.1007/s11548-013-0926-3>
33. Society, A.C.: Colorectal cancer. online (2016)
34. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014)
35. Styner, M., Lee, J., Chin, B., et al.: 3d segmentation in the clinic: A grand challenge ii: Ms lesion segmentation (2008)
36. Tajbakhsh, N., Gurudu, S., Liang, J.: Automated polyp detection in colonoscopy videos using shape and context information. *Medical Imaging, IEEE Transactions on* **PP**(99), 1–1 (2015). DOI 10.1109/TMI.2015.2487997
37. Theano Development Team: Theano: A Python framework for fast computation of mathematical expressions. arXiv e-prints **abs/1605.02688** (2016). URL <http://arxiv.org/abs/1605.02688>
38. Tieleman, T., Hinton, G.: rmsprop adaptive learning. In: COURSERA: Neural Networks for Machine Learning (2012)